

LLIN Visualizer

Visualizing Distributions of Long-Lasting Insecticide-Treated Mosquito Nets in
the Epidemiological Context

Denis Drescher

denis.drescher@fu-berlin.de

4280818

August 18, 2016

Freie Universität Berlin
Department of Computer Science
Databases and Information Systems

Master's Thesis

Supervisors

Prof. Dr. Agnès Voisard

Dr. Kashif Rasul

Mentor

Dr. Kashif Rasul

Executive Summary

The LLIN Visualizer makes transparent many of the factors that guide the prioritization decisions of the Against Malaria Foundation. Thereby it fosters trust between the foundation – one of the funders of distributions of long-lasting insecticide-treated bednets for malaria vector control – and altruistic investors who donate toward the intervention, and thus addresses a coordination problem that poses a hurdle toward a more efficient market for altruistic impact.

Statement of Authorship

I hereby attest that I have composed this thesis with no other help than the indicated references and resources. Furthermore, I attest that this thesis has not been submitted for any other degree or professional qualification.

Denis Drescher

Contents

1	Introduction	8
2	Previous Work	12
2.1	Previous Research on Altruistic Coordination	12
2.1.1	Willingness to Pay for Increased Moral Goal Satisfaction	12
2.1.2	Diversity of Moral Preferences	14
2.1.3	Identifying Moral Preferences	15
2.1.4	Values Spreading	17
2.1.5	Trust	21
2.2	Previous Research on Altruistic Prioritization	21
2.2.1	Finding the Best Charities	22
2.2.2	Motivations of Donors	24
2.2.3	A Market for Impact	27
2.3	Prioritization Research	28
2.3.1	Constraints	28
2.3.2	Methodologies	30
2.4	Strategy of the Against Malaria Foundation	35
2.5	Distributions of Long-Lasting Insecticide-Treated Mosquito Nets	39
2.5.1	Key Epidemiological Facts	39
2.5.2	Brief History	40
2.5.3	Role of the Against Malaria Foundation	42
3	Requirements and Modeling	44
3.1	Forming the Idea	44
3.2	Data Exploration	47
3.2.1	The Survey Form	48
3.2.2	The Survey Data	49
3.2.3	Forthcoming Data	50
3.2.4	Other Data	50
3.2.5	Insecticide Resistance	51
3.2.6	Incidence and Endemicity	51
3.2.7	Net Gaps	51
3.2.8	Indicators	52

3.3	Functional Requirements	54
3.3.1	The Original Set	54
3.3.2	The Extension Set	54
3.4	Coordination with AMF	55
4	Architecture and Implementation	57
4.1	Architecture	57
4.1.1	Models	57
4.1.2	Controllers	60
4.1.3	Views	60
4.1.4	Management Commands	60
4.2	Technology	64
4.2.1	Core Technologies	64
4.2.2	Other Important Technologies	66
5	Evaluation and Learnings	69
5.1	Functional Optimization	69
5.1.1	Evaluation	69
5.2	Nonfunctional Optimization	71
5.2.1	Evaluation	73
5.3	Learnings	74
5.3.1	Dates	74
5.3.2	Health Areas	75
6	Appendices	77
6.1	IMA 2014 Predistribution Registration Survey	77
6.2	Concept for Donor Coordination	83
6.2.1	Donor Coordination	83
6.2.2	High-Level Goals	86
6.2.3	Core Concepts	86
6.2.4	Interest Groups	87
6.2.5	Functional Requirements	88
6.2.6	Challenges and Proposed Solutions	88
6.2.7	Remaining Challenges	90
	Bibliography	92

Acknowledgements

I would like to express my gratitude to Rob Mather and Dr. Andrew Garner of the Against Malaria Foundation, who supplied me with much data and answered all my questions about net distributions, and my mentors and supervisors for their feedback and encouragement. I would also like to thank all my beta testers, Sara Nowak for her invaluable quality assurance, Martin Wilson for English copyediting, and Djahbaï Chonville for French copyediting. Finally I want to thank all that have shaped my ethical and epistemic worldview, without whom I would not have had the idea of writing this thesis.

Source Code and Demonstration

Source code. The complete LLIN Visualizer project is open source and hosted on Bitbucket.¹ At the time of writing, the latest commit was fbdbf2e (July 14, 2016).

LLIN Visualizer demonstration instance. I am maintaining a demonstration instance of the project on my server at <http://nets.claviger.net/>.²

LLIN Visualizer demonstration video. I have also produced a video demonstrating the use of the LLIN Visualizer.³

-
1. Drescher, "Telofy / LLIN Visualizer – Bitbucket."
 2. Drescher, "Map – LLIN Visualizer."
 3. Drescher, "LLIN Visualizer Demo."

“Homines ad deos nulla re propius accedunt quam salutem hominibus dando.” (In nothing do humans more nearly approach the gods than in giving health to humans.)

Cicero, Pro Ligario, XII

1

Introduction

Just as there is a market for monetary profit – which is an agent-relative goal since one person’s profit is not automatically another person’s profit – there could also be a market for altruistic impact – for maximizing happiness, minimizing suffering, or any other such agent-neutral goals that people find worthy of optimization. An organization that manages to make efficient progress toward such a goal would receive lavish investment until it reaches its limits of scalability and its marginal cost-effectiveness drops below that of another organization. Startups in this space would receive such lavish investment if they promise rapid growth that will render them highly cost-effective in the long run.

Unfortunately there are a few hurdles left to overcome toward such a market. Some agents are not or not yet intent on investing in such a market; failure modes of this category will be addressed briefly below. They have been subject to research in social psychology and behavioral economics that is outside the purview of this text. But even agents that are eager to drive positive change in the world can face problems. Here, two important ones:

1. When investors and recipients are two distinct groups, the investors have an even harder time assessing the impact they have had than they would have had anyway, so that, in the absence of rigorous impact evaluations, vast inefficiencies result.
2. The investments that are most cost-effective in expectation require highly risk-neutral investors, but the learned heuristics for the diminishing marginal utility of money for personal investments skew expectations for altruistic investments.

Only the first of these is subject to the LLIN Visualizer discussed in this text. The second is addressed in the appendix 6.2. So let us consider some of the inefficiencies – or injustices – that these and other problems allow to continue unabated.

The organization Giving What We Can has compiled data from three sources to get an idea of the global income distribution.¹ According to those data, a person in Germany with an income of €42,000 is already within the richest 1% of the world population.

This inequality leads to large numbers of poverty-related deaths,² a rough lower bound of which can be estimated at over 30,000 per day.³

If just a fraction of the richest in the world redistributed a portion of their wealth to the poorest – directly and by investing into research how to alleviate poverty most effectively as well as scaling up the solutions – this death toll could decrease drastically.

And in fact we are seeing significant flows of international aid and philanthropy. In 2013, aid and philanthropy totaled \$197 billion,⁴ while the Brookings Institution estimates that only \$66 billion per year can suffice to eliminate extreme poverty globally.⁵

MacAskill also estimates that in view of purchasing power parity, average income, and diminishing marginal returns to wealth, “for those of us living in rich countries, you should expect to be able to do at least one hundred times as much to benefit other people as you can to benefit yourself.”⁶ Yet, using the example of US philanthropy in 2014, only 4% of donations were invested outside the US.⁷ These figures are similar for other developed countries and would indicate that citizens of a country value their compatriots about $\frac{100\%}{4\%} \cdot 100 = 2500$ times as highly as a person in extreme poverty. Still, the guiding mantra of the Bill & Melinda Gates Foundation – that every life has equal value – does not commonly stir up controversy.

There are many reasons why we are seeing this odd allocation of resources, and most of them are well without the purview of this text, such as the psychology behind surprising investment decisions that seem to indicate that some people value a week in space higher




Figure 1.1: Global income distribution.

1. “The figures for between the richest 1% and the richest 21% are based on micro data from national household surveys carried out in 2008, kindly provided by Branko Milanovic, author of *The Haves and the Have-Nots*. The figures for the poorest 73% are based on the 2008 data from PovcalNet, adjusted based on the approximation that the surveys covered unbiased samples of the poorest 80% of the world’s population. The figure of \$70,000 for the top 0.1% is from Milanovic’s book. All figures have been adjusted for the Consumer Price Index measure of inflation.” Giving What We Can, “How Rich Am I?”

2. MacAskill, *Doing Good Better: How Effective Altruism Can Help You Make a Difference*, 431.

3. The author adds up the death tolls of the leading poverty-related causes of death, which I divided by 365: $\frac{12.7 \text{ million}}{365} \approx 34795$.

4. Adelman et al., “The Index of Global Philanthropy and Remittances, with a Special Report on Emerging Economies,” 6, 12.

5. Chandy and Gertz, *Poverty in Numbers*, 13.

6. MacAskill, *Doing Good Better: How Effective Altruism Can Help You Make a Difference*, 55.

7. Giving USA Foundation, *Giving USA 2015*.

than saving the lives of approximately 9,600 people,⁸ or that people value saving one dog over saving approximately 150 dogs or almost 290,000 other animals,⁹ to cite just two examples.¹⁰

While the specific reasons are manifold, they can be roughly categorized into three categories:

Moral differences. The following chapters will give a brief insight into the research into the different moral preferences people have, as well as the debate over whether these are differences of taste or of belief. Usually, these differences are fairly subtle, but that is not to say that there are not also highly outré moral preferences.¹¹

Lacking education. Educating oneself comes at a price that is usually paid in time.¹² Hence it is rational for agents to forgo education if they expect the value of information of the activity to be lower than the price they would have to pay for it. However, such an estimate of the expected value of information may be grossly inaccurate, in which case agents may make decisions that are contrary to their moral preferences simply out of ignorance. Apart from lacking information, cognitive biases such as scope neglect,¹³ the identified victim effect,¹⁴ and the bystander effect,¹⁵ are often involved.

Failures of coordination. Even when groups of agents are value aligned and well informed, they can fail to minimize the overall total of the opportunity costs of their investment decisions if they lack critical ingredients of coordination, such as communication and trust.

8. BBC News, “Profile: Tito the Spaceman”; GiveWell, “Against Malaria Foundation (AMF).”

9. The Week, “The Little Dog Lost at Sea”; Animal Charity Evaluators, “Impact of Donations.”

10. The estimate are based on GiveWell’s and Animal Charity Evaluators’ cost-effectiveness estimates of bednet distributions, animal shelters, and farmed animal advocacy programs respectively. The costs are inflation-adjusted according to the Consumer Price Index.

11. An example from a personal correspondence: “I’d happily condemn a hundred thousand strangers in a developing country to death in exchange for just one person I know to get better, simply because I couldn’t care less about the life of a complete stranger – I don’t wish him any harm, but I am hard pressed to find a reason why I should help him. He doesn’t affect my life in the slightest; he’s just void.”

12. Time as well as money in the case of formal education in many countries. Access to information can also cost money in cases where it requires computer hardware and Internet access.

13. Kahneman et al., “Economic Preferences or Attitude Expressions?,” 212; Desvousges et al., *Measuring Nonuse Damages Using Contingent Valuation*, 66.

14. Erlandsson, Björklund, and Bäckström, “Emotional Reactions, Perceived Impact and Perceived Responsibility Mediate the Identifiable Victim Effect, Proportion Dominance Effect and In-Group Effect Respectively”; Kogut and Ritov, “The ‘Identified Victim’ Effect”; Jenni and Loewenstein, “Explaining the Identifiable Victim Effect”; Small and Loewenstein, “Helping a Victim or Helping the Victim.”

15. Hudson and Bruckman, “The Bystander Effect”; Darley and Latane, “Bystander Intervention in Emergencies.”

The contributions of this project only address agents that feel a sufficient moral alignment with the goals of the Against Malaria Foundation (AMF), and the problem of lacking education is being addressed by several organizations of the effective altruism spectrum and hence left outside the scope of this text. Rather, the contributions of this project aim to improve trust into the prioritization decisions of AMF through enhanced transparency, a crucial ingredient of the coordination of altruistic efforts.

In the following sections of this chapter I will discuss these hurdles in more detail in order to locate the contributions of this project within their context, introduce methods of prioritization research, and then focus on the specific prioritization decisions relevant to the intervention that the LLIN Visualizer addresses: the distribution of long-lasting insecticide-treated mosquito nets for malaria vector control. On that foundation I will describe the process of conception and implementation of the LLIN Visualizer: previous work that it builds upon, the requirements it was conceived of to satisfy, the considerations that have gone into its architecture and implementation, and finally the evaluation of the results and important learnings drawn from the process.

Aside 1: Identified Victim Effect

The *identified victim effect* is closely associated with the *proportion dominance effect* and *in-group effect*,^a and thus makes for a rich example of the psychological phenomena that will not be covered any further. Three factors that influence people's readiness to help are their sympathy with the victim, the perceived utility of their help, and their feeling of responsibility.

Sympathy can be aroused by giving the name and possibly a picture of the victim rather than citing a statistic – even though these donors most likely know that each person in such a statistic is also an individual with name and face.

The perceived utility is partially dependent on the reference class, so that people are more excited about helping 20 people out of a group of 24 rather than 20 out of 400 – even though the value of each life does not change as the reference class changes, and reference classes are often chosen arbitrarily.

Feelings of responsibility are stronger the closer or more alike the victim is. Erlandsson, Björklund, and Bäckström, for example, vary the degree of the family relationship but other “in-groups,” such as nationality, ethnic group, or species, are imaginable. This comes at a detriment to those who lack any in-group members with the capacity to help them, for example many of the poor in developing countries.

^a Erlandsson, Björklund, and Bäckström, “Emotional Reactions, Perceived Impact and Perceived Responsibility Mediate the Identifiable Victim Effect, Proportion Dominance Effect and In-Group Effect Respectively.”

"I could have got more out. I could have got more. I don't know. If I'd just... I could have got more."

Oskar Schindler, played by Liam Neeson in *Schindler's List*, 1993

2

Previous Work

2.1 Previous Research on Altruistic Coordination

2.1.1 Willingness to Pay for Increased Moral Goal Satisfaction

Moral preferences. A basic metric in economic analysis is the willingness to pay (WTP) of an agent. Agents have different preferences, and as these preferences are met, the marginal benefit of further preference satisfaction along the same dimension decreases, which is termed the *principle of decreasing marginal benefit*: As the marginal benefit of a good to an agent decreases, so does the agent's WTP for that good.¹

Usually agents consider some of these preferences *moral preferences* – some typical examples are a preference for feeling beings to suffer as little as possible or a preference for people not to lie. I will leave it to the intuition of the reader to decide what should constitute a *moral* preference as opposed to any other preference, as any definition I have seen leaves a significant gray area and would not be necessary for comprehension of this text.

Rather I will distinguish goals according to the degree to which they are shared by agents. Here agent-relative goals (such as personal profit, neighborly help, patriotism, etc.) have an inherent upper bound on the extent to which other agents can share them, while agent-neutral goals (reducing suffering, maximizing justice, maximizing the number of paperclips in the universe, etc.) do not have such a limit.

To aid comprehension, however, I will make an artificial distinction of moral *preferences* and moral *goals* that becomes meaningful in the case of agent-relative preferences: two people with a personal profit motive share the same preference for profit but their goals

1. Parkin, Powell, and Matthews, *Economics*, 40.

are different ones since they are different agents. If they also share the agent-neutral preference for minimizing global suffering, then they also share the same goal of reducing it.

At least agent-neutral goals are pure public goods in that they are *nonrival* and *nonexcludable*: One agent's enjoyment of achievement of a moral goal does not lessen another agent's enjoyment, and it is impossible (or prohibitively costly) to bar someone from enjoying it.²

Analogously, I will not distinguish donations from purchases or investments, except insofar as I view them all as different types of investments.

Measuring the willingness to pay. In some contexts it is hard to tease apart the influence of agent-neutral and agent-relative motives on someone's revealed preference. People may be concerned about climate change because of the potential harm it can cause for all sentient life or because they own property in coastal regions.

One area where researchers have found this to be more feasible is in farmed animal welfare. Norwood and Lusk determined the revealed WTP of a fairly representative sample of informed consumers for (1) pork and eggs from different systems of animal farming³ and (2) changes of systems of animal farming for animals whose produce the subjects could not consume.⁴ Particularly the second part is of interest here, as such welfare increases constitute a pure public good.

The authors could not simply study point of sales data and infer people's preferences and WTP from these because most people are uninformed of practices in animal agriculture. This parallels a wider problem in altruistic investment that will be the subject of the following subsections.

Norwood and Lusk thoroughly informed a sample of consumers about advantages and disadvantages for animal welfare of production systems such as battery cages and the free range system for hens, took measures to minimize social incentives, and then auctioned animal welfare changes affecting set numbers of animals at university farms far from the research locations. The study relied on actual payments rather than hypothetical questions, which usually yield exaggerated valuations. Thus the influence of environmental concerns associated with animal agriculture is minimized since the number of animals in question is unaffected, and selfish interests, such as health benefits, are minimized since the subjects are highly unlikely to ever get in touch with the farms' produce.

Nonetheless, two thirds of the participants were willing to pay for a single hen to be moved from the cage to the free range system for an average (including the one third of bids of zero) of \$1.08.⁵

Results like these show that moral goals, absent selfish interests, have real monetary value to a majority of people.

2. Robert E. Hall, *Microeconomics: Principles and Applications*, 477.

3. Norwood and Lusk, *Compassion, by the Pound*, 266.

4. *Ibid.*, 294.

5. *Ibid.*, 298.

2.1.2 Diversity of Moral Preferences

As will be detailed in section 2.2.2, the benefits from seemingly altruistic actions like donating are for many people either dominated by other benefits than moral preference satisfaction or their moral preferences are hard to distinguish from the self-interested preferences of other people. For other people their moral preferences are paramount. In either case moral preferences are of different nature and importance for different people.

The reasons for this disagreement are subject of the philosophical discipline of metaethics. The two metaethical views that are typically distinguished are moral realism and moral antirealism. According to the first doctrine, there are objective moral facts while according to the second morality is subjective.

Singer argues that this distinction, while fundamental, is yet of little practical relevance. While moral realists would attribute the diversity of moral views to differences in beliefs about the world, antirealists attribute it to differences about personal preferences. Consequently, realists have failed to establish how any set of moral goals follows from circumstances of the observable world while antirealists, by nature of their claim, fail to establish why any set of moral goals should be binding for anyone.⁶ Since neither view has been substantiated by evidence, there is instead a considerable debate over which view bears the burden of proof in the first place.

As these doctrines use different language to describe the same facts, I am forced to choose lest I need to present every argument from two perspectives. My weakly held perception is that antirealism is the more parsimonious theory, so that I will use the language of antirealism throughout this paper. I want to be clear, however, that this choice is motivated not by evidence for the view but by an absence of evidence for the competing view. Either view of course acknowledges that differences in revealed moral preferences are real.

Often these differences are due to insufficient information or biases. A person may, for example, be willing to pay significant amounts for the building of schools in a region, but may then discover that the availability of schools is not the bottleneck to achieving better education in that region. Plausibly, for example, the children's health or the parents' and teachers' beliefs about the benefits of education may be bottlenecks.⁷ Now the person may be willing to pay much less for more schools and much more for improved health and improved access to information in the same region. The preference for more schools was not an informed preference and thus not indicative of a *terminal* moral goal.

Still somewhat unpragmatically, a terminal moral goal can be defined as any moral goal that a person would still hold even if they were arbitrarily intelligent, perfectly rational, and had available to them all relevant information.⁸ In the following I will use *moral goal* (or *moral preference*) to mean *terminal moral goal* unless stated otherwise.

6. Singer, "The Triviality of the Debate over 'Is-Ought' and the Definition of 'Moral'."

7. Banerjee and Duflo, *Poor Economics*, 79.

8. These thoughts are inspired by Bostrom's orthogonality thesis. Bostrom, "The Superintelligent Will"

2.1.3 Identifying Moral Preferences

As in the case of animal agriculture, we again face the problem that we cannot simply study donation statistics to infer people's moral preferences since most of these revealed preferences are likely far from being terminal moral preferences.

The Money For Good II study provides some quantification of the problem. It found that only 6% of donations from individual donors (as opposed to advisors and foundations) are informed by any comparative research, and 33% are informed by any research, with the additional 17 pp. of donors only doing research on the one charity they seek to donate to.⁹ Without doing comparative research, these donors are not informed of whether there are charities that realize their moral goals much more directly and efficiently, and regarding the motivations that Money For Good II found for donations, we will also see that many are unrelated to moral goals.

One limited means for approximating terminal moral goals are surveys of philosophers such as that of The PhilPapers Foundation, following the assumption that philosophers are likely to have had to question their moral intuitions for some time. The survey reveals that the three most popular classes of moral theories, deontology (25.9%), consequentialism (23.6%), and virtue ethics (18.2%), share the space with no single one clearly dominating the others.¹⁰ These do not allow inferences to the popularity of individual moral goals, but they do make implications about the relative popularity of classes of moral goals at least within the surveyed cohort.

One caveat is that this survey appears to have a fairly informal character, which makes it hard to assess its quality. Another caveat is that philosophers of domains other than the branches of ethics may have no informed opinion on moral questions but that limiting the sample to ethicists greatly reduces the sample size. Finally there is some evidence to contest the assumption that education about ethical questions reduces a person's susceptibility to bias.¹¹

Another interesting branch of research is that associated with Moral Foundations Theory.¹² Researchers in this space try to find moral foundations – which may be analogous to what I refer to as moral goals – with high predictive power for explaining cultural, often political, differences. The researchers that proposed the theory proceeded from a set of first five, later six, and in the future possibly more moral principles. The original set of five principles derived from early work by Shweder¹³ and has since been expanded. More recent papers tested the principles' predictive power through self-report surveys, measures of reaction time, facial micro-expressions, event-related potentials, neuroimaging, and text analysis.¹⁴ In all cases that I am aware of, the set of principles was predetermined, and in the case of the surveys the researchers had developed questions that aimed to iso-

9. Hope Consulting, "Money For Good II," 14.

10. The PhilPapers Foundation, "Preliminary Survey Results."

11. Schwitzgebel and Cushman, "Expertise in Moral Reasoning?"

12. Graham et al., "Moral Foundations Theory."

13. *Ibid.*, 16.

14. *Ibid.*, 15.

late single principles in a hypothetical situation that was simple enough not to require further education of the subject.

Here is the latest set of foundational principles according to the researchers' website:¹⁵

“ Each culture then constructs virtues, narratives, and institutions on top of these foundations, thereby creating the unique moralities we see around the world, and conflicting within nations too. The five foundations [plus one “very good candidate”] for which we think the evidence is best are:

Care/harm This foundation is related to our long evolution as mammals with attachment systems and an ability to feel (and dislike) the pain of others. It underlies virtues of kindness, gentleness, and nurturance.

Fairness/cheating This foundation is related to the evolutionary process of reciprocal altruism. It generates ideas of justice, rights, and autonomy. Note: In our original conception, Fairness included concerns about equality, which are more strongly endorsed by political liberals. However, as we reformulated the theory in 2011 based on new data, we emphasize proportionality, which is endorsed by everyone, but is more strongly endorsed by conservatives.

Loyalty/betrayal This foundation is related to our long history as tribal creatures able to form shifting coalitions. It underlies virtues of patriotism and self-sacrifice for the group. It is active anytime people feel that it's “one for all, and all for one.”

Authority/subversion This foundation was shaped by our long primate history of hierarchical social interactions. It underlies virtues of leadership and followership, including deference to legitimate authority and respect for traditions.

Sanctity/degradation This foundation was shaped by the psychology of disgust and contamination. It underlies religious notions of striving to live in an elevated, less carnal, more noble way. It underlies the widespread idea that the body is a temple which can be desecrated by immoral activities and contaminants (an idea not unique to religious traditions).

Liberty/oppression This foundation is about the feelings of reactance and resentment people feel toward those who dominate them and restrict their liberty. Its intuitions are often in tension with those of the authority foundation. The hatred of bullies and dominators motivates people to come together, in solidarity, to oppose or take down the oppressor. We report some preliminary work on this potential foundation in this paper, on the psychology of libertarianism and liberty.

15. MoralFoundations.org, “Moral Foundations Theory.”

While the hypothetical scenarios used in at least one study did seem sufficiently simple to me to readily comprehend their implications, the isolation of the principles was less successful in several cases. For example, one question that tests for readiness to subvert authority asks what monetary incentive would be sufficient for the respondents to “curse [their] parents, to their face,” answers to which may just as soon be indicative of the respondents’ readiness or reluctance to inflict harm.

The theory has been criticized for being less parsimonious than established taxonomies that rely only on harm¹⁶ or only on authority.¹⁷ Conversely I could not find any mention of whether principles such as the proliferation of knowledge or beauty can be reduced to the known principles or whether they might represent principles of their own right. The question is suggested by the significant levels of philanthropy to support universities and the arts. Perhaps the research is forthcoming.

2.1.4 Values Spreading

Agents interested in one or several of the same moral goals will soon notice that they can reach these more easily when they cooperate on realizing them, convince more agents to join their cause, avoid zero-sum competition, and possibly even trade with groups with different goals. I have put these modes of values spreading into the following order so that the cheapest method (requiring minimal resources and sacrifice) has first priority, while the rest require increasing effort and compromising on one’s goals.

1. **Cooperation.** Agents can agree on some of their (not necessarily terminal) moral goals and cooperate toward their realization.
2. **Education.** Education comes in two subforms, both of which can generally be thought of as cooperative behaviors.
 - a) **Correcting ignorance.** Educating other agents on topics they had not considered or did not know enough about.
 - b) **Correcting misinformation.** Helping other agents realize that some of their assumptions about the world are mistaken.
3. **Trade.** When it has either been established that both sides disagree on preferences rather than beliefs or when both sides find that the other is resistant to education to the point that it would make it prohibitively costly, then both sides can usually gain from avoiding a zero-sum fight in favor of compromise.

These are explained in more detail below:

16. Gray, Young, and Waytz, “Mind Perception Is the Essence of Morality.”

17. Kugler, Jost, and Noorbaloochi, “Another Look at Moral Foundations Theory.”

Cooperation. When agents have several moral goals each, the situation resembles a stag hunt in that there are two Nash equilibria – when the agents cooperate and when they defect. (In the prisoner’s dilemma, for comparison, the only Nash equilibrium is the case of both agents defecting.)

To cooperate, they will need to agree on pursuing the moral goal that they share. A concrete example may be the following: One activist is interested in both the proliferation of knowledge and the reduction of suffering. Another activist is focused solely on the reduction of suffering. Through dialogue, they can come to agree on the shared goal of reducing suffering. A variation is the situation where at least one agent shares only an instrumental goal with the other agent. Thus, for example, a queer activist and an anti-HIV/AIDS activist can both agree to support secularization efforts in Uganda. Of course, the agents can enter into several such cooperative efforts if their working relationships and resources allow.

Education. Education comes in two different forms depending on whether the recipient is mistaken or merely oblivious.

Correcting ignorance. When people do not seem to share any moral goals, it may be the case that they simply never considered some key question. They may not be aware that there is a decision that they have constantly been making by omission because they were unaware of it. Telling such people about the options they have may greatly enrich their thinking and their lives. It may also come at a benefit to the internal consistency of their moral system or their identity and thus reduce cognitive dissonance. Some people may never have considered philanthropy, for example, so simply by telling them about this possibility one can create new cooperation partners.

Aside 2: Game Theory Terms

Stag hunt. The stag hunt is a scenario where at least two agents have the option to cooperate on a common goal that would be optimal for both of them or defect in favor of some smaller gain, thereby also ruining the chances of the other to attain more than the smaller gain. The payoff matrix of all options (cells) for both players (tuple elements):

$$\begin{array}{cc} 2, 2 & 0, 1 \\ 1, 0 & 1, 1 \end{array}$$

Prisoner’s dilemma. The prisoner’s dilemma is different from the stag hunt in that the cases in which exactly one agent defects are optimal for the respective agent. Since the outcome of both defecting is slightly better than cooperating while the other defects, there is a double incentive to defect. The objectively best outcome, however, is that of both agents cooperating. The payoff matrix of all options (cells) for both players (tuple elements):

$$\begin{array}{cc} 3, 3 & 0, 5 \\ 5, 0 & 1, 1 \end{array}$$

Nash equilibrium. In our examples, the Nash equilibria are the scenarios in which neither agent has anything to gain by changing their own strategy when they learn the strategy of the other agent.

Correcting misinformation. One variation on this theme is the case where one party has badly chosen instrumental goals but compatible terminal goals, whether they are aware of them (thinking their instrumental goals are well chosen) or not (mistaking their instrumental goals for terminal goals). In such cases, one cost-effective and cooperative way of winning them for the cooperative effort would be to provide them with the information they are lacking.

For example, a person might think that “charity begins at home,” and that when everyone supports their neighbors, friends, and family, everyone is supported and even efficiently so, because we know how to effectively support those close to us. But then they may learn that in some countries poverty is so rampant, and the rich are so well insulated from the poor, that any neighbors, friends, and family of the person in need are just as much in need themselves. They may learn that there are neglected interventions whose success is much more certain and cost-effective than that of any interventions addressing common afflictions in our society. Realizing this, the person may update and abandon the mantra.

Trade. Other people, however, may have the same convictions as the misinformed person, but may not actually lack information. They may be well aware of all one’s arguments and still not be swayed by them. Here the education approach will fail. What remains is trade.¹⁸

When these people have a strong interest in spreading their values, we encounter a situation that resembles the prisoner’s dilemma. Much research has been invested into this scenario. Axelrod, for example, has boiled down his insights from a competition of algorithms in the repeated prisoner’s dilemma into four rules:¹⁹

“

1. Don’t be envious.
2. Don’t be the first to defect.
3. Reciprocate both cooperation and defection.
4. Don’t be too clever.

”

One very simple algorithm that implemented this behavior was *Tit for Tat*, which cooperated unless the other agent chose to defect in the last round. Thus it quickly forgave defection, cooperated by default, punished and rewarded defection and cooperation respectively, and was perfectly predictable.

18. We should note, however, that the distinction between ethical and epistemic differences can be a fuzzy one. When someone’s (epistemic) beliefs are very different from one’s own, it can be very costly to convince the other person all the way to one’s own view even if one is objectively right, so that compromise can again become more efficient.

19. Axelrod, *The Evolution of Cooperation: Revised Edition*.

Tomasik has investigated what this could mean for altruists. He compares the case of (1) deep ecologists, which make a distinction between humans and other animals and care terminally about diversity of life and noninterference, and (2) animal welfarists, which care about the well-being of all individual sentient beings.²⁰ These two are often opposed. Humans have been able to reduce the suffering of humans and domesticated animals due to disease, starvation, adverse climate, and many other sources by developing technical, medical, and economic remedies. Other animals are still exposed to all these to maximal degrees and cannot change these conditions for themselves, centrally due to their lower intelligence. Hence, improving animal welfare will require intervention in ecosystems, presumably by humans directly or proximately.

When these opposed factions both want to control the future to implement their moral preferences, they would have to fight each other, a fight they would win with some probabilities less or equal to one that sum up to one in this simplified example.

However, both sides know aspects of the future that they care more strongly about than about others, while the opposing faction does not necessarily care more about the same aspects. Animal welfarists, for example, may value a world with few animals suffering greatly over a world with many animals suffering, but deep ecologists may be indifferent to the number of animals per species so long as diversity is maintained.

Based on their estimates of how likely each faction would be to overpower the other in a fight and the value they put on different aspects of their utopias, they can arrive at a compromise that is better with certainty for both than the expected value of a fight.

This is already happening. One example is the cooperation between VEBU (“Vegetarian Union”) and the sausage producer Rügenwalder on creating vegetarian products. While both sides had to make concessions, Rügenwalder was able to achieve financial success and VEBU reduced the degree of animal suffering experienced by the company’s livestock.²¹

When the factions are less homogeneous than a company or association, then there is the additional risk of subfactions forming and obviating the conditions of the original compromise. Such considerations would complicate the situation further than necessary for this overview since most of the users of the LLIN Visualizer will likely be individuals or tight-knit groups.

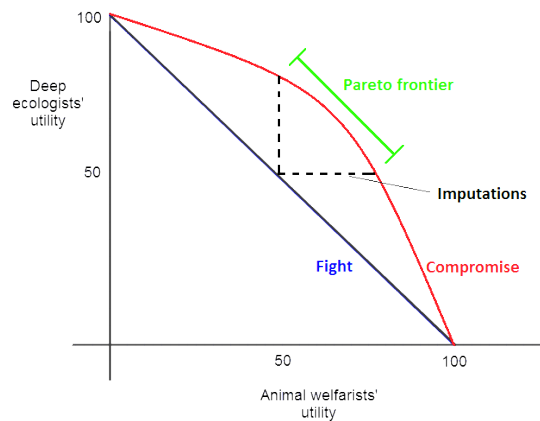


Figure 2.1: Visualization of mutual gains from trade by Tomasik.

20. Tomasik, “Gains from Trade through Compromise.”

21. Dierig, “Warum die Deutschen so gerne Veggie-Wurst essen”; Stragies, “Wursthersteller erstmals fleischfrei.”

2.1.5 Trust

The practicalities of a proliferation of this approach and related forms of trade based on moral differences have been explored in Ord, “Moral Trade.” One central problem that is discussed in the paper is the problem of trust, which takes two forms, factual and counterfactual trust.

To use an example of particular relevancy to the LLIN Visualizer: there are factories that produce WHOPEs II–approved nets as the Against Malaria Foundation (AMF) uses²² in Vietnam, Thailand, China, and Tanzania.²³ These factories are run by for-profit companies.

When altruists want to deploy a number of nets in the fight against malaria, they pay AMF to do so. In practice they pay AMF to pay some of these factories to deliver nets, and AMF then cooperates with a local partner in their distribution. But for simplicity let us view AMF as a black box: the donor pays AMF to deploy nets because AMF is better at it than the donor.

Factual trust now means that the donor can actually trust AMF to deploy the right number of high-quality nets. AMF can prove that by providing extensive documentation in the form of photos, videos, statistics, and reports from independent supervisors.

Counterfactual trust is more intricate; it means proving that the same effect would not have occurred anyway, without the donors’ contributions. This is the central reason randomized controlled trials use randomized control groups, namely to know what would have happened without the intervention. In the case of AMF, for example, one of my worries was that a distribution by the foundation may only replace another distribution by another funder who instead invests into a less cost-effective intervention. As we will see in section 2.4, the intervention is too neglected for that to be a major concern.

Trust is central to the LLIN Visualizer as it gives an overview over the prioritization decisions AMF has made as well as their results as evinced by successfully completed distributions. The specifics of how this is achieved are the matter of chapter 4.



Figure 2.2: “A woman poses for a picture in front of her net.” Photo and caption courtesy of AMF.

2.2 Previous Research on Altruistic Prioritization

Let us assume that the problem of value alignment is solved. For example, one may be allocating only one’s own donation budget or that of a foundation with a value-aligned team, or one might be deciding on one’s own career path. In these cases the next hurdle

22. Against Malaria Foundation, “About Mosquito Nets.”

23. Against Malaria Foundation, “Question to AMF: ‘Where are Your Nets Manufactured?’”

to take is that of deciding on the most viable path to take to realize one's moral goals.

Here, first, three very unlike entities require an introduction, GasBuddy, disability-adjusted life years (DALYs), and the *Disease Control Priorities in Developing Countries* study.

GasBuddy is a website that tracks gasoline prices across the US and Canada. At the time of writing, the lowest average price across all states of the US was \$1.796 per gallon and the highest \$2.802 per gallon ($\sigma = 0.226$). Admittedly, some extreme prices might have gotten normalized by the averaging within each state, but even the range of prices across all cities the website tracks is just slightly greater ($\sigma = 0.279$). The relevance of this observation will become apparent later.

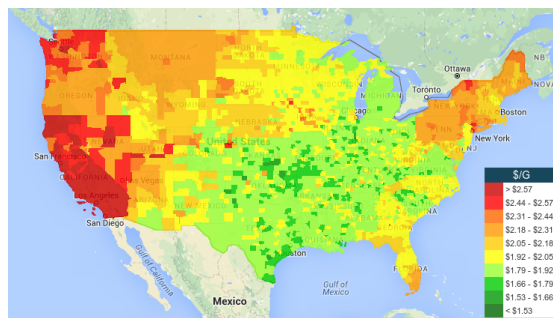


Figure 2.3: Gas prices across the US according to GasBuddy.

For many people it is intuitive to care about their own well-being and that of others. This is evident, for example, in the wide acceptance of government-funded health services in many countries. Health heavily impacts well-being, as does freedom, intellectual edification, the social environment, relationships, and many more factors. Researchers in medicine often employ quality-adjusted life years (QALYs) and disability-adjusted life years (DALYs) as measures for health and thus as proxies, however imperfect, of well-being. DALYs are the sum of years lived with disability (YLD) and years of life lost (YLL). Sassi gives a more comprehensive introduction to the two measures, including their historical background, and contrasts them.²⁴

Disease Control Priorities in Developing Countries, second edition, (DCP2) uses DALY measures for the study of the cost-effectiveness of 309 medical interventions that are typically seen as highly cost-effective by developed-world medical standards.²⁵ It is a four-year-long collaboration of over 350 specialists from around the world, made possible by funding and staff time from the Bill & Melinda Gates Foundation, the World Bank, the World Health Organization, and the Fogarty International Center of the National Institutes of Health.²⁶

2.2.1 Finding the Best Charities

It is studies like DCP2 as well as more recent ones²⁷ that GiveWell is greatly interested in. GiveWell is an organization that has been looking for the most effective charities in the world since 2007. Eight years into its search, it has found four charities that meet their criteria. One of these criteria is that the cost-effectiveness of the intervention the

24. Sassi, "Calculating QALYs, Comparing QALY and DALY Calculations."

25. Jamison et al., *Disease Control Priorities in Developing Countries*.

26. Bill & Melinda Gates Foundation, "New Initiative Will Assess Disease Control Priorities In Developing Countries."

27. The third edition of DCP2 is still forthcoming at the time of writing.

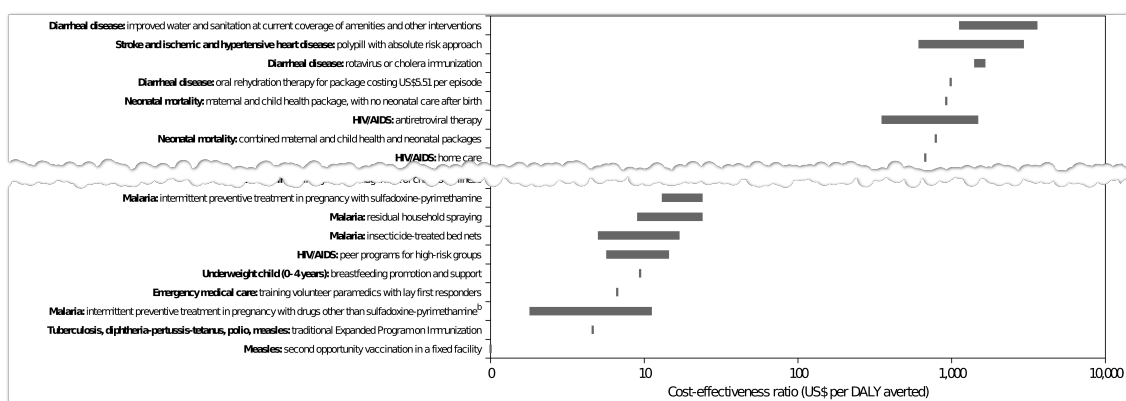


Figure 2.4: DCP2 cost-effectiveness estimates of interventions related to high-burden diseases in sub-Saharan Africa.

charity implements has to exceed that of unconditional cash transfers, the reasoning being that simply giving cash to people was such a simple and scalable intervention that any more complex and thus risky programs need to prove that they are at least better than unconditional cash transfers in expectation.

Several such programs exist, but only two – the distribution of long-lasting insecticide-treated bed nets and mass deworming in schools – have implementing charities that GiveWell can confidently recommend.²⁸ These charities have funding gaps of over \$100 million for 2016–2018, meaning that tens of millions of people will end up untreated or unprotected if the funding goals cannot be met. If the Against Malaria Foundation could reach everyone who needs a bednet, that funding gap could easily be five times as high.²⁹

According to GiveWell’s estimates, both these priority programs are about ten times as cost-effective as unconditional cash transfers. Expressed in dollars per DALY averted, the authors of DCP2 estimate that in sub-Saharan Africa, averting one DALY from malaria through distributions of insecticide-treated bednets costs \$5–17; averting one DALY from HIV/AIDS through condom promotion and distribution costs \$52–112; and averting another DALY from HIV/AIDS through antiretroviral therapy costs \$350–1,494.³⁰ Even though the cost-effectiveness estimates span three orders of magnitude, GiveWell has at least considered all of these interventions, and there are charities implementing them.

For completeness it should be noted that GiveWell’s latest estimates at the time of writing put LLIN distributions at \$2,838 per life saved.³¹ People who die of malaria usually die at age five or younger before they can build natural immunity.

28. In order to recommend a charity, GiveWell requires detailed information from the charity, so that the cooperation of the charity in the process is paramount. Hence there may be excellent charities besides GiveWell’s top charities, but these charities merely declined to cooperate with GiveWell. One reason for that can be that they are already receiving enough funding, so that extra funding is not worth the staff time for them.

29. GiveWell, “A Conversation with Marcy Erskine on September 18, 2013.”

30. I am ignoring deworming in this comparison because GiveWell found errors in the DCP2 estimates for the intervention. GiveWell, “Errors in DCP2 Cost-Effectiveness Estimate for Deworming”

31. GiveWell, “Against Malaria Foundation (AMF).”

For simplicity I will assume that QALYS can be converted to DALYS in the following sample calculations, that the life expectancy is around 57 years,³² that discounting for age and uncertainty can be disregarded, and that the average life incurs less than 0.1 YLD per year.³³ Then I arrive at about \$58 per DALY ($\sigma = 8$).³⁴

2.2.2 Motivations of Donors

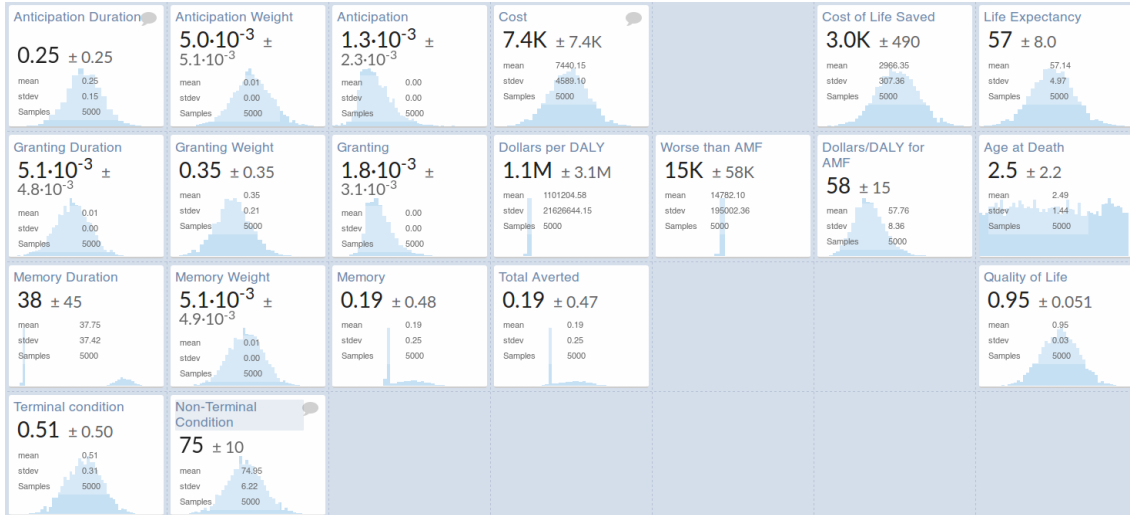


Figure 2.5: A Fermi estimate (using Guesstimate) for comparison of AMF and the Make-A-Wish foundation that assumes that anticipation and memory of a granted wish are on average as good as not suffering hearing loss and the wish itself is on average as desirable as being spared severe chest injury.

Make-A-Wish. But there are also charities like the Make-A-Wish Foundation that invest \$7,500 on average to make a wish come true for a child with a life-threatening (but not necessarily terminal) illness.³⁵ Based on the information in the organisation’s FAQ, I have created a simple estimate (using the Guesstimate app) as an example of the rough Fermi calculations that would help almost any donor improve their giving decisions.

In this model, I assume that the desirability of being granted a wish is somewhere on the spectrum of the WHO disability weights up to being relieved of severe multiple sclerosis (disability weight of 0.707,³⁶ one of the most severe disabilities listed). The estimate puts the cost for one DALY averted for the Make-A-Wish Foundation very roughly somewhere in the area of a million dollars.

32. The World Bank, “Sub-Saharan Africa (Developing Only) – Data.”

33. Institute for Health Metrics and Evaluation, “Global Burden of Disease Compare – IHME Viz Hub.”

34. Drescher, “Make-A-Wish Dollars Per DALY Averted.”

35. Make-A-Wish® America, “Frequently Asked Questions | Make-A-Wish® America.”

36. World Health Organization, “Global Burden of Disease 2004 Update.”

Since such quantitative estimates are rough and murky, one can instead do a thought experiment to compare the intervention of Make-A-Wish to a life-saving one like that of the Against Malaria Foundation. One can imagine that there is a mother with a dying child. She has only \$7,500 and can spend it either to grant her child a last wish or to try an experimental treatment that might save the child's life. How unlikely would the success of the treatment have to be before the mother would choose to grant the wish instead? The probability could probably be so small as to be intuitively indistinguishable from being infinitesimal.

Yet, in 2014 and in the US alone, people chose to donate almost \$66 million to Make-A-Wish.³⁷ \$66 million that would have easily filled the funding gap of all GiveWell top charities in 2014, and \$66 million that had the same impact as a donation of a few thousand dollars to AMF would have had, so that they have come at the opportunity cost of more than 23 thousand lives lost to malaria, lives that could have been saved.

Make-A-Wish is just one example. As noted in the introduction, the impact of 96% of US donations³⁸ could have been multiplied by an approximate factor of one hundred by investing them in a developing country.

Types of donors. The Money For Good II study³⁹ provides some quantification and explanation of the problem. The authors investigated the motivations of donors and clustered them into six categories, which they describe with some key statements from the respective donors:

“repayer” (23%) “I give to my alma mater”; “I support organizations that have had an impact on me or a loved one.”

“casual giver” (18%) “I give to well known nonprofits because it isn't very complicated.”

“high impact” (16%) “I support causes that seem overlooked”; “I give to nonprofits I feel are doing the most good.”

“faith based” (16%) “We give to our church”; “We only give to organizations that fit with our religious beliefs.”

“see the difference” (13%) “I think it's important to support local charities”; “I give to small organizations where I feel I can make a difference.”

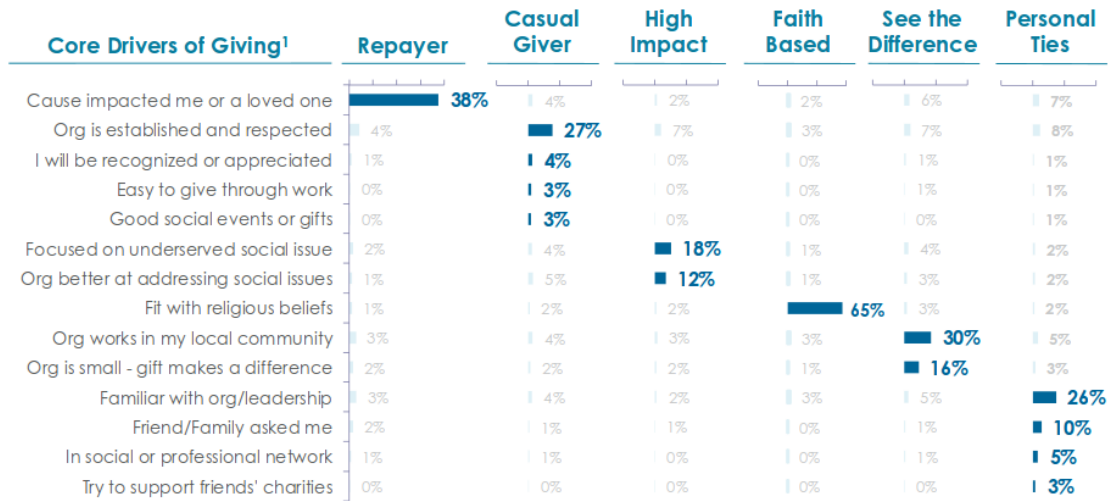
“personal ties” (14%) “I give when I am familiar with the people who run an organization.”

For each of these categories, they investigated the driving motivations of the donors for specific donations.

37. ProPublica, “Nonprofit Explorer – Make-A-Wish Foundation of America.”

38. Giving USA Foundation, *Giving USA 2015*.

39. Hope Consulting, “Money For Good II.”



1. The segments were derived by grouping individuals who had similar priorities across these "Core Drivers" of giving. We tested for multiple segmentations (from 3-9 groupings) and found this breakout of six segments to be the most robust. The %'s represent the relative importance of each variable to each segment's decision making for charitable giving. "I care deeply about the cause" was important to all segments so was removed from the analysis (it's more of a table stake than a driver of segment-specific decision making).

Figure 2.6: Key drivers of donations.

These categories suggest links to the moral foundations that Graham et al. have found. For example, the faith-based and personal ties donations may be motivated by feelings of loyalty. This may also explain the strong propensity of people to donate to causes within their country despite the higher price. The high percentage of donations made as quasi repayment may be an artifact of the high evolutionary value of not defecting in prisoner's dilemma-like situations, only that here these donations come at a high opportunity cost, leading to a tragedy of the commons.

That said, the categories leave some questions unanswered, for example what portion of "faith-based" giving goes to organizations that work on furthering religious ends, such as churches, and what percentage goes to any charity. The breakdown does make clear that no other factor that was studied had significant influence on these donors' choices,⁴⁰ but "fit with religious beliefs" is the only attribute that establishes this category, so it may be that a significant number of them donated to any nondescript charity that happened not to conflict with their religious beliefs.

Similarly, the "personal ties" category leaves open what portion of these people is doing a favor to friends and family by donating and what portion is merely so risk-averse that they only trust organizations that are run by people they know personally. The leading attribute, "Familiar with org/leadership," indicates the second while the runner-up attribute, "Friend/Family asked me," indicates the first. A more interesting grouping may be one into "personal ties" and "risk averse," where the first only captures the factors "Friend/Family asked me" and "Try to support friends' charities."

40. Hope Consulting, "Money For Good II," 134.

The “GiveWell market.” The study also found that only about 6% of donors do any comparisons between charities before donating. On a market where the cost-effectiveness of products varies by several orders of magnitude, this indifference is striking. The top reasons for not conducting research (any research, not just comparative research) are that donors felt they were familiar with the organization, that it is a well-known organization, that they are involved with the organization, that the organization was recommended to them by a trusted person, that it is a religious institution, and that the donation was small. Few of these could double as arguments against comparative research. Inaccessibility or absence of information was only cited as a problem by few donors.⁴¹ On the other hand they also identified effectiveness as the key unmet need of individual donors that were interested in conducting research.⁴² Taking into account that better research would sway some donors toward not donating and some toward donating to better organizations, the researchers estimated that better information could move about \$10 billion per year from individual donors toward higher-performing organizations (about \$15 billion in total, taking foundations and advisers into account).

GiveWell requested the raw data of the study to determine in a more fine-grained fashion the size of the market that it serves. It presents three estimates. Broadly, the donations from all donors already interested in comparing and evaluating multiple organizations total \$4.1 billion according to its estimates. When it narrowed the facet down by excluding donors that ranked what it considered to be factors irrelevant to impact higher than relevant factors, the total shrunk only slightly to \$3.8 billion. Only when they narrowed the facet to donors that both fit the first group and ranked expected impact highest did the amount shrink to \$554 million.

However, it also notes that the market may be still much bigger since successful businesses often create solutions to problems that only a fraction of its future customers had previously been aware of, or in the author’s terms, “for most successful businesses I can think of, it’s still the case that most people aren’t customers of them.”⁴³

2.2.3 A Market for Impact

Once a significant number of people are explicitly donating where they see the greatest potential for impact, something else will happen: a new market will emerge. The reason the prices that GasBuddy tracks are as close to each other as they are is that there are market pressures that do not allow some companies to sell the same product at a multiple of the price – no one would buy it. Conversely there is a strong incentive for companies to become more cost-effective in their production so that they can hand down the lower prices to their customers to attract more of them.

Analogously, as soon as charities notice that there is money to be made in becoming more effective, be it even just millions rather than billions of dollars, they will be incentivized to become more cost-effective too. With 86% of all donations coming from loyal

41. Hope Consulting, “Money For Good II,” 46.

42. *Ibid.*, 50.

43. GiveWell, “The Money for Good Study.”

donors,⁴⁴ there is great value not only in redirecting donations to more effective organizations but also in nudging established organizations to become more effective.

One variation of this idea is to simply pay people, organizations, or countries for the impact they have already achieved, thus incentivizing impact and providing means for further impact, which, in turn, can be purchased. Countries like Britain and Norway, and the Gates Foundation are experimenting with this approach,⁴⁵ and impact certificates are an example of such an effort by individual funders.⁴⁶

2.3 Prioritization Research

Having established the importance of prioritization in the introduction, this chapter will introduce to methodologies, constraints, and the interest groups of prioritization research.

2.3.1 Constraints

Purview of prior research. One early effort at creating a consolidated overview of cost-effectiveness measures for priorities setting was the first edition of *Disease Control Priorities in Developing Countries* (DCP1) published in 1993. It surveyed 52 interventions, or more than 100 when grouped by cause, as some have multiple objectives. For this research they drew on studies dating back into the 1970s.⁴⁷

The second edition expanded this set to 319 interventions whose cost-effectiveness measures were compiled by a staff of 350, more than four times the staff of DCP1.⁴⁸ The authors are quick to acknowledge the specificity of their purview:

“ A frequent, often justified, criticism of cost-effectiveness analyses is that they address only one of many criteria that could be used to evaluate health interventions. Epidemiological, medical, political, ethical, and cultural factors often also play important roles in the decision to allocate resources to a specific health condition or intervention; however, determining how one might weigh cost-effectiveness ratios alongside these other considerations when setting priorities for spending is difficult. Musgrove (1999) shows how to take some of these connections into account, including circumstances in which cost-effectiveness is an adequate criterion by itself. One approach is for the policy maker to think of cost-effectiveness ratios as the relative “price” of purchasing a unit of health (a DALY, for instance) using different interventions. These costs, along with the budget constraint,

44. Hope Consulting, “Money For Good II,” 140.

45. The Economist, “It’s Not What You Spend.”

46. Impact Purchase, “The Impact Purchase.”

47. Jamison et al., *Disease Control Priorities in Developing Countries*.

48. Jamison et al., *Disease Control Priorities in Developing Countries*.

can help determine the optimal allocation of resources among a given set of interventions.

”

This clarification omits what exactly is meant by cost-effectiveness analysis in this context. If, for example, cultural factors make condom distributions infeasible, all this would mean is that a policy maker would incur high costs due to the marketing necessary to change cultural norms or that they would incur low adoption rates and thus low effectiveness. In either case the effect could again be expressed in terms of cost-effectiveness. The implementing party might also lose acceptance among the population, incurring a different type of cost in addition to the monetary one. The effectiveness might suffer, too, in that the intervention may be cut short after the next election so that it will have less time to earn back some $O(1)$ startup costs. Or if an epidemic is about to break out, a treatment for the disease that used to be less cost-effective could become highly cost-effective in view of the rapid rate of future infections that containment of the disease can forestall.

But that is not what is meant by cost-effectiveness analysis in the context of DCP2. Cost-effectiveness analyses of communal, medical interventions are hard to create, and usually require randomized controlled trials (RCTs) at scales much smaller than the interventions policy makers want to implement. It is RCTs like these that DCP2 compiles and cost-effectiveness analyses like these that it is referring to throughout.

As GiveWell has found repeatedly, such estimates can be flawed.⁴⁹ Even more importantly, however, their external validity is highly limited. A recent paper published by the IDinsight team – and endorsed even by a staunch critic of RCTs in development economics⁵⁰ – introduces the taxonomy of knowledge-focused evaluations (KFEs), “those primarily designed to build global knowledge about development interventions and theory,” and decision-focused evaluations (DFEs), those “driven by implementer demand, tailored to implementer needs and constraints, and embedded within implementer structures,” that serves as basis for their recommendation: “Where the primary need is for rigorous evidence to directly inform a particular development programme or policy, DFEs will usually be the more appropriate tool. KFEs, in turn, should be employed when the primary objective is to advance development theory or in instances when we expect high external validity, *ex ante*. This recalibration will require expanding the use of DFEs and greater targeting in the use of KFEs.”⁵¹

Measures at scale. While RCTs have low external validity, self-evaluations of NGO that are scaling programs cannot feasibly measure their impact.⁵² Doing so would be very costly, they could only serve half the population because of the randomized control group, the results would likely be of poor quality because of the tremendous effort it would take

49. GiveWell, “Errors in DCP2 Cost-Effectiveness Estimate for Deworming.”

50. Pritchett, “Using ‘Random’ Right: New Insights from IDinsight Team.”

51. Shah et al., “Evaluations with Impact.”

52. Evidence Action has published an article contrasting these sources of insight, but the article seems to sell a standard practice as something controversial and contrarian. Nassar, “Why We Do Not Measure Impact At Scale”

to monitor the operation, and while randomization of individual intervention on village level may avoid problems of inequity,⁵³ it may also produce less reliable results.

What NGOs do instead is that they measure proxies that the prior research allows them to translate to impact fairly reliably. AMF, for example, measures net condition and use in intervals of six months after each distribution. These are data that can then be used to extrapolate what the impact in terms of DALYS averted may be.

Limited value of information. Another factor that will crop up repeatedly in this text is that the expected value of information sometimes does not warrant the attention of the researchers. If the same time can be used to find interventions that are more effective or have larger funding gaps at a sufficient level of effectiveness, then these should be prioritized within the prioritization process. There may be more hurdles of the same kind depending on the context, as the next section will explain.

2.3.2 Methodologies

2.3.2.1 Government Aid

Priorities setting in the political domain is a highly complex process that emerges as a multitude of agents, some willing to compromise to greater degrees than others, struggle to realize their preferences, whether agent-neutral or agent-relative.

Dead Aid arrives at an unfavorable verdict on the impact of governmental development aid over the past decades.⁵⁴ At the same time, many invaluable research efforts (such as many cited in this text) were partially funded by governmental or government-funded institutions, and similar institutions also implement interventions with strong evidentiary bases. Whether all the government activities over the past decades add up to a net plus or minus is impossible to determine.

Instead of trying to answer such sweeping questions as whether multilateral development aid is beneficial or harmful, Banerjee and Duflo, as well as the many outputs of the Abdul Latif Jameel Poverty Action Lab in general, make many concrete and highly specific recommendations for interventions that require action from governments or large funders but remain relatively neglected.⁵⁵

Even though *Poor Economics* was published in 2011 and many of the studies it cites much earlier still, at least some of the interventions – such as school-based deworming – remain underfunded. Hauck and Smith investigated the concrete reasons for the hold-ups, which resulted in an analyses of typical problems of the political process itself. They identified the following explanatory models in the literature:⁵⁶

53. Haushofer, Reisinger, and Shapiro, “Your Gain Is My Pain: Negative Psychological Externalities of Cash Transfers.”

54. Moyo, *Dead Aid*.

55. Banerjee and Duflo, *Poor Economics*.

56. Hauck and Smith, “The Politics of Priority Setting in Health: A Political Economy Perspective.”

“

1. Models of competing interest groups, based on the assumption that powerful interest groups – for example provider organizations – may seek to skew decisions in their own favour at the expense of less organized stakeholders.
2. Voting models, such as the median voter model, (Ahmed and Greene 2000), which assert that political decision-makers will seek to develop policies that attract particular voter groups, in an effort to maximize political support. The implication of this insight for priority setting is that the size and contents of a public benefits package may be skewed towards the preferences of key voting groups.
3. Bureaucratic decision-making models, which assert that ‘bureaucrats’ may make decisions in their own interests, rather than the interests of the population as a whole.
4. Institutional economics to explain how governmental institutions and more generally the organisation of the political system influence decisions
5. Decentralization and contracting-out to non-governmental organizations, and implications on provision of public services.

”

Since the processes of priorities setting are so deeply embedded in a such a much more complex system, it is not the ideal context in which to explain any methodologies.

2.3.2.2 Foundations

The Good Ventures foundation cooperates closely with the Open Philanthropy Project for prioritization research. It is highly transparent about its processes and thinking, and I am familiar with its work, two advantages that make it a good exemplar to explain one methodology that foundations employ to allocate funding.⁵⁷

Given the findings of “Money For Good II” that only 39% of foundations seek to maximize their impact with their grants and 36% never question whether some set of cause areas they are associated with is in fact the most impactful focus of their grants (often, perhaps, because their by-laws stipulate it),⁵⁸ the methodology of the Open Philanthropy Project is probably not representative of foundations in general, but it is informed by interviews with scores of other foundations, and, crucially, it is closely aligned with the thinking that AMF applies in its context.

Learning from other foundations. One basic first step is to tap the knowledge that already exists out there. Since other foundations may not be as transparent about their

57. Open Philanthropy Project, “Our Process.”

58. Hope Consulting, “Money For Good II,” 64.

decision making but may well have much valuable experience in the field, co-funding of projects is one opportunity to access that experience.

Exploring potential causes. *Focus areas*, as the Open Philanthropy Project calls them, or cause areas are often fixed for foundations by decree of the founder or funder. Unless this person has invested significant research into this decision and updates it as the changing circumstances require, it is a matter of luck whether the cause area harbors any impactful investment opportunities.

The founder of AMF has boiled this down to a credo to the effect that the goal of every charity should be to shut down,⁵⁹ because it is founded to solve a specific problem, and needs to solve it even while that means that it will make itself obsolete.

Foundations should take the same approach to choosing cause areas. Hence the Open Philanthropy Project tries to gain an overview of the whole spectrum of cause areas before them and ranks them according to three core criteria:

“

Importance. How many individuals does this issue affect, and how deeply?

Neglectedness. All else equal, we prefer causes that receive less attention from others, particularly other major philanthropists.

Tractability. We look for clear ways in which a funder could contribute to progress.

”

Given the sheer number of such fields, the funder needs to conduct many fairly shallow investigations. These investigations aim to answer the questions of what the problem is and how important it is; what avenues there are for addressing it; and who else is working on it.

Having gained this cursory overview, some causes will stand out and thus warrant deeper investigation: interviews with scores of experts of the field, extensive reviews of the academic literature, and so-called learning grants that open doors to the networks of grantees.

Selecting and prioritizing cause areas. Having investigated these top cause areas, more factors are taken into account to filter and rank them, such as the capacities and aptitudes of the staff at the time.

Hiring and grantmaking. Finally a process of delegation is initiated whereby one or several domain experts are hired to assume responsibility of the area in the long term. Thanks to sometimes decade-long experience in the field, these people bring knowledge to the organization that by far exceeds that of the foundation staff even after their deep

59. Paraphrased from a verbal conversation.

investigations of the areas. These new hires are gradually integrated into the organization with the goal of them eventually running the research processes in their domain nearly autonomously.

These are the stages that the Open Philanthropy Project has completed for at least⁶⁰ two focus areas so far. Within these areas the domain experts then recommend grants to the funders, especially Good Ventures.

When organizations receive grants and in response scale up their operation, there is a lock-in happening that makes the organization dependent on continued grants. Hence it is important to research individual grants rigorously and, if that should become necessary, to phase out giving with exit grants.⁶¹ In any case, the opportunity costs of a less than optimal decision are high, but a less than optimal selection of cause areas would incur the combined opportunity cost of all grants within it, so it is a much greater wager.

2.3.2.3 Individual Donors

The methodology of the Open Philanthropy Project is a useful template for understanding the processes for individual donors, but it is joined by a few problems that lead to key characteristics that are probably hard to avoid.⁶²

One key problem is that donors transfer their level of risk aversion from investments into highly agent-relative preferences to investments into agent-neutral investments.

Foundations may fund hundreds of projects hoping that one or two will succeed to a degree that makes their investment cost-effective. In areas where experience is scarce, this is often the only process of knowledge creation there is.

Individual donors are usually only able to fund a fraction of one such project and when it fails, as it is likely to, such a donor would be disillusioned and possibly cease donating all together. This is unlikely to change, so the recommendations prioritization research can make to individual donors are limited to comparatively low risk–low return opportunities.⁶³

Strong focus on individual organizations. While foundations can think in terms of making grants toward cause areas, individual donors are much more inclined to think in terms of donating to individual charities.⁶⁴ For the foundation the most important part of the research is in selecting the cause areas to focus on, but for individual donors who cannot devote much time to the process, having to review charities even within one cause

60. At the time of my writing, the process page lists two, but they have hired an expert on farmed animal welfare as well, Open Philanthropy Project, “Introduction To Risk Aversion” implying that they have completed these stages for a third area.

61. Open Philanthropy Project, “Giving Now vs. Later.”

62. GiveWell, “Process for Identifying Top Charities.”

63. The informal experience of the Open Philanthropy Project, however, has shown that taking away this constriction does not necessarily yield higher expected return but rather merely greater funding gaps, and these are not critical for most individual donors, so that they are not necessarily missing out. Matthews, “You have \$8 Billion. You Want to Do As Much Good As Possible. What Do You Do?”

64. The donor coordination system introduced in appendix 6.2 aims to address this problem.

area is not feasible. Hence organizations in this space – such as GiveWell, Giving What We Can, and Animal Charity Evaluators – need to make recommendations of specific charities and make them with high confidence.

Funding gaps. Foundations are often accountable to a small set of highly engaged donors that exert control over them. In these cases they are not accountable to thousands of donors, all with their own opinions or preferences for how their money should be used, the way most other charities are. However, when a foundation wants to invest its funding highly effectively, and one opportunity it has found is one that many individual donors are also donating to, it runs into two problems.

First, it cannot fill the funding gap of that giving opportunity completely or it would deprive the individual donors of their giving opportunity. Some may not respond by doing research and finding a new target for their donations but simply by stopping to donate, lost donations that may as well count as costs the foundation incurred through its grant. Hence the foundation needs to make sure that it projects the donation influx of the charity correctly and then grants only so much and so often that its grants do not fill the charity’s funding gap.

Second, however, even that is not enough since such projections are by necessity mostly based on past data, so that donations that the charity receives at any time will influence future projections upward and thus foundation grants downward. This relationship is not a very direct one, but it is direct enough for GiveWell to go to some lengths to avoid it, because individual donations would become partially fungible with foundation grants, so that the foundation suddenly morally ought to be accountable to those donors. That is not an option for a foundation that wants to make experimental grants without a complex democratic decision process, or, vice versa, drive away donors that are attentive enough to notice this dynamic.⁶⁵

Harmonizing grants and funding gaps such that other donation flows do not influence them and there are still no contingencies of over or underfunding an intervention is a delicate balancing act.

Marketing. A foundation rarely needs to worry about its public perception, but an organization that makes public recommendations has a strong interest in its own growth, be it in terms of donors, in terms of money moved to its recommended charities, or any other such metric. Hence a foundation has considerable freedom in its grant making which the first organization lacks.

The result is that the tractibility criterion (cf. the taxonomy of importance, neglectedness, and tractibility in section 2.3.2.2) is heavily influenced by how uncontroversial the recommendations are in the current cultural or political discourse of the donors. That is a significant limitation.

65. Open Philanthropy Project, “Giving Now vs. Later.”

2.4 Strategy of the Against Malaria Foundation

AMF's prioritization strategy is a consequence of its hybrid position both as a fundraiser that receives donations from a wide range of small and major donors and as a grant maker. Donors to AMF are, by necessity, agnostic about the location of the net distribution that they fund. AMF needs to first hold significant funding before distribution partners will engage with them in the lengthy proposal process. Hence donors can only afterwards know what distribution their donation has been assigned to.

The decision to focus on the prevention and thus elimination of malaria via the most cost-effective means available at the time and now is one that must have been made prior to the founding of the foundation in 2004. It may have been that the process that led to the decision was as powerful as that of GiveWell and thus yielded similar results, but it is probably more parsimonious to assume that there was an amount of luck involved and that I would be writing about another organization now had that luck not favored the founders of AMF.

Apart from this problem of cause selection, however, many of the themes introduced above also mark the decision processes of AMF. The founder has provided me with a prioritized list of the factors that they routinely consider as part of this process. I will explain them in order.

Malaria burden. Of course, in order to be considered, a region needs to have a high malaria burden. AMF used to demand from its distribution partners that they collect data on the rates of malaria cases from health centers in the region, but the quality of the data, when it was even available, was consistently low and the effort of collecting it high, so that AMF has more recently made exceptions to this requirement. GiveWell has also ceased to track these data.⁶⁶ There is no question that malaria needs to be a “material problem for public health”⁶⁷ in a region for AMF to consider it, but the value of information of the most precise data available on the rate does not warrant its collection.



Figure 2.7: “The process includes awareness through banners, speeches and performances on the perils of malaria.” Photo and caption courtesy of AMF.

66. GiveWell, “Against Malaria Foundation (AMF),” bullet point “Malaria case rate data”.

67. Rob Mather, email conversation with the author, January 2016. Cited with permission. All further quotations are from the same conversation unless marked otherwise.

Net gap. The second most important criterion is the magnitude of the net gap, the number of nets needed in a region. Here AMF distinguishes two categories of distributions, strategic and tactical ones. Strategic means that the national malaria control program of a country seeks funders for distributions two to four years down the line; tactical means that a distributions by other funders have left regions unprotected, usually for lack of funds, so that AMF fills these gaps to provide closer to universal coverage. These distributions are more short term.

Accountability. Another critical criterion is whether the national malaria control program welcomes the enhanced accountability measures that AMF demands. For a discussion of a related problem, see the paragraph on “trusted partners” below.

Non-net costs. Transport and distribution of the nets is of course associated with significant costs, not even taking into account the cost of the nets themselves. In the experience of the foundation, it is psychologically beneficial to ask donors to pay for something very concrete like the nets, even though the non-net costs may well be higher leverage given that they are essential and yet constitute but a fraction of the net costs. Hence the website of AMF is geared toward fundraising only for nets – when making a donation one can choose whether to enter a monetary value and have the number of nets it funds calculated or whether one wants to enter a number of nets and have the monetary value calculated.

The founder told me that this has not been a problem for them in the past as distribution partners are often able and willing to carry non-net costs and because the foundation also contacts major donors and individually seeks permission to convert their contributions into unrestricted funding; these donors are usually eager to give their permission.

Certainly since the unrestricted donations of several tens of millions of dollars from Good Ventures, AMF has the means to cover non-net costs when its distribution partners can not, but historically, this has been a critical factor. Since the recommendation by GiveWell the decision to fundraise for restricted grants has sometimes been criticized as

Aside 3: Counterfactual Validity

There are large, multilaterally endowed funders of net distributions, such as the Global Fund, and to an individual donor it is not obvious that such a funder will not observe that a region where it planned to conduct a distribution has already been served by AMF so that it will cancel its plans.

This would not be a problem if the donor could assume that, in turn, the use of the Global Fund’s money in the case where it does not benefit the distribution is roughly as good as AMF’s use of the donation. But while the Global Fund is highly transparent and invests into a portfolio of research-backed medical interventions that are considered highly cost-effective by developed-world standards, it yet contains interventions that are markedly less cost-effective than LLIN distributions even under the unrealistically favorable conditions of randomized trials.

veiled complicity in perpetuating the overhead myth,⁶⁸ but AMF avoids any more explicit endorsement of it.

Trusted partners. The last major factor is whether there is a distribution partner in the country with whom AMF is comfortable working. In 2013 this posed a problem when AMF tried for the first time to finalize a distribution of several millions of nets, an endeavor for which they had to cooperate with several other parties, some of whom did not fully comply with AMF’s requirements during the proposal phase. The distribution fell through.⁶⁹ At the time of writing, however, talks for distributions of similar scale are progressing well.

Factors of lesser importance. Besides the above factors, I encountered lesser factors that could potentially influence prioritization decisions. Rob Mather, however, told me that they have little overall influence.

First, since natural immunity can be lost after years of not being exposed to malarial mosquitoes, I imagined that redistributions in regions where it can be assumed, after more than three years, that most nets are in bad condition might take precedence over fresh distributions in regions where no nets had previously been distributed. But according to AMF this has not been an issue that they considered so far. Natural immunity does not fully account for the drop in the fatality rate around that age. Rather it is due to many factors “including better physiological reserves, immunity, and access to treatment,”⁷⁰ so that the malaria case rate may increase (though the consensus view is that it tends not to) but without an accompanying increase in mortality.

Second, the problem with counterfactual validity that I also consider in section 2.1.5 only becomes relevant when there is significant funding available compared to the funding gap. For a hypothetical explanation, see the aside on counterfactual validity.

Following that reasoning, I considered that it might also be relevant for AMF to make sure that the regions it serves would not have been served by any other funder had AMF not intervened.



Figure 2.8: “Distribution team sets an iron ledge over unstable bridge to help people cross by foot.” Photo and caption courtesy of AMF.

68. The widespread misconception that the ratio of program funding to administrative funding says anything about the effectiveness of a charity.

69. Against Malaria Foundation, “AMF Update – 28 November 2013.”

70. GiveWell, “A Conversation with Professor Christian Lengeler, August 20, 2015.”

However, one important assumption mentioned above, that the funding is significant compared to the funding gap, is not given in the case of LLIN distributions, so that AMF considers this a minor factor even in strategic distributions. (That it is not a factor in tactical distributions holds by definition.) AMF estimates the annual funding need at \$5 billion and the available funding from all funders at \$2.5 to \$3 billion,⁷¹ so it is unlikely that such a replacement effect will occur. AMF, however, notes that this might happen at lower geographic levels, for example within countries. Funding, however, is usually budgeted at the large funders that fund net distributions, so that even if such replacement takes place, the funds would be spent with roughly equivalent cost-effectiveness on another distribution.

Third, insecticide resistance has not yet been a factor to decide where to conduct a distribution or whether to conduct one in a specific region but rather to decide on the type of net to use. Rob Mather specifically pointed out that in regions with higher levels of insecticide resistance they use nets treated with piperonyl butoxide (PBO) in addition to the insecticide, but also noted the experimental nature of this very recent technology.⁷²

“ PBO is not an insecticide by itself, but acts by inhibiting certain metabolic enzymes (e.g. mixed-function oxidases) within the mosquito. These enzymes detoxify or sequester insecticides before they can have a toxic effect on the mosquito. Therefore, a PBO LLIN would in theory have an increased killing effect on malaria vectors that harbour such resistance mechanisms compared to a pyrethroid treated LLIN. However, the entomological and epidemiological impact of PBO LLINS is expected to vary according to the bioavailability and retention of PBO, and the level, intensity and mechanisms of insecticide resistance for local vectors, as well as across different transmission settings. ”

PBO is crucial because while there are different insecticides (notably permethrin and deltamethrin), both of them are insecticides of the same class, pyrethroid, and mosquitoes typically form resistance to a class of insecticide rather than to individual chemicals. Research on a new class of insecticide is ongoing.

Fourth, lacking infrastructure in a country has not thwarted distributions in the past. A good partner, good planning, and patience are key. “For example, the Democratic Republic of Congo (DRC) is one of the most challenging places in the world to work given ten years of war in recent times, the scale and geography of the country amongst other reasons. However, we recently demonstrated with considerable (but not perfect) success, that it is possible to use smartphone technology to collect data from 250,000 households and produce accurate data to support a net distribution, including knowing the location of all

71. But not that Rob Mather added that “over the years it has been difficult, for those attempting to collect data and get a clear picture, to know numbers with great certainty. Collating information from all funders, including governments both in and outside Africa, seems challenging.”

72. Organization, “Conditions for Use of Long-Lasting Insecticidal Nets Treated with a Pyrethroid and Piperonyl Butoxide.”

households to within 6 metres.”

Fifth, AMF agreed that the political situation in a country can at times convince a partner that it would not be safe for them to conduct a distribution, so that it needs to be postponed.

2.5 Distributions of Long-Lasting Insecticide-Treated Mosquito Nets

This section will give a brief overview of the intervention the Against Malaria Foundation implements.

2.5.1 Key Epidemiological Facts

Malaria is caused by the parasite *Plasmodium*, of which the form *Plasmodium falciparum* is the most common one in Africa. It is not contagious but rather is spread almost exclusively by female *Anopheles* mosquitoes, which bite mainly during the night. While there are strains of malaria that affect other animals and humans can be infected with these strains in rare occasions, there are no known cases of these zoonotic strains then spreading to another human.⁷³ Furthermore, the *Anopheles* species that are the primary malaria vectors in Africa are strongly anthropophilic, that is, have a strong preference for human blood over the blood of other animals.⁷⁴

The most at-risk groups are pregnant women and children under the age of five because in

Aside 4: Mortality and Fertility

Some people object to AMF’s program on the basis that saving lives in general anywhere in the world is, supposedly, of negative net value.

It is highly unclear whether the expected negative effect, if any, of an additional person on all other sentient beings is sufficient to offset any positive net value the life has for the person, and the relative importance of human overpopulation is doubtful as well.^a But these critics are correct in that an increase in population may result – an increase, however, that is likely to be smaller and less certain than intuitively apparent.

Roodman has conducted a metastudy of fourteen studies of the influence of life-saving interventions on population developments. His conclusion is that: “The best interpretation of the available evidence is that the impact of life-saving interventions on fertility and population growth varies by context, above all with total fertility, and is rarely greater than 1:1.”^b

Other highly effective interventions, such as deworming, improve rather than save lives, and are thus immune to this criticism.

a. Monbiot, “There’s a Population Crisis All Right. But Probably Not the One You Think.”

b. Roodman, “The Impact of Life-Saving Interventions on Fertility.”

73. WHO, “Malaria.”

74. CDC, “Malaria – About Malaria – Biology – Mosquitoes – *Anopheles* Mosquitoes.”

both cases children are in danger that have not yet developed natural immunity to malaria. This immunity usually forms around the age of five and leads to few lethal infections at older ages. It is, however, important to note that this immunity can also be lost over time when no reinfections occur.⁷⁵

While there are a few chemicals that can be used as prophylaxis against malaria, there is no market-ready vaccine to date. The development of vaccines against malaria has proved challenging for numerous reasons, among them the high degree of polymorphism of *P. falciparum* proteins and the high rate of reproduction, which leads to a rapid evolutionary development of resistance.⁷⁶ As a result, the history of malaria control has largely been a history of malaria vector control.

2.5.2 Brief History

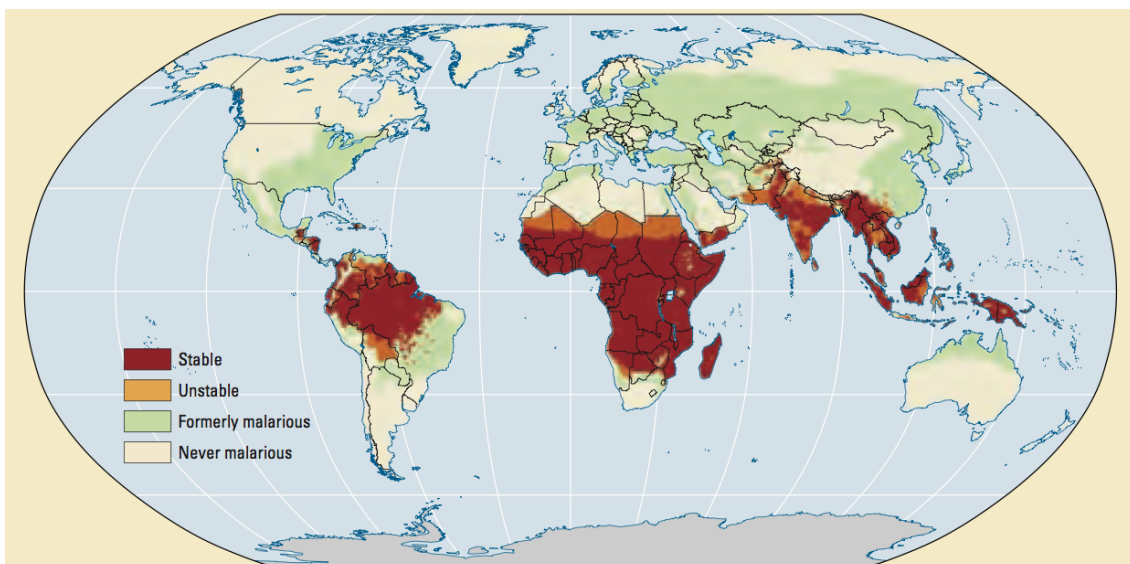


Figure 2.9: Many developed countries were formally malarious. (Image courtesy of Rigg et al.)

Malaria has existed for at least one hundred thousand years, but it only spread to humans a few thousand years ago. It was mentioned in ancient Chinese texts of 2550 BCE, in ancient Egyptian texts of 1325 BCE, and ancient Indian texts of around 500 BCE. The last author already correctly identified insects as the vectors of malaria. Hippocrates disagreed and thought weather, dirty drinking water, or poor health led to malaria.

The ancient Romans finally gave rise to the term *malaria* that we use today. They correctly associated malaria with marshes and swamplands, but thought it was the fetid gases

75. Webb, "The Long Shadow of Malaria Interventions in Tropical Africa."

76. Crompton, Pierce, and Miller, "Advances and Challenges in Malaria Vaccine Development."

from these areas that caused the disease, so that their term for “bad air,” *mala aria*, gained currency. Malaria took a heavy toll on the Roman empire, and by weakening Rome’s soldiers and civilians, may have contributed to its collapse.⁷⁷

In the 1940s and 1950s, large-scale and well-funded campaigns eliminated malaria in the US and much of Europe. They involved the drainage of stagnant water, installation of window screens, and spraying of DDT.⁷⁸ DDT was later deemed too dangerous and banned, but mosquitoes also developed resistance to it, so even a lifting of the ban would have limited positive effects.⁷⁹

In the following, I want to focus on the one form of vector control relevant to this thesis, the distribution of long-lasting insecticide-treated bed nets. Nets have been used for protection against mosquitoes since around 500 BCE, but only during World War II the practice of impregnating nets gained currency.

Mosquitoes would often crawl underneath the rims of the nets, find rents in them, bite body parts that were touching the net, or would even be trapped underneath the net from the start. Insecticide-treated nets do not have these disadvantages because mosquitoes would either die or at least be repelled by the insecti-

cide, and in the latter case the insecticide may curtail their life expectancy to the point that the malaria parasite will not have time to fully develop before the death of its host.

These insecticide-treated nets formed a key part of the Roll Back Malaria program of the World Health Organization (WHO), the United Nations Children’s Fund (UNICEF), the

Aside 5: Insect Suffering

Critics sometimes object that killing mosquitos to save humans is speciesist, especially when enough mosquitoes have to die to save one human that the years of life lost for the mosquitos combined outweigh the years of human life saved. According to a quick “back-of-the-envelope” estimate based on data from the CDC, GiveWell, and WHOPEs,^a this is typically the case, with the years of mosquito life lost outweighing human life years saved by three orders of magnitude – and that is only counting the directly affected first generation, not their lost offspring.

What the view above ignores, however, is that mosquitoes may not, on balance, have positive lives, because they are very short and usually end painfully. Tomasik, an advocate for wild animals including insects, thinks that the jury is still out on whether insecticides as well as more human methods of vector control are net good as they reduce populations and thus spare future generations of the species a net negative life, or whether the direct suffering and life lost are worse. Hence, he argues for the usage of *humane* insecticide.^b

a. Drescher, “Mosquito YLL Per Human \neg YLL.”

b. Tomasik, “Humane Insecticides.”

77. Goldsmith, *Battling Malaria: On the Front Lines against a Global Killer*, 9–13.

78. Marcus, *Malaria*, 72; Rigg et al., “The World Development Report 2009 ‘Reshapes Economic Geography’,” 117.

79. Marcus, *Malaria*, 68.

World Bank, and the United Nations Development Programme (UNDP).⁸⁰ Yet they had the problem that they needed frequent retreatment with fresh insecticide.

This situation was greatly improved when LLINs, long-lasting insecticide-treated bed-nets were developed and the WHO shifted their investments toward the new models in starting in 2002. These nets stay impregnated for their full lifetime of about two to three years. This improvement instated them clearly as the currently most cost-effective means of malaria prevention.⁸¹

In 2011, however, the Global Fund, one major distributor of LLINs, was hit by the global financial crisis as well as its discovery that \$25 million had vanished from some of its community programs.⁸² Such fluctuations are particularly dangerous because the protection adults had enjoyed up until then may have undermined their natural immunity, so that timely redistributions are critical.

Another important development was that originally the WHO recommended targeted distributions where the groups most at risk of malaria, pregnant women and children under five, would receive nets in distributions. In 2007, however, the WHO made the switch to recommending universal distributions.⁸³ In addition to the protection LLINs lend to the individual, the WHO also views their distribution as a vector control intervention with considerable communal effects. Humans can be seen as bait that lures mosquitoes to the protected sleeping spaces where the insecticide kills them. Surviving mosquitoes, at least the anthropophilic species, who avoided contact with LLINs or have developed resistance to the insecticide, may still die of lack of feeding. Thus universal distributions can reduce the mosquito population below some critical threshold that will bar the parasites from completing their life cycle, thus eliminating malaria. The mechanism requires that almost all people in a region be protected by LLINs. Hence, universal distributions have to be conducted with great attention to detail.

For further reading I recommend *The Long Struggle against Malaria in Tropical Africa* as the most comprehensive volume on recent malaria elimination efforts that I have found during my research.⁸⁴ *The Making of a Tropical Disease* gives a wider perspective on the history of the disease.⁸⁵

2.5.3 Role of the Against Malaria Foundation

AMF, founded in 2004, has been a rather minor funder of LLINs, but it is now rapidly scaling up. Its small scale has enabled it to pilot a monitoring methodology that was unprecedented by distributions outside scientific trials. This is a crucial innovation for three reasons.

80. Packard, *The Making of a Tropical Disease*, 228–229.

81. Masum et al., “Africa’s Largest Long-Lasting Insecticide-Treated Net Producer.”

82. Terra Daily, “Global Fund Faces Billion-Dollar Gap.”

83. World Health Organization, “Insecticide-Treated Mosquito Nets.”

84. Webb, *The Long Struggle against Malaria in Tropical Africa*.

85. Packard, *The Making of a Tropical Disease*.

First, adults tend to underestimate the harm of the disease; since it is rarely deadly for them, they neglect prevention and rather buy treatment, which is more expensive in expectation. The psychology of these nonoptimal decisions has been analyzed in *Poor Economics*.⁸⁶ Hence, education is a critical part of every net distribution, and so is regular monitoring of the actual behavior of the recipients afterwards to detect weaknesses in the education model.

Second, corruption is widespread in many countries that suffer from a great malaria burden. Agreements on strong monitoring regulations ward off corruption by making the distribution unappealing to corrupt decision makers in potential partner organizations.

And third, the rigorous monitoring yields comprehensive measurements on many proxy variables for impact – data with high informational value for prioritization research.



Figure 2.10: “GPS coordinates of each house that received nets were recorded and uploaded each day.” Photo and caption courtesy of AMF.

86. Banerjee and Duflo, *Poor Economics*, 59 ff.

“As our own species is in the process of proving, one cannot have superior science and inferior morals. The combination is unstable and self-destructing.”

Arthur C. Clarke

3

Requirements and Modeling

3.1 Forming the Idea

The motivation behind the specific approach I applied to deciding on the topic of this thesis was that I wanted to find the idea that provided the best fit for the requirements of the university while having the maximal positive impact on the world. To this end, I contacted twelve organizations working on pressing and neglected problems and described my circumstances to them so they could determine whether there were any projects on their backlogs that might fit my requirements.

Relevant attributes. The first step toward forming the idea for this thesis was to define the relevant attributes of any proposal along which they would be prioritized. The final set was one consisting of six factors, whose estimation was necessarily highly informal, informed by the proposal of the organization and my own critical assessment. The estimates distinguish three levels indicating high fit, low fit, and uncertainty.

Impact. The potential for positive impact that I thought the proposal promised after hearing the arguments of the proposing charity.

Replaceability. The ease with which the charity could realize the project without my help. If they would be able to do it themselves at a relatively small cost, it would be incompatible with a high estimate for impact.

Solidity. A measure for how certain the project is to succeed. A software development project is comparatively likely to succeed while an exploratory data analysis might yield no interesting insights and a project that depends on forming a community of users might easily fail to materialize.

University fit. The degree to which the proposal fits the requirements of the university.

Personal fit. The degree to which the proposal fits with my abilities.

Personal potential. The degree to which the proposal fits with my future plans.

Four states emerged from the correspondence with the charities:

No reply. Two of the organizations did not reply despite reminders.

Replied. Four of the organizations replied, but the discussion did not advance further. One, for example, told me they had a project with good fit in their pipeline, but asked me to answer a few questions first. I did, but the contact person left the organization before replying, which I found out when I tried to contact other people of the organization later on.

No need. Three organizations discussed the idea in their team but found that they had no projects they could use my help for.

Delivered. Three organization had proposals for me and invited my feedback.

One hurdle had been that most of these organizations work on problems where volunteers are generally not worth the staff time necessary to train them, and that insofar as they do draw on volunteer work, it is for clearly delimited tasks that require either very little training or training that has been captured in writing.

Among the three responses was one pure software project that would not have satisfied university requirements. The most interesting proposals came from Animal Charity Evaluators and the Against Malaria Foundation.

The first group presented me with eight ready proposals that they were looking for volunteers to realize. Three of them scored low on solidity as they relied on data that would have to be obtained from corporate entities. It seemed highly uncertain to me whether I would be able to obtain these data even after extensive correspondence with a wide range of them. Three more were proposals for studies with human subjects and required significant amounts of contact with these subjects. Since important decisions will be based on the results and I have great respect of the difficulties of this study format combined with little prior experience, I felt a low personal fit. Finally two studies scored high on all counts, a study of historical influence of advocacy campaigns on legislation and a study on the effectiveness of specific types of social media advocacy. Below the gists:

“ **Legislation:** Study the impact of legislative advocacy by determining the total amount of funding spent lobbying and advertising for pro-animal legislation in successful and unsuccessful campaigns and comparing this to the estimated amount of suffering saved by successful campaigns and the proportion of funding that went to successful campaigns.

Social Network Spread: Track the flow of common animal advocacy messages through online or in-person social networks to compare speed of dispersal. Possible axes of comparison include positivity/negativity, system/individual focus, and specificity.

”

The Against Malaria Foundation (AMF) provided the proposal for analysis of their electronically collection distribution data, whose nature – in its developed state – will become evident from the remainder of this section.

Additionally, I generated an idea for a donor coordination system that piqued the interest of the Centre for Effective Altruism, a UK nonprofit. This idea is subject of appendix 6.2.

I investigated all three remaining proposals. With the study on legislative changes I identified the following problems:

1. It was unclear to me what degree of legal knowledge would be necessary to assess the impact of legislative changes.
2. Even if the required legal knowledge was minimal, assessing levels of animal well-being is fraught with uncertainties from domains including psychology, biology, medicine, economics, and philosophy. I feel like I have a sufficient grasp on the extent of these uncertainties not to make premature judgments on related issues, but conversely the uncertainties will only in rare cases allow me to make confident judgments. Since most legal changes that I am aware of from the recent past were of small, incremental nature, finding the ones that do allow for unequivocal assessments with small margin of error would be hard and might not be achievable in a reasonable timeframe.
3. Even if I manage to locate unequivocally successful changes in legislation that I can quantify with low margins of error, the process of attributing these to specific advocacy campaigns will be hard or impossible in most cases since there are usually different groups working towards subtly different ends with different levels of efficiency, all of which would require assessing.
4. In those cases where all previous requirements are met, including a relatively clear attribution to one or few campaign efforts, I would still need to obtain data on the funding that benefited the campaign or campaigns, data that may be confidential or, after a longer time, hard to reconstruct.
5. Even if there are campaigns that meet all these requirements, conclusions from this small and highly selective set would probably not extrapolate well to future campaigns.

The study on social media spread seemed easier to realize. The only problems I discovered were:

1. There is already a significant amount of data and research on social media spread that might extrapolate fairly well to the specific types of messaging that the study would have prioritized.

2. Running the study would require access to social media accounts with sufficient reach and a representative base of followers. The organization might be able to provide me with the first, but its followers are by no means representative of the larger population.

The following are concerns I had about the proposal of the AMF:

1. It addresses a highly specific set of problems, which limits my optimism about its impact.
2. Many of its features can only be used if there is sufficient data at a level that so far few local partners of AMF have agreed to collect. If this highly detailed level of data collection should remain an exception in the future, many features of the software will find little use.
3. In extension of the previous point, some features are still not in use because the data is not yet available.

My own idea has the following disadvantages, the reason I am only developing the plan pending such a time when I have the support of an established organization to drive it further:

1. It depends on the continued cooperation of a sizable community of donors.
2. It depends on the continued cooperation of many nonprofits.
3. Some of its failure scenarios are silent, meaning that it is hard to know whether they have occurred.

Given these considerations, I decided to prioritize the project of AMF and complement it with my plans for the donor coordination solution.

3.2 Data Exploration

Having obtained the first set of distribution data from AMF, I investigated the nature of the data with a specific focus on its collection, format, peculiarities (possible errors), and unequivocally erroneous entries.

The first set of data was predistribution registration survey (PDRS) data of the distribution of 676,000 long-lasting insecticide-treated nets (LLINs) in Kasai Occidental, Democratic Republic of the Congo (DRC), funded by AMF in August to November 2014. Their local partner that conducted the distribution, IMA World Health, collected the survey data using an OpenDataKit-based electronic system.

AMF conducts universal distributions. (The distinction between universal and targeted distributions is explained in section 2.5.2.) To gauge the demand of nets per region and thus the nets that need to be available at each distribution hub, they have their local partner organizations conduct predistribution registration surveys, which determine the number of sleeping spaces per household and the number of existing and usable LLINs.

3.2.1 The Survey Form

In the present case, the distribution partner asked for a number more data points in the survey, which are not of primary interest to AMF. Nonetheless these are interesting data, and AMF informed me that the government has indicated interest in the data to augment its dated census data.

I performed my initial analyses in R, but soon transferred my findings to Python, where they are now implemented as part of cleaning and analysis routines.

Language. The survey data consists of 57 variables, of which at least 10 are determined automatically. The French versions are authoritative but since I do not know French and AMF also prefers to work with English labels, they created a spreadsheet of translations; many of the English labels seemed automatically translated to me, so I asked a French native speaker to kindly correct them for me. The original French versions and the English translations are given in appendix 6.1.

A native speaker confirmed that “Est ce que votre ménage a au moins une Moustiquaire” asks whether there is “at least one mosquito net” as opposed to “more than one.” In the spreadsheet, “moins_une_MILD_HH” is translated as “More than one net,” but the same native speaker said that “moins_une_MILD_HH” alone is too compressed to make sense either way. The meaning of this variable is hence almost certainly “at least one.”

Missing questions. In several cases the questions associated with certain variables are not given in the spreadsheet, which may be due to them being answered in advance by the interviewer rather than the household representative. The interviewer may know what to answer from a preparatory briefing.

Restrictive response options. The final problem that I see with the survey is that it asks for the source and the brands of the preexisting nets, but if a household has several preexisting nets from different sources, then they may also be of different brands, and both questions will become hard to answer. This is true of 106,824 responses in the sample.



Figure 3.1: Open Data Kit storing coordinates (of St. Albans, Hertfordshire, UK).

Oddly, the question asking for the sex of the respondent had response options in English rather than French in the original and did not allow for intersex options.¹ It also only asked of the representative of the household, so that the sex of the other members remains unknown.

3.2.2 The Survey Data

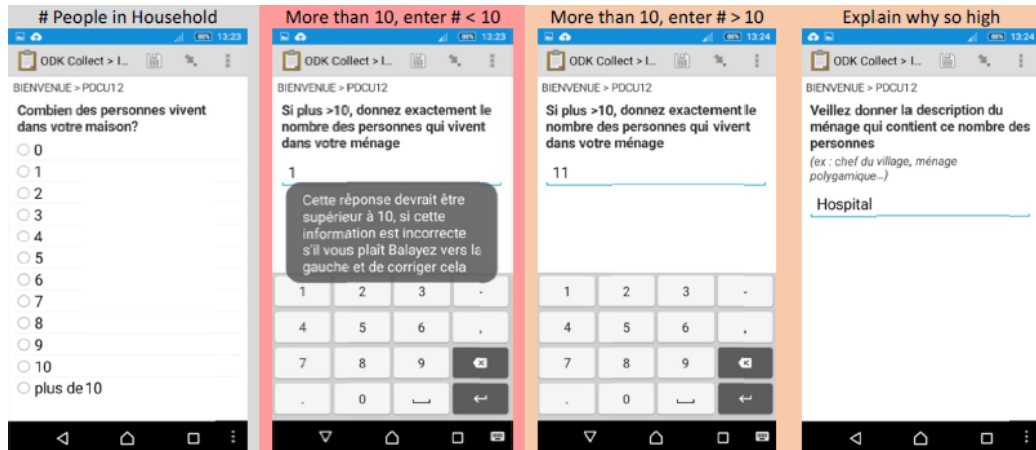


Figure 3.2: The software tries to prevent, catch, or ask for confirmation for unusual inputs.

Although there were some problems with the data, as detailed below, most of the data were of sufficient quality. The government of the DRC even indicated interest in the data as they are more up-to-date than the only census data that it has available.

Invalid rows. I found that when date, one of the GPS coordinates, the confirmation of the legal script, and the day of the distribution were missing from rows, that they were either empty or invalid.

Most of the 9938 rows fell into the first category, but four seemed to be aggregates of other rows, claiming to be samples of households with several hundreds of thousands of tenants. I filter all of these rows during the import.

Dates. The dates of the rows are broken in so many ways that it seems likely that several problems coincided here to produce them. My findings are detailed in section 5.3.1.

1. Rather intrusively, the survey did not ask for people’s gender but their sex, and about 111,000 (45%) of the respondents answered that with “woman” while about 124,000 (51%) responded “man.” Almost 10,000 (4%) did not answer, maybe because they found neither appropriate, the question too private, or for various other reasons.

Opt-outs. Only 516 (0.2%) of the respondents opted out of the survey and thus probably the distribution.

Age. The age distribution was significantly skewed upwards in comparison to that of the DRC overall² with a median age of 33 rather than 18, surely because the respondents were the people the household sent to represent them and not a random sample.

People per household. There are a few errors up to almost 241,000 people per household, so the mean is skewed upward, but the median of the people per household is 5 with a median absolute deviation of 3.

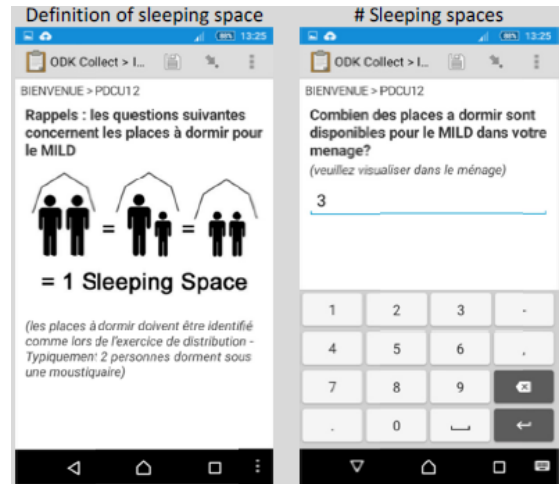


Figure 3.3: The software provides guidance to avoid errors it would be unable to catch.

3.2.3 Forthcoming Data

Planning the LLIN Visualizer, I already made provisions for postdistribution check-up (PDCU) data to be imported into the database, but AMF has not yet been able to send me the raw data.

The data is intended to shed light on the actual usage of the nets. To this end, PDCUs are conducted in six-month intervals after every distribution. Five percent of the recipients are randomly selected and surveyed, a process during which condition and usage of the nets are determined.

In order to find the households and to match them to the respective rows in the PDRS data set, the houses are marked. This matching may be imperfect if the marks get removed or the village moves due to changes in seasonal labor. A separate household relation in the LLIN Visualizer is meant to provide a flexible fuzzy mapping between survey rows to link responses of the same household at different points in time.

3.2.4 Other Data

3.2.4.1 The World Map

All the data in the LLIN Visualizer is shown on a map whose data is provided by one of two sources that the user can choose from, the standard map of OpenStreetMap and the humanitarian version of it.³

2. Central Intelligence Agency, “The World Factbook.”

3. MapQuest unfortunately shut down access to its map, which I had used as default because of its geographic features.

3.2.5 Insecticide Resistance

The vector species *Anopheles* has a quick reproductive cycle, so that it can produce resistance to insecticide quickly. The mosquitoes never cross any wide, barren areas, so that resistance forms in independent localized areas, but it is nevertheless a significant and insufficiently solved problem.

Various organizations conduct tests for insecticide resistance, and the website IR Mapper⁴ provides an overview of the last sixty years of such tests. the LLIN Visualizer displays this data and allows for filtering by country, insecticide class, and year range, since AMF only uses insecticides of the Pyrethroids class, and because insecticide resistance tests that were older than a decade at the time of the distribution are probably obsolete.

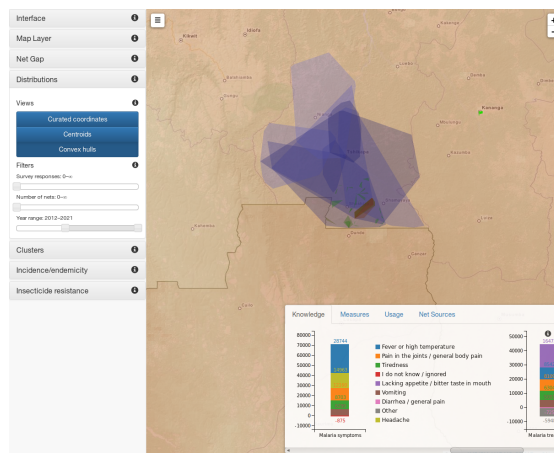


Figure 3.4: Distributions visualized as convex hulls in the LLIN Visualizer.

3.2.6 Incidence and Endemicity

The LLIN Visualizer is indebted to IR Mapper also for the data on *Plasmodium falciparum* and *Plasmodium vivax* incidence rates and the entomological inoculation rate – the number of infective bites per person per time. IR Mapper, in turn, obtained these data from the Malaria Atlas Project (MAP). The data has a high resolution even though the project could draw on fewer measurements. It achieved this by augmenting the measurements with a model that took into account “environmental covariates [such as] rainfall, temperature, land cover, and urban/rural status.”⁵

3.2.7 Net Gaps

The net gap is the difference between the number of nets that are needed in a given country and year and the number of nets that are already funded for that country and year. These are laboriously compiled by entities such as the Roll Back Malaria collaboration.⁶

AMF provided me with two of these tables, one with estimates from 2013 for 2013–2016 – including nets needed, funded, and their difference – and one newer one from 2015 with estimates for 2015–2017 but for fewer countries and only including net gaps.

4. IR Mapper, “IR Mapper.”

5. Kessler and Tschuggnall, “Malaria Atlas Project.”

6. Roll Back Malaria, “Malaria Programmatic Gap Analysis.”

The LLIN Visualizer allows for the essential filtering by year, but it also provides two different view modes. Malaria does not respect country borders, so from the medical point of view, the aggregation by country is arbitrary, so usually what users will be interested in are the absolute net gaps, answering for example questions as to how much bigger the net gap was than a distribution in a country in a certain year. But malaria eradication is also fought in the political area, and distribution partners of AMF are only active in some countries, so the relative gaps of countries are also important. Countries with a larger population will tend to have greater absolute net gaps, but maybe it is also better served by NGOs or by its government, so that the relative net gap is not unusually high. The populous Nigeria and Ethiopia have large absolute and relative net gaps, for example, but according to the 2015 data, no nets at all were funded for Malawi resulting in a net gap of about 1.7 million which would not have been noticeable on the absolute overview. That net gap is tiny in comparison to Nigeria and Ethiopia, but at the time of the estimate, AMF had never conducted a distribution of even half that size. Yet they have a particularly well-trusted partner in Malawi, Concern Universal, so that they are in a perfect position to close this gap – in a much better position, surely, than in Nigeria, where they conducted a single tiny distribution in early 2008 with a partner they have only cooperated with once.

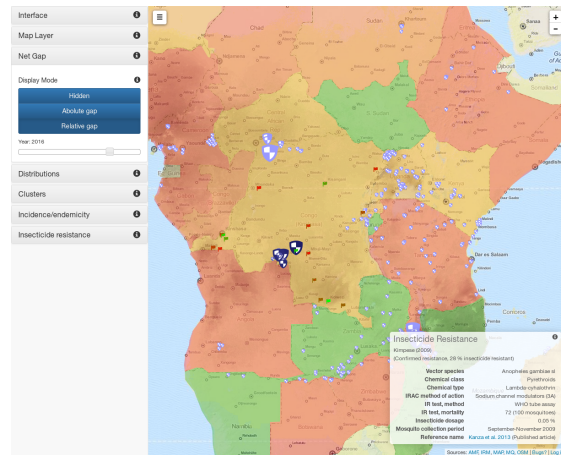


Figure 3.5: Net gaps visualized in the LLIN Visualizer.

3.2.8 Indicators

One piece of feedback I have gotten on the LLIN Visualizer is that users would like to compare the endemicity and LLIN-related data with generic, frequently-used indicators. In particular the mean GDP per capita was cited as an interesting one, but I have found many more interesting measures recorded or calculated by the World Bank, Gallup, the Center for Global Development, the World Happiness Report, and others. In response, I extended the LLIN Visualizer to display them.⁷

Confidence in government A measure of the people’s confidence in their respective national government (not government in the abstract).

7. Gallup, “How Does the Gallup World Poll Work?”; United Nations Sustainable Development Solutions Network, “World Happiness Report 2016 Update”; World Bank, “GINI index (World Bank estimate)”; Gallup, “Understanding How Gallup Uses the Cantril Scale”; Diofasi and Birdsall, “The World Bank’s Poverty Statistics Lack Median Income Data, So We Filled In the Gap Ourselves”; Hillebrandt, “Median GDP per capita: how much does the typical person earn in different countries?”

Delivery quality Delivery of government services, comprising indicators of government effectiveness, regulatory quality, rule of law, and control of corruption.

Democratic quality Quality of the democratic process, comprising indicators of voice and accountability, and political stability and absence of violence.

Freedom to make life choices Answers to the Gallup World Poll question “Are you satisfied or dissatisfied with your freedom to choose what you do with your life?”

Generosity Gallup World Poll data on donation behavior in proportion to GDP, so that poorer people are not penalized for being proportionally less giving.

Gini index The Gini index is a measure commonly used to measure income inequality.⁸ It provides an important addition to the GDP as explained below.

Life expectancy at birth It goes without saying that the life expectancy is not necessarily an age at which many people die, because especially in many sub-Saharan countries frequent child deaths skew the life expectancy downward even though people are much less likely to die once they cross a certain age threshold (around age five).

Life Ladder The Life Ladder or Cantril Self-Anchoring Striving Scale is one type of well-being assessment tool used by Gallup,⁹ who provide the data for the whole past decade an almost all recognized countries.

Mean GDP per capita The mean gross domestic product per capita is such a well-known metric that it hardly needs introduction. The data imported into the LLIN Visualizer at the time of writing are the latest World Bank estimates.

Median GDP per capita One criticism of the *mean* GDP per capita is that it both hides income and consumption inequality,¹⁰ and skews the result upward.¹¹ The latter is a problem that the median avoids. For a response to the first problem, see the Gini index above. The Median GDP per capita data are available for all the same countries and years that the mean GDP per capita figures are also available for.

Most people can be trusted The proportion of people who think that most people can be trusted as opposed to “Can’t be too careful.”

Negative affect The affect measures comprise self-reported worry, sadness, and anger on the previous day.

8. World Bank, “GINI index (World Bank estimate).”

9. Gallup, “Understanding How Gallup Uses the Cantril Scale.”

10. The current figures are consumption based.

11. Diofasi and Birdsall, “The World Bank’s Poverty Statistics Lack Median Income Data, So We Filled In the Gap Ourselves”; Hillebrandt, “Median GDP per capita: how much does the typical person earn in different countries?”

Perceptions of corruption The mean of questions from the Gallup World Poll about perceived corruption in different parts of society.

Positive affect The affect measures comprise self-reported happiness, laughter, and enjoyment on the previous day.

Social support Responses to “If you were in trouble, do you have relatives or friends you can count on to help you whenever you need them, or not?” from the Gallup World Poll.

World Happiness Report score The World Happiness Report publishes a score that is the weighted aggregate of many indicators selected to give an idea of the degree of happiness of the denizens of 157 countries. The current figures are the ones from the 2016 update, which also cover earlier years.

3.3 Functional Requirements

3.3.1 The Original Set

The architecture decisions underlying the LLIN Visualizer were shaped by the following requirements:

1. Any user will be able to view aggregated data of the net distributions.
2. Privileged users will be able to view maximally detailed net distribution data.
3. Users will be able to view relevant statistics of the aggregated data.
4. Administrators will be able to import new data into the system.
5. AMF will be able to embed the software into its website.

3.3.2 The Extension Set

I extended this set by the following optional requirements:

1. Users will be able to view past distributions even when only aggregated data is available on them.
2. Users will be able to put distribution data into the context of geographical patterns of malaria incidence, insecticide resistance, and unmet need for nets (net gaps).
3. Users will be able to filter distributions, incidence data, insecticide resistance tests, and net gaps by their date and various data type-specific attributes.
4. Maintainers will be able to extend the statistical analyses.

My reasoning was the following:

1. Even just superficial meta data of distributions – such as their date, size, and geographic location – are helpful to understand prioritization decisions of AMF and also how these changed over the years as the organization gradually scaled up its operations.

2. Malaria incidence, insecticide resistance, and net gaps are important elements in prioritization decisions. There are other factors that are harder to get data on, such as the degree to which AMF trusts various potential partner organizations in the respective countries (see section 2.4 for more details), but visualizing these three factors already provides value to the user.
3. These filters are important because time tends to gradually drive these data toward obsolescence. IR tests that, at the time of the distribution, were decades old should have little bearing on the distribution. Similarly, the chemical classes of the insecticides used have changed and will continue to change. As a result, resistance to one class of insecticide has little bearing on the feasibility of a distribution if AMF uses a different one.
4. The analyses that I implemented are those that seemed interesting to me as well as the participants of a straw poll I conducted among interested users, but it is likely that over time new insights will be desired. Then the maintainer (likely me or someone at AMF) needs to be able to implement these in some fashion that is perpendicular to most of the LLIN Visualizer, so that performing these changes will not require extensive knowledge of the code base.

3.4 Coordination with AMF

I tried to minimize the time that my questions diverted from the staff of the foundation, and so collected my questions so they could answer them all at once or, even better, I found answers to them myself before I even sent them off.

The questions fell into the categories of (1) questions about the data that I had already received, (2) technical questions about the realization of the software, and (3) content questions about topics that I needed to address in this text.

Questions about the data. The question about the odd behavior of the health area variable addressed in section 3.2.2 was one topic of discussion. A related problem was the reverse one, that in some cases names, such as village and health area names, were misspelled. I proposed a metric based on geographic distance of the point clusters and edit distance of the names, but the problem was not critical and we did not pursue it further. Another factor was that we saw no way to obtain any canonical spellings of the village names; in many cases there probably are no agreed-upon versions.

Other topics concerned the availability of further data that I could import or rather the reasons for the delays they are facing. Finally, there was a data analysis task that we discussed so that I could perform and integrate it into my work, but AMF's head of technology eventually completed the analysis himself.

Technical questions. One discussion topic was that of the distinction of access groups; we decided that three groups, anonymous visitors, logged-in users, and administrators were sufficient.

The distinction between anonymous visitors and logged-in users is necessary so that only an anonymized version of the data is publicly available but AMF can manually give special access rights to employees or people it trusts. In terms of the anonymization I proposed two methods, either just barring anonymous users from accessing the most fine-grained household-level layer for anonymous users or doing the first and making sure that k -anonymity and l -diversity are maintained on the cluster level (we did not discuss t -closeness). AMF confirmed that prohibiting access to the household layer was sufficient.



Figure 3.6: “The distributors and local children pose for a photo after a successful distribution and hanging of the nets.” Photo and caption courtesy of AMF.

Maintenance was another topic of discussion. There were two possible scenarios the choice between which we had initially left open: (1) AMF takes over maintenance of the software and runs it on its servers, or (2) I continue to maintain it and run it on one of my servers. Option 1 has the advantages that the AMF staff – mostly their head of technology – has full control over the software and can embed or adapt it in any way they see fit without being dependent on me for it. The main disadvantage is that the head of technology is not familiar with the software framework. Option 2 avoids this disadvantage at the price of some level of dependence. In the end it was option 2 that we agreed on.

Finally I made a number of suggestions for features that I envisioned for the LLIN Visualizer and asked for critique, prioritization, and additional ideas. The results are listed as part of section 5.1.

Content questions. In this category of question I have mostly corresponded with the founder and CEO of AMF on the topic of the foundation’s strategies for prioritization. I have already covered the insights from these emails in section 2.4.

“We’ll continue tomorrow – if I live.”

Paul Erdős

4

Architecture and Implementation

4.1 Architecture

the LLIN Visualizer exploits the model-view-controller design pattern to maximally separate logic, so that each layer can be discussed in turn.

4.1.1 Models

Authentication. The authentication system is provided by the Django framework. It allows for fine-grained access management and for easily administering the permissions of users or groups of users. Finally, management commands provide functionality that is only run explicitly by the administrator or by a cron daemon.

Survey responses. The analytics app is the core of the LLIN Visualizer. The SurveyResponse relation retains all the data of all surveys that have been imported into the system. These data never change but reads need to be fast to provide a good experience for the user, so I chose not to fully enforce the third normal form on this table with regard to the columns that describe the location (since one location, however defined, functionally determines village name, health area name, district name, etc.). Doing so would have been fraught with error anyway, as we will see in section 3.2.2.

The location data was “four dimensional” in the sense that it comprised latitude, longitude, altitude, and accuracy. At first I tried to use all of these coordinates for the geometric component of the table, but storing the accuracy in that fashion seemed to be unusual and for 2.5-dimensional data, including the altitude, I would have had to continually override defaults throughout the code, so I found it simpler to restrict the geometric component to two dimensions and store the rest as floating point number columns.

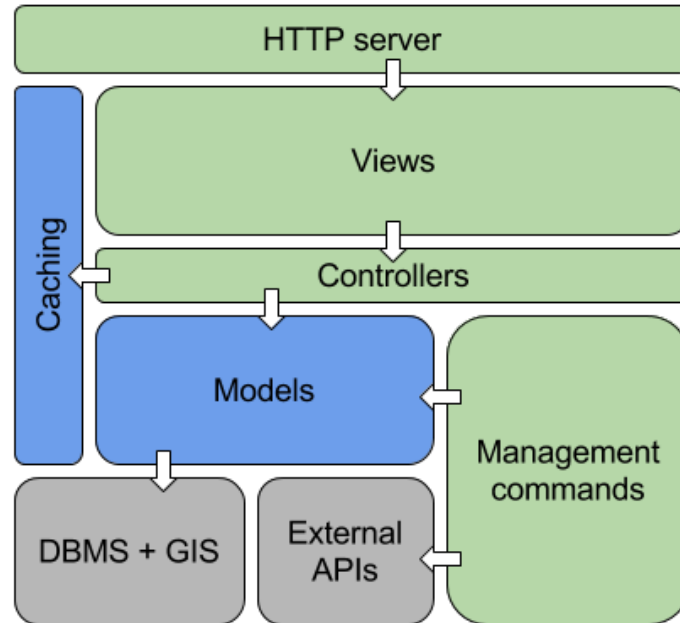


Figure 4.1: The architecture of the LLIN Visualizer. The arrows indicate what code uses which resources. Gray components are used as black boxes; blue components are concerned with data access; and green components are concerned with data manipulation.

Countries. The only location variable that allowed for normalization according to the third normal form and proved necessary for other performance-related reasons¹ is the country. It is separated out to its own relation.

Clusters. Since anonymous users must only see aggregated information for reasons of data privacy, the Cluster relation provides such an aggregation on a flexible level. Depending on the data it can sometimes be useful to provide aggregations on the level of the village, the district, or country-specific regions, such as the health areas of the Democratic Republic of the Congo.

Using the Cluster relation, nested aggregations on several such levels can be prepared, each with geographic information describing its location and expanse.

Households. The household relation only serves to link survey responses that belong to the same household at different points in time. It is less a result of the data management

1. Several types of features displayed by the JavaScript-driven map application are automatically filtered by country to limit the number of features that are displayed at the same time. A larger number would reduce the responsiveness of the application.

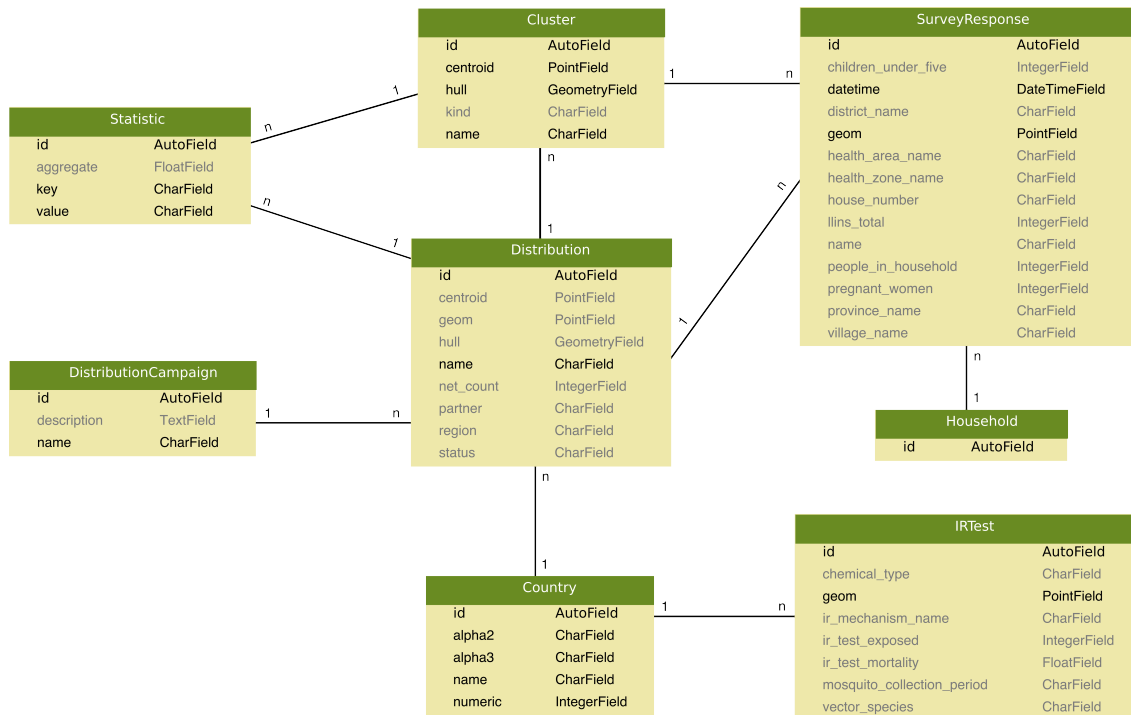


Figure 4.2: The models of the LLIN Visualizer. Some columns are omitted.

than the code layout, since it can neatly contain any logic that will be necessary to match these entries.

Distributions. Every survey response belongs to a distribution, but there are many historical distributions that do not have any survey responses associated with them. Just like clusters, distributions also carry geographic information about their location and expanse as well as various additional meta data.

Distribution campaigns. When AMF conducts distributions, they usually make one contract with a local partner stipulating several parallel or consecutive distributions in the same country or region. By “distribution campaign” I refer to the overall contract that comprises one or several distributions. Distribution campaigns have corresponding pages on the website of AMF.

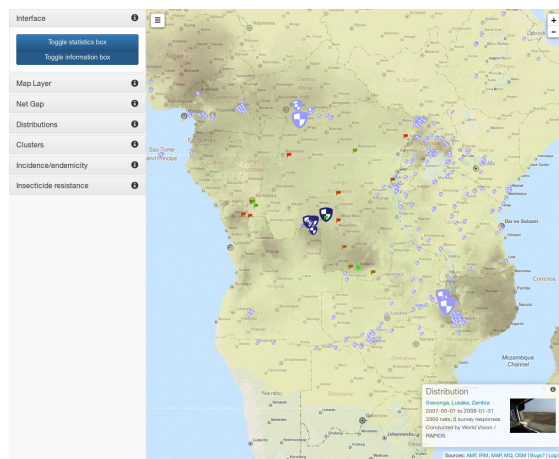


Figure 4.3: Distributions visualized in the LLIN Visualizer.

Statistics. Finally the statistics relation is a flexible aggregation of data that can be visualized for the user. The specifics of this module will be described in section 4.1.4.

All of the data in all relation can be imported or generated by management commands but also individually administered through a web-based backend.

4.1.2 Controllers

The controllers (called “views” in Django) are the most minimal layer of the software since most computations are so time-consuming and unchanging that they are prepared once by management jobs. The results are written to the database and merely served by the controllers.

Most controllers implement JSON API endpoints that are read by the JavaScript frontend.

4.1.3 Views

The views (called “templates” in Django) are centrally based upon the Bootstrap framework, CoffeeScript that compiles to JavaScript, the Leaflet map software, and the D3 data visualization library.

A light scaffold provides the basis for the application code and the tests. Upon this scaffold hang the NetLayer classes, each representing one layer of the data aggregation. The first, the distribution layer, represents all distributions, with or without survey data. Those distributions that have survey data associated with them allow for a drill-down to the flexible cluster layer – for example health zones in the case of the DRC. At the same time the corresponding data on insecticide resistance tests for the same country is displayed; it is represented by a layer of its own right. Finally, authenticated users can drill down to the household level and view individual survey responses.

The data on insecticide resistance tests and malaria incidence is provided by the Insecticide Resistance Mapper.

4.1.4 Management Commands

Most of the logic of the LLIN Visualizer is in the views and in the management commands. These are jobs that are executed very occasionally, for example, when new distribution data is imported. Most of the data whose processing is discussed below have been introduced in section 3.2.4.

Distribution import. An API on the website of AMF provides distribution data as XML dialect in a semistructured fashion. The job parses the XML, extracts the structured data, and then uses a heuristic approach to extract more data from the embedded description.

Any dates are either given as one pair of month and year, as a range of months and a year, or as a range of two pairs of month and year. A significant part of the code is concerned with date parsing.

Finally the data are written to the database atomically for performance and so no clean-up of an incomplete, aborted import is necessary.

This file, however, lacks some data, which are then scraped from the HTML of the AMF website itself, which again necessitates some heuristic parsing.

IR Mapper import. IR Mapper² provides its data as JSON, but unfortunately it is still not perfectly structured. Many numeric values are saved as strings, sometimes without apparent reason, but sometimes also because some values are in different units than others, significantly sometimes in absolute units and sometimes in percent. The import job tries hard to resolve these inconsistencies and save the data in the same unit.

It also maps idiosyncratic names for countries to their official spellings.

The data on Plasmodium falciparum and Plasmodium vivax incidence rates and the entomological inoculation rate, the number of infective bites per person per time, are also sourced from an API provided by IR Mapper in the form of raster layers. The data themselves are the product of the “Malaria Atlas Project.”³

Net gap import. The net gap data (see subsection 3.2.4) have come in tables that can be converted to CSV files. A dedicated job imports data on nets needed, financed, and net gaps severally. The user needs to specify the type of data (out of those three types) and the year of the estimate, so that the LLIN Visualizer can choose the newer estimate if there are several for a given year, country, and type.

Indicator import. The indicator data described in subsection 3.2.4 needs to be imported from various sources in several different formats. The indicator import handles this task. The user supplies the data in CSV files and specifies the type of data they are importing, and the job augments the variables known for each country-year tuple with the new data.

Survey import. The survey import job is probably the module that was most time-consuming to implement.

First, it needs to be flexible. The data AMF has provided up to this point have come in two different formats, but there is no limit to the number of different raw data formats the LLIN Visualizer needs to be able to handle. Hence, a csv parser class provides a generic interface to all the cleaning methods I have written while individual csv profiles inherit these and determine to which fields they need to be applied. At the same time they provide the flexibility of mapping fields of the input data to arbitrary database fields, not only one to one, but also by merging fields or splitting them up.

Second, the cleaning of the data requires flexibility too. Each resulting field needs to have custom cleaning routines associated with it but there also need to be higher-level cleaning routines that can decide whether the data from several fields in combination make sense or whether the whole row is invalid. Due to some intricacies of the order

2. IR Mapper, “IR Mapper.”

3. Kessler and Tschuggnall, “Malaria Atlas Project.”

in which type conversions are applied, I could not draw on the form validation tools that Django provides. The key difference was that these escalate field-level errors to the user to correct or re-enter the data. In our case, this is either impossible or the job of the cleaning routines.

Many of the cleaning routines perform type conversions but there are a few more sophisticated operations involved as well. The date parsing draws on contextual data from the AMF website to sanity-check the dates and discard them if they seem erroneous. Many cleaning routines apply translations that convert fields with values of “oui” and “non” to Boolean types, and normalize various strings to key values that are standardized within the LLIN Visualizer. Finally, names of locations are normalized using a sequence of heuristics that are optimized to merge many different spellings of the same name while introducing few errors.

The module allows the administrator fine-grained control over its behavior by specifying the parser and distribution to be used, specifying some properties of the csv file, and specifying the mode in which the database writer should operate. There is also a mode that allows for the testing of new parser profiles.

The writer is resilient in that it can recover after a failed insertion and continue processing. This is only possible when it commits after every row, so while the atomic writer has a significant performance advantage, it is not applicable when the data contain errors. It is not used by default, but when the user knows that the data contain no errors, the option can be activated.

Cluster generation. The cluster generation job at the same time generates the statistics associated with the clusters and runs the same calculations for distributions, which are structurally identical to clusters. In detail:

1. The first job groups all responses to the surveys by their respective distributions and then uses PostGIS functions to calculate the centroids and convex hulls of these points. These data are then used to augment the distribution relation.

Below a simplified select query in relational algebra using a syntax based on “RelaX”⁴ by way of illustration of the actual update query used in the management command. Note that the relational algebra syntax resembles pivoting more than the group by syntax I chose in the code, so that the structure is slightly different from the SQL version.

```
 $\pi$  distribution_id, st_centroid(collection)  $\rightarrow$  centroid, st_convexhull(collection)  $\rightarrow$  hull  
 $\gamma$  distribution_id;st_collect(geom)  $\rightarrow$  collection  
analytics_surveyresponse
```

With the currently available data, this query takes about 20 s to execute on my computer, aided by the index on the primary key column. These measurements have

4. Kessler and Tschuggnall, “RelaX.”

no absolute relevance but provide comparative values that allow me to optimize them, aided by PostgreSQL explain analyze command, and compare them to the others.

2. The second job is parameterized with `kind` and `fields`, where `kind` is a name for the type of cluster, such as “village” or “health area,” and `fields` are the grouping fields that describe this type. Here it is important to note that for a grouping by village it is better to group by village name, health area name, health zone name, and even larger categories rather than just by village name, since there are many villages with the same name.

The query performs the same PostGIS operations as the previous one, but then also concatenates the values of the grouping fields to one comma-joined string, such as “Shakunga, Kabungu,” as a name for the cluster. These names are by necessity unique.

A simplified select query in relational algebra using the `kind` “village” and the `fields` “village_name” “health_area_name” by way of illustration of the actual parameterized insert query used in the management command:

```
 $\pi$  'village',
    concat(',', ' ', village_name, health_area_name),
    distribution_id,
    st_centroid(collection)  $\rightarrow$  centroid,
    st_convexhull(collection)  $\rightarrow$  hull
 $\gamma$  village_name, health_area_name, distribution_id;
    st_collect(geom)  $\rightarrow$  collection
analytics_surveyresponse
```

This query was significantly sped up – from 17 s to 10 s execution time on my computer – by indices over the typical grouping columns “village_name” and “health_area_name” in addition to the index on the primary key column due to an external merge sort that PostgreSQL perform over the tuple of these three columns. The choice which columns to use for cluster generation, however, is one that is up to the administrator, so it makes sense to add indices like these according to the actual usage patterns rather than in advance during development.

Finally, the job updates all survey responses with relations to their respective clusters, using the generated names for identification.

3. The third job generates the cluster and distribution statistics. The type of the grouping is again parameterized through the foreign key field and the referenced relation. The job creates descriptive statistics such as averages, maximums, standard deviations, counts, and sums of numeric and geographic data in scalar and vector formats.

It heeds the definitions of the data structures from the survey importer and interprets them as consisting of the names of keys and the names of values, to which it assigns aggregate results. An example will clarify the operation.

One simple case is the cluster centroid: The centroid is assigned various attributes to enrich it, namely the average distance of the points it describes, their maximal distance, the standard deviation of their distance, and their number. PostGIS functions are used to determine these values for the centroids of every cluster. In total there are currently 60 such tuples, which are computed for each of 8 distributions and over 10,000 clusters.

This abstract approach to the generation of cluster statistics has become necessary because requirements for statistical analyses can change (expand) quickly and should only require minimal changes to configuration parameters. At the same time the data are scalar in some cases and vectorial in others, and the parameter combinations that require calculation are sometimes flat and sometimes nested (so that each path tuple requires a separate calculation). This more abstract system is sufficiently powerful to limn all these use cases.

This query is heavily parameterized and uses many PostgreSQL features that have no equivalent in any relational algebra syntax that I am aware of, such as array operations and the unnest operation on arrays.

This giant query takes about 14 s to execute on my computer, aided by the existing indices, and I could not find any leverage points to further optimize it.

4.2 Technology

4.2.1 Core Technologies



Figure 4.4: Git, Python, and Django logos.

Git CVS⁵ and SVN,⁶ the other version control systems that I have worked with, hold no sway over Git⁷ in terms of the useability that I have experienced. Only Mercurial⁸ was similar to Git, but I have not gathered much experience with it.

5. Free Software Foundation, “CVS.”

6. Apache Software Foundation, “Apache Subversion.”

7. Software Freedom Conservancy, “Git.”

8. Mercurial community, “Mercurial SCM.”

Python The decision to implement the project in the programming language Python⁹ was motivated by my more than ten years of experience using the language, which in turn is thanks to the elegant syntax of the language, its versatility, and consequently the rapid pace of development that it lends itself to.

Django the LLIN Visualizer, by virtue of its functional requirements, needed to have a frontend exposed to the web. There are a number of web frameworks for Python, some that aim to be minimal and thus light and highly flexible and others that aim to be feature rich.

Typically I would use a minimal framework to reduce the complexity of the project whenever possible, but experience has taught me that unless I can see with certainty that the project will remain small, it is usually the case that its complexity will increase drastically over time. In expectation of the more complex project that will emerge half a year down the line, my strategy is now usually the opposite, namely, to go with a feature-rich framework unless the project clearly does not call for it, or some of the features clash with the developers plans for the program architecture. CherryPy,¹⁰ Flask,¹¹ and Pyramid¹² are examples of such lighter frameworks; Django¹³ is an example of a feature-rich one.

Thanks to Django, the LLIN Visualizer has database migration support, an object-relational mapper, a user authentication system, and much more built in. One place where the assumptions of the framework clashed with my preferences for the LLIN Visualizer, however, was the templating engine. Here I greatly prefer Jinjaz¹⁴ over Django's built-in templating engine. Luckily, it has become easy to select alternative templating engines in Django.

PostgreSQL The main alternative to the database management system PostgreSQL¹⁵ would have been MySQL,¹⁶ but I have had more positive experiences with PostgreSQL. Both systems are well supported by Django.

PostGIS PostGIS¹⁷ is an extension to PostgreSQL that adds support for geographic objects and forms the basis for many calculations in the LLIN Visualizer. It is also well supported by Django.

9. Python Software Foundation, "Welcome to Python.org."

10. CherryPy, "CherryPy."

11. Flask, "Flask."

12. Pylons Project, "Pylons Project: Pyramid: About."

13. The Web Framework for Perfectionists with Deadlines, "Django."

14. Ronacher, "Jinjazz."

15. PostgreSQL Global Development Group, "PostgreSQL."

16. Oracle Corporation, "MySQL."

17. PostGIS Project Steering Committee, "PostGIS."

4.2.2 Other Important Technologies

Buildout Buildout¹⁸ makes sure that when one can build an application with two commands in 2016, one will still be able to build it with the same commands in 2018. In practice it sometimes falls short of this goal when Python libraries have dependencies on libraries or C bindings provided by the operating system, but it goes a long way toward making builds repeatable.

Another advantage is that one can easily integrate packages from outside the Python realm, such as NodeJS or Ruby packages.

Leaflet When I did research on what map library to employ in the LLIN Visualizer, Leaflet¹⁹ and OpenLayers²⁰ were the ones that were recommended most often. OpenLayers was praised as the well-established standard solution, Leaflet as the new and simpler runner-up. Both seemed to be well suited for the purposes of the LLIN Visualizer. What swayed me toward Leaflet is that I found that several companies had switched from OpenLayers to Leaflet when the latter software had sufficiently matured. Given the small differences between the libraries in functional terms and the feasibility of switching within the first days of development, I did not want to invest more than a few hours into the comparison and decided on Leaflet.

OpenStreetMap OpenStreetMap²¹ is a collaborative effort to create a free map of the world. This map – which, unfortunately, is fairly uninformative for many of the most interesting distribution regions – supplies the background layers upon which the LLIN Visualizer displays the net data.

QGIS This GIS application allows for graphical insight into the data stored in PostgreSQL.²² I have been using it to prototype new analyses, which I would later make part of the software proper.

CoffeeScript There are a number of humorous presentations on the Internet that introduce various absurdities of the JavaScript language design.²³ When writing JavaScript code, one option is to learn these and learn how to work around them, another is to use CoffeeScript,²⁴ which compiles to JavaScript. I also enjoy its syntax.

Jinja2 The default templating language of Django takes the idea of separating presentation from logic to an extreme by going to some lengths to force the developer to keep any logic out of the templates. But the purpose of the separation used to be to keep the layers of the MVC model orthogonal, and when the developer is forced to

18. Zope Foundation, “Buildout.”

19. Agafonkin, “Leaflet.”

20. OpenLayers community, “OpenLayers 3.”

21. OpenStreetMap, “OpenStreetMap.”

22. QGIS, “QGIS.”

23. Bernhardt, “Wat.”

24. CoffeeScript, “CoffeeScript.”

implement clearly view-related logic in the controller because the templating language is not powerful enough to handle it, the orthogonality suffers in the opposite direction. Admittedly, there is the option of writing custom template tags to capture custom template logic, but then the language is also limited in ways that do not seem to be purposeful.

Jinja2,²⁵ on the other hand, is a very powerful templating language that follows more in the spirit of Python in that it trusts the developer to make sane decisions in regard to the separation of view logic from controller logic.

Sphinx The documentation of the LLIN Visualizer is written in such a way as to enable other developer to understand and extend the software. As such it takes the form of code comments, which Sphinx²⁶ can automatically compile into a structure of web pages that forms a complete manual to the software. Reading natural language is a cognitively different process from reading code, so that excessive code comments can seem intrusive. In Python they are contained within one string at the start of a module or method, and with Sphinx one can even read them completely separately from the code.

GeoDjango Django ships with a component that extends the abstraction layers of the framework to spacial data. GeoDjango²⁷ provides extensions to the object-relational mapper, the validation framework, and the managements commands, and also contains a host of utilities for the manipulation of spacial data outside the database.

Handlebars JavaScript prior to ECMAScript 6²⁸ did not have its own string interpolation syntax, and while ECMAScript 6 and CoffeeScript do provide this feature, especially longer strings, such as long sections of the HTML of a website, can be formatted much more comfortably with a templating language as Handlebars²⁹ provides it.

D3 The website of D3³⁰ gives an overview of the enormous versatility of the library. It is a JavaScript library that uses SVGs on top of the HTML and CSS stack to visualize data of just about any structure imaginable. the LLIN Visualizer is not using D3 directly but through the graphing library C3.³¹ All statistical visualization are rendered using this library.

Jasmine While Django ships with its own testing framework for the Python code, much of the LLIN Visualizer is implemented in CoffeeScript, which needs to be unit-tested as well, and for that the LLIN Visualizer employs Jasmine,³² a testing framework that

25. Ronacher, "Jinja2."

26. Brandl and the Sphinx team, "Sphinx."

27. Django Software Foundation, "GeoDjango."

28. Ecma International, "ECMAScript."

29. Katz, "Handlebars."

30. Bostock, "D3.js."

31. Bostock, "C3.js."

32. Pivotal Labs, "Jasmine."

runs in the browser alternative to the frontend the visitor sees but drawing on the same code base.

Fabric Since the LLIN Visualizer is not a distributed system, the deployment does not pose a challenge, but even so it is greatly simplified with Fabric³³ scripts that execute common tasks programmatically on the remote server.

33. Forcier, "Fabric."

5

Evaluation and Learnings

5.1 Functional Optimization

In December 2015, I initiated the public beta phase of the LLIN Visualizer. Early users made the following suggestions:

1. Tooltips with explanations of specialized terms.
2. Saving of at least the distribution layer of the map state to the URL so that one can link to specific distributions.
3. Tests for insecticide resistance against chemicals that are not used for LLINs can be hidden by default.
4. Data of distributions of other funders if it can be obtained.
5. The visualized dimension of the data (size of the map icon) can be explained by highlighting the respective item in the info box.
6. The info boxes of distributions can be improved by including photos.

5.1.1 Evaluation

The map that AMF currently still uses to indicate the location of distributions is fairly limited in features to the point where a comparison of features would seem unfair. It is clearly designed as an unobtrusive widget that is ancillary to a distribution page, a widget that adds some value to it but is by no means central to it.

Accordingly it is small in size and limited in interactivity. It displays the endemicity of malaria in various regions around the world as a binary, consistent with AMF's prioritization practice (see section 2.4), either applying a relatively low threshold for endemicity

or using data that has not been updated since the launch of the website.¹

Distributions are indicated by yellow dots. The user can click on continents to zoom in, which opens AMF's distribution list filtered by that continent.

the LLIN Visualizer improves upon this in many ways:

1. it adds all the data that comes with using OpenStreetMap,
2. allows a drill-down into the distributions down to village level or even the individual household if the user is authorized to see it,
3. allows users to filter distributions and villages (or *clusters* in general) according to various criteria,
4. can view distributions and villages in different modes, for example, as centroids or as convex hulls,
5. visualizes statistical data on the distributions and villages,
6. shows a wealth of other dimensions of the data that are not visualized in tabular info boxes,
7. shows aggregate statistics on malaria knowledge, malaria prevention measures, their usage, and the sources of nets for villages and distributions,
8. visualizes statistics (also provided by AMF) on the nets needed and the net gap of many relevant countries and allows users to filter them,
9. displays incidence in addition to endemicity (and both in higher detail) for *Plasmodium falciparum* and *Plasmodium vivax* (using data provided by IR Mapper),
10. draws on IR Mapper for data on insecticide resistance tests from the past decades, which it visualizes and allows the user to filter, and
11. provides guidance to the user with visual feedback, animations, and many info boxes that explain functions and specialized terms.

During my research I have found several other maps that are more or less centrally concerned with malaria. They have all been mentioned throughout this text and are linked from the LLIN Visualizer, since my software draws on some of this existing infrastructure or data. the LLIN Visualizer being concerned with the prioritization of net distributions, it is again hard to compare these solutions.

The main achievement of the Malaria Atlas Project² lies in the estimation of the spacial distribution of *Plasmodium falciparum* incidence, net coverage, indoor residual spraying coverage, and artemisinin-based combination therapy coverage using sparse data augmented with model-based extrapolations. The interactive map the project uses to showcase these data allows for filtering them by year. the LLIN Visualizer draws heavily on these data as described in section 3.2.4.

The same is true of the Insecticide Resistance Mapper.³ This app draws on some data from the Malaria Atlas Project to visualize insecticide resistance tests of the last sixty years,

1. According to IR Mapper, the entomological inoculation rate of Syria, Iran, and Iraq is negligible outside very small regions to the south of Iran, and the same applies to most of China.

2. Kessler and Tschuggnall, "Malaria Atlas Project."

3. IR Mapper, "IR Mapper."

including many for insecticides that have not been used for a long time. To exclude such chemicals, it allows for fine-grained filters over the test results. As detailed in section 3.2.4, the LLIN Visualizer draws its insights into insecticide resistance from this application.

5.2 Nonfunctional Optimization

The cluster layer. Intuitiveness was an important factor in the optimization process. The errors in the data that are discussed in section 3.2.2 lead to tradeoffs where the most intuitive approaches lead to more or less disruptively damaged data.

Another factor is the user experience of working with the map. Several thousand of map features displayed at the same time make the map feel sluggish even in modern browsers on most computers. Once canvas-based solutions become mature enough to be usable as orthogonal abstraction layers on top of the existing library code, this problem will fade into the background, but until then, some degree of windowing, lazy loading, and optimized structuring of the data will be necessary.

If windowing can be thought of as a horizontal limiting of the data that is displayed, then the optimization of the data structure is a vertical measure. This additional measure has become necessary because windowing alone was only in some cases effective in constraining the number of active map features while allowing full usability. For example, when users are interested in a distribution in the DRC, then only insecticide resistance tests in that country need to be displayed since the distribution they are interested in does not stray over country borders either. The same principle, however, does nothing to constrain the number of data points of the distribution itself that are displayed, so new categories for windowing are needed. The optimized data structure provides those.

Consider the two trees in figure 5.2. The first consists of three layers. A user that drills down to the last layer has to click twice to select windows of the lower layers. In total 19 map features are displayed. In the second case, the user has to click thrice, and only 12 map features are displayed in total.

This technique can backfire. Let us call the new layer the cluster layer, as it is called in the LLIN Visualizer. When there are too few clusters, they have hardly any effect; too many clusters also add overhead on the cluster layer that reduces the effectiveness of the optimization, but if the clusters are chosen to present the perfect balance between these extremes, they are likely to be unintuitive for the user. Furthermore, the aforementioned ambiguities and errors also meant that some of the clusters were, given the world knowledge that made them intuitive in general, highly unlikely. Hence I experimented with different types of intuitive clusters, measured the number of clusters, the average cluster

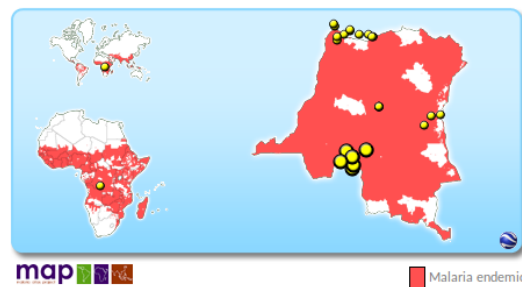


Figure 5.1: One of AMF's current maps.




Figure 5.2: A tree symbolizing the original structure of the data and one symbolizing the optimized structure. A user has drilled down to the lowest layer, so that the highlighted vertices are active.

size, and the proportion of them that seemed erroneous, according to heuristics such as the cluster elements' standard deviation from the geographic cluster centroid.

The results are only interesting for the data that I currently possess, the predistribution registration survey of the distribution Kasai Occidental 2014 in the DRC, and the LLIN Visualizer allows for choosing different types of clusters for different distributions. The winner was a clustering by tuples of normalized village name and health area name. Adding the sanitized health zone name into the tuple had little effect except on the runtime of the query. The resulting solution led to a total of 10,090 clusters with 1,261 clusters per distribution on average ($\sigma = 516$). Many clusters are errors due to misspellings, but with further sanitization steps applied, geared toward precision rather than recall, the reduced set of 676 clusters contains an average of 227 households ($\sigma = 210$).

Precomputation. I have already touched on the query optimizations I have performed in sections 4.1.1 and section 4.1.4. In summary: There are a number of calculations that would normally take several seconds to complete, so that they are unsuitable for execution at request time. Rather, the results are precomputed in the LLIN Visualizer, leading to runs of management commands that can take several minutes but queries on the precomputed

data that only take split seconds. This is one important optimization. The optimizations I was able to make to the long-running queries are described in section 4.1.4; they are mostly related to applying indices to relevant columns.

ORM-level optimizations. The quick queries that are executed at request time, however, are executed through the object-relational mapper interface that Django provides, so it pays to learn the optimization tips that the Django developers provide.

There are two situations where material gains can be realized. One relates to an object (a row in a table) that has a foreign key relationship to another object and is associated with the `select_related` method. By default, only the first object is fetched from the database, so that an additional query has to be executed to obtain the related object if and once it is needed. Two queries usually incur greatly more overhead than one more complex query, so `select_related` allows the developer, who knows what related objects are likely to be needed later on, to specify which of them should be fetched right away.

The other optimization relates to many-to-many relationships and reverse foreign key relationships and is associated with the `prefetch_related` method. From the developers perspective it behaves similar to `select_related`, but whereas `select_related` has the database perform a join operation, `prefetch_related` performs separate lookups for all the related objects and then performs the join in memory in Python.

Response caching. Another important improvement relates to the response caching of the LLIN Visualizer. Even though the precomputation and the query optimization above already reduced the response time significantly, I was able to speed it up further by caching the responses. With the LLIN Visualizer I am in an especially auspicious situation for caching as the data changes rarely. Hence I opted for no automatic expiry and explicit cache invalidation in addition to the unavoidable eviction due to memory restrictions of the server. (With the amount of data currently imported into the LLIN Visualizer, these are not reached.)

5.2.1 Evaluation

Clusters. Since I find “sluggishness” of a map hard to quantify and measure, I will use the number of map features as a proxy. Introducing clusters requires sanitization filters on this level with the useful side effect of further limiting the number of map features that are displayed by default. The numbers ignore unrelated map features, for example from insecticide resistance tests.

The initial view is what shows up before user interaction, and the drill-down is the state after the user clicks one distribution for which data is available. These measurements are averaged over all distributions for which this hold, a sample of eight at the moment.

	Without clusters	With clusters	Default filters
Initial view	634	–	117
Drill-down	29352 ($\sigma = 9853$)	1261 ($\sigma = 516$)	80 ($\sigma = 44$)

I am satisfied with these optimizations.

Precomputation. Response times are more straightforward than “sluggishness,” so I am measuring them directly. Here, the precomputation was an architectural decision I made early on, so the “without precomputation” measurements are those of queries I created purely for comparative purposes. The sample is again one of eight distributions.

	Without precomputation	With precomputation
Distributions	196.5 ms ($\sigma = 124.9$ ms)	1.5 ms ($\sigma = 1.4$ ms)
Clusters	298.0 ms ($\sigma = 159.6$ ms)	8.25 ms ($\sigma = 2.8$ ms)
Distribution statistics	1320.2 ms ($\sigma = 501.9$ ms)	5.14 ms ($\sigma = 8.8$ ms)
Cluster statistics	1676.1 ms ($\sigma = 434.5$ ms)	402.6 ms ($\sigma = 186.3$ ms)

Most of these optimizations had a great impact, except for the precomputation of the cluster statistics, since even the retrieval of the precomputed numbers requires one join. But since the infrastructure was already in place, the comparatively small optimization in this case also came at a small marginal cost.

ORM-level optimizations and caching. This measurement averaged over ten identical requests to the distributions endpoint and all cluster endpoints (one request each). The “with caching” column measures cache hits.

	Without optimizations	With optimizations	With caching
Distributions	3927 ms ($\sigma = 145$ ms)	1660 ms ($\sigma = 161$ ms)	1 ms ($\sigma = 0$ ms)
Clusters	12412 ms ($\sigma = 4832$ ms)	11986 ms ($\sigma = 4522$ ms)	1 ms ($\sigma = 0$ ms)

The impact of the query optimizations was small in this case, almost negligible in the case of the cluster queries, but the caching had a tremendous impact, making it indispensable. To grant the fast cached experience to as many users as possible, the LLIN Visualizer preloads the clusters when the browser is idle, so that they are cached on the server side and in the browser if they are not already there.

5.3 Learnings

5.3.1 Dates

AMF had already noticed that many of the dates were incorrect. When implementing the cleaning jobs, I observed the errors more closely to form theories of their origin. The dates of the rows are broken in so many ways that it seems likely that several problems coincided here to produce them. Two possible reasons are that (1) phones used in the distribution had some default date set at the various times when they were first activated instead of the real one or (2) data was lost on various data conversions.

Only looking at responses from the first days of distributions, about 47,000 (22%) are dated 2013, a year early. Almost 76,000 (36%) of this subset were dated 1899. About 83,000 (40%) were in the correct year. In total 56 different years were given. Less than 1% of the samples fell into each of all other years.

The year 1899 seems like an unlikely default for a phone, so I think it is more likely that here data was lost during a conversion. The countless different and often ambiguous date formats make these notoriously hard. The dates that are only off by a year may well be factory default dates of the phones. Differences of a few months between the erroneous dates can be explained by people first activating the phones at different times.

The conversion problems are particularly likely because instead of keeping and working with the raw data, the responses different teams collected were aggregated and then only saved in a format specific to the program “Microsoft Excel.” It has become possible to read the format without this software, but conversions to and from such a format are virtually bound to lead to data loss.

AMF has taken care to alert the distribution partner to this problem so that it can take greater care to set the correct date and time on the phones and so that it will avoid lossy conversions of the data going forward.

5.3.2 Health Areas

Experimenting with different ways of clustering the responses into geographical units (the process described in section 5.2), I noticed something strange about the level of the health area.

One of my options was clustering by the name of the health area. Looking at the data in QGIS, I noticed that the health area names seemed to be ambiguous, so I used the next larger unit, the health zone, to disambiguate. Even that, however, was unsuccessful in some cases. As the screenshot from the LLIN Visualizer shows, I eventually resorted to clustering by village so that I could use the combination of all three variables, health zone, health area, and village, for disambiguation. (The current implementation uses a simplified version of this approach that provides better performance.)

There are over 10,000 distinguishable villages in the data set, so it is unsurprising that many villages share the same name. The health areas, however, seem like centrally ordained administrative regions, and there are only about 200 of them, so I was surprised that they were ambiguous even within the next greater unit, the health zone.

I reported these findings to AMF who investigated the phenomenon and eventually formed the hypothesis that because the distributions in the different health zones happened consecutively, their distribution partner must have used the same smartphones for them. The health area is a property that is set once so that it does not have to be re-entered for every household. Hence some distribution volunteers must have forgotten to reset the property from a previous phase of the distribution leading to erroneous meta data, which becomes evident whenever there are villages with the same name in health areas that, erroneously, received the same identifier.




Figure 5.3: The village Nsoka, Muladila, has households that are over 170 km away from its centroid, an error in the data.

I suggested that I would be able to fix the data given shape files of the actual health areas, but none of us had access to any such files, and I have not been able to find them online, just as I have only been able to find very minimal data on the names and locations of villages in the DRC.

On usability grounds, to avoid that the convex hulls of some villages cover up other villages, I chose a cutoff at 35 km for the maximum distance of any household from the cluster centroid, the most parsimonious filter that allowed almost all cluster hulls to be immediately visible for the user. Centroidal Voronoi tessellation (CVT) came to mind too, but as the problem only concerned very few clusters, CVT would have destroyed much information of the convex hulls of all the correct clusters. For example it is interesting to derive from the data an idea of how densely or sparsely populated some regions are.

Finally, I considered a geographic hierarchical clustering⁴ to split up the erroneous clusters, but given the pragmatic cutoff, the problem only concerned 68 out of the over 10,000 clusters, and I expect for future distributions the problem to be resolved on the side of the distribution partner, so that this particular cleaning step would only be of very temporary value.

4. *k*-means clustering would be less apposite since it would take much experimentation to find a suitable number of clusters for the algorithm.

“Je n’ai fait celle-ci plus longue que parce que je n’ai pas eu le loisir de la faire plus courte.” (I made this [letter] very long only because I have not had the leisure to make it shorter.)

Blaise Pascal

6

Appendices

6.1 IMA 2014 Predistribution Registration Survey

Field identifiers, questions, and answers with English translations in parentheses.

1. **Identifiant** identification:GPS_HH:Latitude (identification.GPS.latitude),
identification:GPS_HH:Longitude (identification.GPS.longitude),
identification:GPS_HH:Altitude (identification.GPS.altitude),
identification:GPS_HH:Accuracy (identification.GPS.accuracy)
Question “Activer votre GPS pour chaque ménage” (“Enable GPS for each household”)
2. **Identifiant** identification:date_heure (identification.dateTime)
Question “Entrez la date et l’heure” (“Enter the date and time”)
3. **Identifiant** identification:phone_id (identification.phoneID)
Question “Entrez votre telephone identification (N° Reco)” (“Enter your phone identification (ID Reco)”)
4. **Identifiant** identification:aire_de_sante_province
(identification.healthArea.province)
Question Not given.
5. **Identifiant** identification:aire_de_sante_district (identification.healthArea.district)
Question Not given.
6. **Identifiant** identification:aire_de_sante_hz (identification.healthArea.healthZone)
Question Not given.
7. **Identifiant** identification:aire_de_sante (identification.healthArea.type)

- Question** “Selectionner votre Aire de Santé” (“Select your Health Area”).
8. **Identifiant** identification:nom_village (identification.village)
- Question** “Tapez le nom du village” (“Type the name of the village”).
9. **Identifiant** identification:LQAS_enum:enumerer_jour
(identification.dayOfDistribution)
- Question** “Quel jour de la distribution est-ce?” (“Which day of the distribution is it?”)
- Answers** (1) “Jour 1” (“Day 1”), (2) “Jour 2” (“Day 2”), (3) “Jour 3” (“Day 3”), (4) “Jour 4” (“Day 4”), (5) “Jour 5” (“Day 5”)
10. **Identifiant** identification:LQAS_enum:enumerer_HH
(identification.houseNumber)
- Question** “N° Ménage” (“House number”)
11. **Identifiant** identification:nomdechef_HH (identification.name)
- Question** “Ecrire le nom de chef de ménage” (“Write the head of household name”)
12. **Identifiant** identification:telephone_HH (identification.phone)
- Question** “N° téléphone Chef ménage” (“Phone number of head of household”)
13. **Identifiant** identification:mentions_legales (identification.legalScript)
- Question** “Bonjour. Votre participation à l’enquête de distribution est souhaitée. Le but de l’enquête est de recueillir les informations démographiques et déterminer si oui ou non vous avez besoin de moustiquaires. Nous vous demandons de répondre à quelques questions et entrer dans votre maison pour installer la MILD. Notre entretien prendra environ 10 minutes. Si vous souhaitez que votre image soit posté en ligne, nous pouvons prendre une photo avec la MILD; si non on prendra l’image de la MILD sans vous. Êtes-vous d’accord? Si oui, signez.” (“Hello. Your participation in the survey distribution is desired. The purpose of the survey is to collect demographic information and determine whether or not you need mosquito nets. We ask you to answer a few questions and enter your home to install the LLINS. Our interview will take about 10 minutes. If you want your image to be posted online, we can take a picture with the LLINS; if not, we will take the image of the MILD without you. Do you agree? If so, sign.”)
- Answers** (1) “Oui” (“Yes”), (2) “Non” (“No”)
14. **Identifiant** identification:signature_of_chef (identification.signature)
- Question** “Si vous acceptez de participer à l’enquête, s’il vous plaît signer ici:” (“If you agree to participate in the survey, please sign here:”)
15. **Identifiant** legales_non:socio_demographique:age_repondant
(sociodemographic.age)
- Question** “Age du répondant (ou donné complète)” (“Age of respondent (or given complete)”)

16. **Identifieur** legales_non:socio_demographique:sexe_repondant
(sociodemographic.sex)
Question “Sexe du repondant” (“Sex of respondent”)
Answers (1) “Male,” (2) “Female”
17. **Identifieur** legales_non:socio_demographique:total_personnes
(sociodemographic.peopleInHousehold)
Question “Nombre de personne dans le ménage” (“Number of people in household”)
18. **Identifieur** legales_non:socio_demographique:nbr_femmes_enceinte
(sociodemographic.pregnantWomen)
Question “Nombre de femmes enceinte” (“Number of pregnant women”)
19. **Identifieur** legales_non:socio_demographique:nbr_enfant_sous_5
(sociodemographic.childrenUnderFive)
Question “Nombre des enfants en dessous de 5 ans” (“Number of children under 5 years”)
20. **Identifieur** legales_non:connaissance_perception:malaria_transmise
(knowledge.malaria.transmission)
Question “Comment attrape-t-on la malaria?” (“How do you get malaria?”)
Answers (1) “en étant piqué par une moustique femelle infectée” (“by being bitten by an infected female mosquito”), (2) “en prenant un repas ou de l’eau sale” (“taking a self-catering or dirty water”), (3) “étant mouillé out frappé par soleil excessif” (“wet out being hit by excessive sun”), (4) “un environnement sale” (“dirty environment”), (5) “sorcellerie” (“witchcraft”), (6) “en mangeant les mangues ou les fruits verts” (“eating mangoes or green fruit”), (7) “en vivant avec les personnes malades” (“living with sick people”), (8) “autre (question suivante)” (“other (question)”), (9) “n’est pas connaitre / ignoré” (“does not know / ignored”)
21. **Identifieur** legales_non:connaissance_perception:malaria_transmise_autre
(knowledge.malaria.otherTransmission)
Question “Entrez autre raisons” (“Enter other reasons”)
22. **Identifieur** legales_non:connaissance_perception:malaria_signes
(knowledge.malaria.symptoms)
Question “Pourriez vous nous citez 4 signes / symptômes de la malaria?”
Answers (1) “fièvre / haute température” (“fever / high temperature”), (2) “maux de tête” (“headache”), (3) “douleurs au niveau de l’articulations / douleur général du corps” (“pain in the joints / general body pain”), (4) “fatigue” (“4) “tiredness”), (5) “vomissement” (“5) “vomiting”), (6) “manqué d’appétit / goût amère dans la bouche” (“6) “lacked appetite / bitter taste in the mouth”), (7) “diarrhée / douleur général” (“diarrhea / General pain”), (8) “autre (question

- suivante) (“other (question)”), (9) “n’est pas connaitre / ignoré” (“does not know / ignored”)
23. **Identifieur** legales_non:connaissance_perception:malaria_signes_autre
(knowledge.malaria.otherSymptoms)
Question “Entrez autre raisons” (“Enter other reasons”)
24. **Identifieur** legales_non:connaissance_perception:malaria_prevenir
(knowledge.malaria.prevention)
Question “Que pensez vous qu’on doit faire pour se prévenir de la malaria”
 (“What do you think we should do to prevent malaria”)
Answers (1) “creuser les canalisation d’eau stagnante” (“digging stagnant water pipeline”), (2) “nettoyer les herbes grandissante” (“clean growing herbs”), (3) “dormir sous Moustiquaire imprégnée d’insecticide” (“sleep under mosquito net impregnated with insecticide”), (4) “utilisez les contraceptif contre le moustique” (“contraceptive use against mosquitoes”), (5) “bruler les herbes” (“burning herbs”), (6) “prendre les médicaments anti malaria” (“take anti malaria drugs”), (7) “utiliser les insecticides” (“use insecticides”), (8) “autre (question suivante)” (“other (question)”), (9) “n’est pas connaitre / ignoré” (“does not know / ignored”)
25. **Identifieur** legales_non:connaissance_perception:malaria_traitement_simple
(knowledge.malaria.treatment)
Question “Quel est le médicament recommandé pour le traitement de la malaria?” (“What is the recommended drug for the treatment of malaria?”)
Answers (1) “SP/Fansidar” (“SP/Fansidar”), (2) “chloroquine” (“chloroquine”), (3) “amodiaquine/camoquine” (“amodiaquine/camoquine”), (4) “quinine” (“quinine”), (5) “artémisinine-base sur combinaison” (“artemisinin-based combination”), (6) “paracetamol” (“paracetamol”), (7) “aspirine/ibuprofène” (“aspirin/ibuprofen”), (8) “autre” (“other”), (9) “Je ne sais pas” (“I do not know”)
26. **Identifieur** legales_non:connaissance_perception:malaria_traitement_autre
(knowledge.malaria.otherTreatment)
Question “Entrez autre traitement” (“Enter another treatment”)
27. **Identifieur** legales_non:prevention:moins_une_MILD_HH
(prevention.atLeastOneNet)
Question “Est ce que votre ménage a au moins une Moustiquaire” (“Does your household have at least one mosquito net”)
Answers (1) “Oui” (“Yes”), (2) “Non” (“No”)
28. **Identifieur** legales_non:prevention:moins_impregnee_HH
(prevention.atLeastOneLLIN)
Question “Y-a-t-il au moins une Moustiquaires imprégnée d’insecticide dans votre ménage” (“Do you have at least one insecticide-treated mosquito net in your household?”)

- Answers** (1) “Oui” (“Yes”), (2) “Non” (“No”)
29. **Identifieur** legales_non:prevention:nbr_MILD_impregnee (prevention.LLINS)
Question “Combien de Moustiquaires imprégée avez vous dans votre ménage”
 (“How many LLINS do you have in your household?”)
30. **Identifieur** legales_non:prevention:obtenu_MILD_ou (prevention.netSource)
Question “Où avez-vous obtenu votre Moustiquaire?” (“Where did you get your mosquito net?”)
Answers (1) “marché/magasin/mharmacie” (“market/store/pharmacy”), (2) “ONG” (“NGO”), (3) “campagne de distribution” (“distribution campaign”), (4) “centre de santé” (“health center”), (5) “Autres” (“other”)
31. **Identifieur** legales_non:prevention:obtenu_MILD_ou_autre (prevention.otherNetSource)
Question “Entrez autre” (“Enter other”)
32. **Identifieur** legales_non:prevention:tout_personnes_dormi (prevention.peopleUnderNetLastNight)
Question “Combien des personnes dans votre maison ont dormi sous la Moustiquaires la nuit précédente?” (“How many people in your house slept under the nets last night?”)
33. **Identifieur** legales_non:prevention:enceintes_dormi (prevention.pregnantWomenUnderNetLastNight)
Question “Combien de femmes enceintes ont dormi sous la moustiquaire la nuit dernière?” (“How many pregnant women slept under the net last night?”)
34. **Identifieur** legales_non:prevention:IPT_utilisez (prevention.fansidarDuringPregnancy)
Question “Combien de fois a-t-elle pris le Fansidar (TPI) pendant sa grossesse pour prévenir la malaria?” (“How many times did she take Fansidar (IPT) during her pregnancy to prevent malaria?”)
35. **Identifieur** legales_non:prevention:enfants_dormi (prevention.allChildrenUnderNetLastNight)
Question “Est-ce que tous les enfants de moins de 5 ans ont dormi sous la moustiquaire la nuit dernière?” (“Did all children under 5 years of age sleep under the net last night?”)
36. **Identifieur** legales_non:prevention:prevention_dans_HH (prevention.preventiveMeasures)
Question “Quelles sont les mesures de prevention faites vous pour se prevenir de la malaria dans votre menage?” (“What preventive measures do you do to prevent malaria in your household?”)
Answers (1) “avoir enlevé l’eau stagnante” (“removing standing water”), (2) “utiliser la lotion anti moustique” (“Use mosquito repellent lotion”), (3) “pulverisation des maisons” (“spraying houses”), (4) “faucher les herbes”

- (“mowing grass”), (5) “nettoyer l’environnement” (“Clean Environment”), (6) “couvrir les recipients d’eau” (“cover water containers”), (7) “bruler des feuilles, fumier, etc.” (“burning leaves, manure, etc.”), (8) “rein” (“kidney”), (9) “n’est pas connaitre / ignoré” (“does not know / ignored”), (10) “autre (question suivante)” (“other (question)”)
37. **Identifieur** legales_non:prevention:prevention_dans_HH_autre (prevention.otherPreventiveMeasures)
Question “Entrer d’autres méthodes” (“Enter other methods”)
38. **Identifieur** legales_non:traitement:enfant_fievre (treatment.child.feverLastTwoWeeks)
Question “Dans le deux semaines y-a-t-il un enfant en dessous de 5 ans qui a souffert de la fièvre?” (“In the two weeks there-he has a child under 5 years who had fever?”)
Answers (1) “Oui” (“Yes”), (2) “Non” (“No”)
39. **Identifieur** legales_non:traitement:enfant_beneficier (treatment.child.benefitFromDrug)
Question “L’enfant a-t-il beneficier du medicament contre la fievre ou la malaria?” (“Does the child benefit from the drug against fever or malaria?”)
Answers (1) “Oui” (“Yes”), (2) “Non” (“No”)
40. **Identifieur** legales_non:traitement:enfant_medicament (treatment.child.medicine)
Question “Quelle sorte de medicament que l’enfant a pris?” (“What kind of medicine that the child has taken?”)
Answers (1) “SP/Fansidar” (“SP/Fansidar”), (2) “chloroquine” (“chloroquine”), (3) “amodiaquine/camoquine” (“amodiaquine/camoquine”), (4) “quinine” (“quinine”), (5) “artemisinin-base sur combinaison” (“artemisinin-based combination”), (6) paracetamol” (“paracetamol”), (7) “aspirine/ibuprofen” (“aspirin/ibuprofen”), (8) “autre” (“other”), (9) “Je ne sais pas” (“I do not know”)
41. **Identifieur** legales_non:traitement:consulte_enfant (treatment.child.healthcenter)
Question “A-t-on consulté l’enfant au centre de santé pendant sa maladie.” (“Was the child taken to a consultation at the health center during his illness.”)
Answers (1) “Oui” (“Yes”), (2) “Non” (“No”), (3) “Je ne sais pas” (“I do not know”)
42. **Identifieur** legales_non:distribution:place_dormir (distribution.sleepingSpaces)
Question “Nombre de place à dormir” (“Number of sleeping spaces”)
43. **Identifieur** legales_non:distribution:nbr_bonne_mild (distribution.LLINS.good)
Question “Combien de MILD de bonne qualité sont disponibles dans le ménage” (“How good LLINS are available in the household”)
44. **Identifieur** legales_non:distribution:MILD_par_HH (distribution.LLINS.total)
45. **Identifieur** legales_non:distribution:distrbution_summary (distribution.distributionSummary)

46. **Identifiant** legales_non:distribution:MILD_installees (distribution.LLINS.installed)
Question “Entrez Nombre de MILD installées” (“Enter Number of LLINS installed”)
47. **Identifiant** legales_non:distribution:MILD_marque (distribution.LLINS.brand)
Question “Choisissez la marque de moustiquaires installées” (“Choose the brand of nets installed”)
Answers (1) Olyset, (2) PermaNet
48. **Identifiant** legales_non:distribution:MILD_retournees (distribution.LLINS.returned)
Question “Entrez Nombre de MILD retournées” (“Enter number of LLINS returned”)
49. **Identifiant** legales_non:distribution:MILD_emballage (distribution.LLINS.packagingReturned)
Question “Entrez Nombre d’emballage” (“Enter Number packaging collected”)
50. **Identifiant** legales_non:distribution:photo_installee (distribution.LLINS.photo)
Question “Entrez la photo de MILD installées” (“Enter the photo of LLINS installed”)
51. **Identifiant** start (start)
Description Start time
52. **Identifiant** end (end)
Description End time
53. **Identifiant** deviceid (deviceID)
Description Phone serial number
54. **Identifiant** meta:instanceID (meta.instanceID)

6.2 Concept for Donor Coordination

6.2.1 Donor Coordination

When we focus on just those individuals and organizations that would constitute the “GiveWell market” discussed in section 2.2.2, people centrally interested in making their donations and contributions of time maximally impactful, then new coordination problems emerge.

One problem that GiveWell has struggled with emerges when two donors are not fully value-aligned but can agree on wanting to fund one of GiveWell’s top charities. The result is that they wait each other out, a deadlock, both wanting the other to fund the GiveWell charity, because they value the other’s counterfactual use of the donation lower than their own.¹ GiveWell is regularly refining its response to this problem.

1. GiveWell, “Donor Coordination and the ‘Giver’s Dilemma’.”

For this problem to become relevant, there need to be at least two large donors or monolithic groups of donors, where large means that their planned donations are close to – for example within the same order of magnitude of – the funding gap of the charities in question. This is a good problem to have.

More commonly, however, the funding gap is large compared to the potential individual donations (where *individual* is meant to exclude the aforementioned monolithic groups of donors), so that the above problem becomes an edge case while centrally we face a different problem. Donors that focus their contributions on charities that have a significant evidence base and track record for impact – a large part of the “GiveWell market” – are often accused of being too focused on just these established charities thereby missing small high-impact opportunities from nonprofit startups or projects that will stay small or short-lived by design.

The distinction is similar to that between, on the one hand, passive investors that buy exchange-traded funds (ETFs) of, for example, the top 30 (DAX) or 500+ (S&P 500) companies in order to hold them, and, on the other hand, business angels or venture capitalists that invest into startups. The first group has excellent information to make relatively low risk–low return investments; the second group has to rely on rough heuristics, such as their faith into the founders, to make high risk–high return investments – of which they need to be able to make many in order to profit at least fairly reliably.

But a profit motive is an agent-relative goal. Investors (such as donors) with agent-neutral goals that are shared by at least a few others have much better opportunities for cooperation. These have largely not been tapped into. While the LLIN Visualizer is clearly focused on the low risk–low return market, this high risk–high return market also calls for a software solution to its coordination problem.

The central motivating problems are the following:

Irrational risk-aversion. Drops in marginal utility of a resource suggest risk aversion. In that context it is rational to prefer a low return with high probability to a high return with low probability at the same expected value.² In the context of altruistic interventions,³ the utility of marginal donations only noticeably decreases when it reaches the area of millions of dollars, some of the GiveWell top charities’ funding gaps. Since few donors have funds of that magnitude at their disposal, most risk aversion of average donors is disproportionate.

At the same time there are many donors that see a high likelihood that effective interventions are possible in a certain cause area. Unfortunately, these interventions are, by necessity, more speculative than, for example, the interventions GiveWell prioritizes. Yet there are charity startups implementing them.

The funding gaps of these charities tend to be too small for any respectable prioritization organization, like GiveWell, to warrant investing staff time into evaluating them, so donors

2. SDG International, “Incoming Program Officer: Lewis Bollard.”

3. Please note that in the following I will use “intervention” and “program” semantically interchangeably conditional on which terms seems more idiomatic to me in the collocational context.

Competition for exposure. The other side of the medal is that charities are aware of these dynamics. While their values may be aligned, for funding they are yet each dependent on its own pool of donors, and any cross-promotion of another charity among the first charity's own donor base may lead to donors shifting their support to the endorsed organization. This behavior stifles cooperation.

The solution presented here will instead allow all charities in a program area to fill their funding gaps to similar degrees. If a sufficient number of donors come to accept this solution, any incentive for charities to engage in uncooperative behavior will be diminished.

Donor Coordination (working title) is a software system and strategy that fosters cooperation between value-aligned donors by allowing them to make large contributions in teams and donate to whole program areas rather than individual nonprofits. It can improve upon the current state if it is accepted and trusted by a sufficient number of donors.

6.2.2 High-Level Goals

The donor coordination solution can be considered successful when it achieves the following goals:

Team-level atomicity. Donors can choose portfolios with whom they are value-aligned to the point that they perceive their donations as coming from the team of donors that invests into that portfolio rather than them personally.

Program-level atomicity.

1. Donors can choose charity portfolios that, as a whole, represent their moral preferences well enough that they perceive their team as donating to a program area rather than an individual charity.
2. Charities are value-aligned with the organization their donations are fungible with to the point that they make fully altruistic statements about their funding gaps.

6.2.3 Core Concepts

The donor coordination solution should likely take the shape of a web application to enable users of any platform to use it. The idea is roughly inspired by Wikifolio.

Visitor. An unauthenticated person viewing the website.

User. A donor, a charity, or an administrator.

Portfolio. The portfolio is an allocation rule that partitions funds among a set of charities. Every user can create portfolios, favorite or watch portfolios, and donate to portfolios.

Donor. The donor is a user other than a charity.

Charity. A charity is a user that only has the ability to enter some meta data about itself and its funding gap, and participate in discussions.

6.2.4 Interest Groups

The interests of beneficiaries are:

1. Beneficiaries want to maximize the available funding toward their preferences.
2. Beneficiaries want the, at the margin, most effective interventions to receive maximal funding.
3. Beneficiaries want the funding gaps of the most effective interventions to be greater than or equal to the available funding.

In some cases the beneficiaries can give direct input, but in many cases their interests need to be represented by donors and charities because they have insufficient levels of intelligence to express them efficiently or are not yet born.

Hence the interests of donors are:

1. Donors want to maximize the available funding toward their moral goals.
2. Donors want the, at the margin, most effective interventions realizing these moral goals to receive maximal funding.
3. Donors want the funding gaps of the most effective interventions realizing these moral goals to be greater than or equal to the available funding.

Hence the interests of the charities are:

1. Charities want to maximize the available funding toward the charity's moral goals.
2. Charities want the, at the margin, most effective interventions realizing these moral goals to receive maximal funding.
3. Charities want the funding gaps of the most effective interventions realizing these moral goals to be greater than or equal to the available funding.

The main difference between donors and charities as two groups is the direction of the money flow. The main difference between the donors and charities internally is their different moral goal makeup.

From these primary interests follow proximate interests for value-aligned teams of donors (all donors to a program area as defined by a public portfolio):

1. Being value aligned, the members of a team are happy to make their donations fungible with the donations of all other members of the team.
2. Since their value alignment with other teams varies, there may be teams with partially opposing moral goals. Teams will want to minimize fungibility with such teams.
3. Since the funding gaps of charities are limited, teams also want to increase the funding gap of their program area by broadening its scope.

Analogously for charities:

1. The charities of popular portfolios are likely to be highly value aligned and thus happy to calculate their funding gaps cooperatively.
2. Since their value alignment with charities of other program areas varies, there may be portfolios of charities with partially opposing moral goals. Charities will want to increase their scale in order to be able to enter greater funding gaps so portfolio authors can minimize fungibility with such opposing program areas.

Clearly, the last two interests of the donor teams are in conflict. Small donation flows will favor portfolios of small, pure clusters of charities while greater donation flows will necessitate compromise in order to form greater, less pure clusters with larger funding gaps.

For simplicity I assume that all donors are perfectly informed and their only differences are differences of value alignment. This is unlikely to be the case in practice, but the only difference between a donor that is not value aligned and a donor that acts as if they were not value aligned because of lacking information is that the latter can be educated.

This educational mission is without the purview of Donor Coordination, but the software should provide the platform that donors will need to educate each other because this may be important for fostering user activity.

6.2.5 Functional Requirements

1. Visitors can create donor accounts.
2. Administrators can create administrator accounts.
3. Administrators can create charity accounts.
4. Visitors can view public portfolios including their descriptive statistics.
5. Donors can add public portfolios to their watch list.
6. Donors can donate to public portfolios.
7. Donors can author public portfolios.
8. Donors can draft and test portfolios in a private or draft state.
9. Donors can comment on portfolios.
10. Charities can enter new funding gaps for themselves.
11. Charities can enter new system-external donation flows.
12. Administrators have all privileges.

At some point moderator accounts will become necessary, so moderators do not need to enjoy the same level of trust as administrators to contribute to the community maintenance.

6.2.6 Challenges and Proposed Solutions

Descriptive statistics. The functional requirements mention descriptive statistics. These are important for portfolio authors and other donors to decide how to structure a portfolio so not to duplicate very similar ones or which portfolio to donate to. At least two metrics are required:

1. The sum of the funding gaps of the charities in a portfolio P , $\text{gap}(P) = \sum_c^P \text{gap}(c)$.
2. A ranked list of the portfolios with the highest fungibility but lowest similarity. One idea may be the quotient, $\text{compromise}(P, P') = \frac{\text{fungibility}(P, P')}{\text{similarity}(P, P')}$, of the following metrics:

$$\text{a) fungibility}(P, P') = \sum_c^{P \cap P'} \text{gap}(c)$$

$$\text{b) similarity}(P, P') = \left| \bigcup_c^P \text{donors}(c) \cap \bigcup_c^{P'} \text{donors}(c) \right|.$$

The fungibility and similarity metrics should also be displayed in isolation, particularly as a guide for authors of portfolios of new charities when the combined compromise metric is undefined.

It may also become necessary to take weights into account, and the formulas will surely need to be tweaked further once real data become available.

Funding gaps. There needs to be a common definition of a funding gap, so that charities have hard, unyielding guidelines as to what figure to enter for a given year.

Prioritization organizations already face a similar problem: Imagine two charities, charity A with the ability to invest \$100 million with some baseline effectiveness e on average and charity B with the ability to invest \$10 million with an average effectiveness of $10e$ within a given year. Further assume that the charities are value aligned to simplify the problem to one dimension of impact.

A commonly used uncertainty discount is 3% p.a.⁵ and for simplicity we assume that suffering in the world, absent the interventions, remains constant, so that aggregate suffering increases linearly over time.

A donor that wants to invest \$100 million now has the choice to donate it to charity A, knowing that it will be invested in the same year, or to charity B, knowing that \$10 million of it will be invested in the same year, \$90 million of it will wait on the charity's bank account at an interest rate of maybe 1% for another year, \$80 million plus interest will wait for two years, and so on.

Clearly, a definition of funding gaps that only takes into account a charity's ability to invest some amount per year would set very different bars for the marginal impact of the last dollar of that funding gap.

Since the donor coordination solution addresses coordination problems that arise when the funding gaps of the individual charities do not warrant the attention of a prioritization organization, we assume that there are no meaningful differences of their relative effectiveness, so we face a simpler version of this problem.

One solution may be to adopt GiveWell's excess assets policy: "We seek to be in a financial position such that our cash flow projections show us having 12 months' worth of unrestricted assets in each of the next 12 months."⁶

5. Lopez, *Global Burden of Disease and Risk Factors*, 399–400.

6. GiveWell, "GiveWell's 'Excess Assets' Policy."

Allocation. Another open question is the allocation of donations within the portfolio. Conceptually, donors donate to program areas, but factually they will have to transfer their donation to a specific organization. Splitting it up across several organizations would be an unnecessary hassle of the donor, so the algorithm that suggests the specific organization should know some ideal allocation and then recommend a recipient organization such that the actual allocation comes closest to the ideal allocation. It could also take tax deductibility into account as a tie breaker.

The simplest option might be an equitable allocation where the algorithm aims to assign the same level of funding to each charity after taking donations external to the system into account.

Another option may be to prioritize small funding gaps as an additional incentive for charities not to exaggerate their funding gaps in the *moral lies* scenario. However that would have little effect since the charities in a given program area are value aligned and can thus easily conspire with each other, and it may have the detrimental effect that charities would be incentivized to be tardy with entering new funding gaps.

Donation swaps. Donors often agree on donation swaps where each partner donates to the charity of choice of the other partner in order to harness the tax deduction of the charity in the respective country.

In order to help portfolio authors to trade off fungibility against funding gaps, there would need to be a ranking of other portfolios that the given portfolio is most fungible with. However, portfolios whose audience is very similar are least interesting to portfolio authors, so the ranking should be sorted by something like the fungibility per cardinality of the cut set of donors, and here donation swaps would add noise to the calculation.

It needs to be either clear to the donors that they need to enter the donation of their swap partner as their donation or the software should allow them to mark donations as swaps and enter their partner. The first is probably the better solution for an MVP, but the second may be more foolproof.

6.2.7 Remaining Challenges

Moral lies. When there is a pair of program areas such that the teams of each see the team of the other as an opposing team, but there is some set of charities that they can agree on, and the available funding is close to or greater than the available funding gaps of their program areas without the consensus charities, the intended result is that donors compromise and add charities to their portfolios that increase the funding gap at the cost of greater fungibility.

But charities are of course value aligned with these teams. Hence it will be ethical for them to lie about their funding gaps, inflating them, to drive the opposing donors to fund the more fungible funding gaps. Analogously, the opposing team's charities can also inflate their funding gaps; they even have to lest their cause suffer. When one group defects in such a fashion, the cooperation breaks down. A classical example of the prisoner's dilemma.

In practice, the donor coordination solution will be used mostly or at least at first only by donors that are all fairly value aligned at least to the extent that they value the type of moral plurality that exists among them. Hence this problem may not manifest any time soon.

Market research. My experience that such a software would be helpful is based on reports of friends, some of whom are donors and some employees of affected charities. Unless, however, there is a sizable number of prospective users that are interested in the project, charities will not have sufficient faith into the growth of the user base to warrant their time investments.

Apart from surveys among likely prospective users, one central market research tool needs to be a minimal viable product (MVP). Other donors and nonprofit staff have considered opening a group on a social network such as Facebook to bring together all participants in whose actions need coordination. The group would provide a means for communication but would leave any functions beyond that to the participants to be implemented in a manual, ad-hoc fashion. This way it will become clear which processes are in most urgent need of automation. It will also become clear if the community is large enough to sustain a more comprehensive solution like the one proposed here.

Community building. An important strategic and marketing problem is the following: Entering funding gaps will only warrant the effort for the charities if they can expect significant donation flows from the donor coordination solution. For donors the donor coordination solution is only interesting when the program areas they want to donate to are well represented by charities working on them.

One solution may be for administrators to regularly poll information on funding gaps from charities and invite them to claim their accounts themselves. That way, the administrators will have added effort during the startup phase, which will be increasingly outsourced to the charities as donors come to accept the system.

To achieve said donor acceptance, it would be helpful if the project were run by a reputable organization with considerable reach, and if the project collected early signups prior to its launch, both in order for it to launch with momentum. Until such an organization has been found, I cannot consider this challenge solved.

Bibliography

- Adelman, Carol, Yulya Spantchak, Kacie Marano, Jeremiah Norris, Jesse Barnett, Haowen Chen, Darice Xue, and Aubrey Thrane. "The Index of Global Philanthropy and Remittances, with a Special Report on Emerging Economies." *Hudson Institute Report, November* (2013).
- Agafonkin, Vladimir. "Leaflet: A JavaScript Library for Interactive Maps." 2015. Accessed January 12, 2016. <http://leafletjs.com/>.
- Against Malaria Foundation. "About Mosquito Nets." 2015. Accessed December 30, 2015. https://www.againstmalaria.com/FAQ_Bednets.aspx.
- . "AMF Update – 28 November 2013." 2013. Accessed January 29, 2016. https://www.againstmalaria.com/Downloads/AMF_Update_28Nov2013.pdf.
- . "Question to AMF: 'Where are Your Nets Manufactured?'" 2015. Accessed December 30, 2015. <https://www.againstmalaria.com/Newsitem.aspx?NewsItem=Question-to-AMF-Where-are-your-nets-manufactured>.
- Animal Charity Evaluators. "Impact of Donations." 2015. Accessed September 6, 2015. <http://www.animalcharityevaluators.org/research/foundational-research/impact-donations/>.
- Apache Software Foundation. "Apache Subversion." 2015. Accessed January 12, 2016. <https://subversion.apache.org/>.
- Axelrod, Robert. *The Evolution of Cooperation: Revised Edition*. Basic Books, 2009. ISBN: 978-0-78-673488-7. <https://books.google.de/books?id=GxRo5hZtxkEC>.
- Banerjee, Abhijit, and Esther Duflo. *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*. PublicAffairs, 2012. ISBN: 978-1-61-039160-3. <https://books.google.de/books?id=2dlnBoX4licC>.
- BBC News. "Profile: Tito the Spaceman." 2001. Accessed December 5, 2015. <http://news.bbc.co.uk/2/hi/science/nature/1297924.stm>.
- Bernhardt, Gary. "Wat." 2012. Accessed March 7, 2016. <https://www.destroyallsoftware.com/talks/wat>.
- Bill & Melinda Gates Foundation. "New Initiative Will Assess Disease Control Priorities In Developing Countries." 2002. Accessed December 9, 2015. <http://www.gatesfoundation.org/Media-Center/Press-Releases/2002/09/Initiative-to-Assess-Disease-Control-Priorities-in-Developing-Countries>.
- Bostock, Mike. "C3.js: D3-Based Reusable Chart Library." 2016. Accessed March 12, 2016. <http://c3js.org/>.

- Bostock, Mike. "D3.js: Data-Driven Documents." 2016. Accessed March 12, 2016. <https://d3js.org/>.
- Bostrom, Nick. "The Superintelligent Will: Motivation and Instrumental Rationality in Advanced Artificial Agents." *Minds and Machines* 22, no. 2 (2012): 71–85.
- Brandl, Georg, and the Sphinx team. "Sphinx." 2016. Accessed March 12, 2016. <http://www.sphinx-doc.org/>.
- CDC. "Malaria – About Malaria – Biology – Mosquitoes – Anopheles Mosquitoes." 2016. Accessed January 22, 2016. <http://www.cdc.gov/malaria/about/biology/mosquitoes/>.
- Central Intelligence Agency. "The World Factbook: Congo, Democratic Republic of the." 2015. Accessed January 3, 2016. <https://www.cia.gov/library/publications/the-world-factbook/geos/cg.html>.
- Chandy, Laurence, and Geoffrey Gertz. *Poverty in Numbers: The Changing State of Global Poverty from 2005 to 2015*. Brookings Institution, 2011.
- CherryPy. "CherryPy: A Minimalist Python Web Framework." 2015. Accessed January 12, 2016. <http://www.cherrypy.org/>.
- CoffeeScript. "CoffeeScript." 2016. Accessed March 7, 2016. <http://coffeescript.org/>.
- Crompton, Peter D., Susan K. Pierce, and Louis H. Miller. "Advances and Challenges in Malaria Vaccine Development." *The Journal of clinical investigation* 120, no. 12 (2010): 4168.
- Darley, John M., and Bibb Latane. "Bystander Intervention in Emergencies: Diffusion of Responsibility." *Journal of personality and social psychology* 8, no. 4p1 (1968): 377.
- Desvousges, William H, F Reed Johnson, Richard W Dunford, Kevin J Boyle, Sara P Hudson, K Nicole Wilson, et al. *Measuring Nonuse Damages Using Contingent Valuation: An Experimental Evaluation of Accuracy*. Research Triangle Park, NC: RTI Press, 2010.
- Dierig, Carsten. "Warum die Deutschen so gerne Veggie-Wurst essen." 2015. Accessed December 30, 2015. <http://www.welt.de/wirtschaft/article145694267/Warum-die-Deutschen-so-gerne-Veggie-Wurst-essen.html>.
- Diofasi, Anna, and Nancy Birdsall. "The World Bank's Poverty Statistics Lack Median Income Data, So We Filled In the Gap Ourselves: Download Available." 2016. Accessed July 8, 2016. <http://www.cgdev.org/blog/world-bank-poverty-statistics-lack-median-income-data-so-we-filled-gap-ourselves-download-available>.
- Django Software Foundation. "GeoDjango: A World-Class Geographic Web Framework." 2016. Accessed March 12, 2016. <http://geodjango.org/>.
- Drescher, Denis. "LLIN Visualizer Demo." 2016. Accessed August 6, 2016. <https://www.youtube.com/watch?v=S5K3aF21E4Y>.
- . "Make-A-Wish Dollars Per DALY Averted." 2016. Accessed March 5, 2016. <https://www.getguesstimate.com/models/4296>.
- . "Map – LLIN Visualizer." 2016. Accessed July 14, 2016. <http://nets.claviger.net/>.
- . "Mosquito YLL Per Human \neg YLL." 2016. Accessed March 7, 2016. <https://www.getguesstimate.com/models/4302>.

- Drescher, Denis. "Telofy / LLIN Visualizer – Bitbucket." 2016. Accessed July 14, 2016. <https://bitbucket.org/Telofy/llin-visualizer>.
- Ecma International. "ECMAScript." 2016. Accessed March 12, 2016. <http://www.ecmascript.org/>.
- Erlandsson, Arvid, Fredrik Björklund, and Martin Bäckström. "Emotional Reactions, Perceived Impact and Perceived Responsibility Mediate the Identifiable Victim Effect, Proportion Dominance Effect and In-Group Effect Respectively," *Organizational Behavior and Human Decision Processes*, 127, no. March (2015): 1–14. ISSN: 0749-5978. <http://dx.doi.org/10.1016/j.obhdp.2014.11.003>.
- Flask. "Flask: A Python Microframework." 2014. Accessed January 12, 2016. <http://flask.pocoo.org/>.
- Forcier, Jeff. "Fabric." 2016. Accessed March 12, 2016. <http://www.fabfile.org/>.
- Free Software Foundation. "CVS: Concurrent Versions System." 2015. Accessed January 12, 2016. <http://www.nongnu.org/cvs/>.
- Gallup. "How Does the Gallup World Poll Work?" 2016. Accessed July 9, 2016. <http://www.gallup.com/178667/gallup-world-poll-work.aspx>.
- . "Understanding How Gallup Uses the Cantril Scale." 2016. Accessed July 8, 2016. <http://www.gallup.com/poll/122453/understanding-gallup-uses-cantril-scale.aspx>.
- GiveWell. "A Conversation with Marcy Erskine on September 18, 2013." 2013. Accessed January 1, 2016. [http://files.givewell.org/files/conversations/Marcy%20Erskine%209-18-13%20\(public\).pdf](http://files.givewell.org/files/conversations/Marcy%20Erskine%209-18-13%20(public).pdf).
- . "A Conversation with Professor Christian Lengeler, August 20, 2015." 2015. Accessed February 7, 2016. [http://files.givewell.org/files/conversations/Christian_Lengeler_08-20-2015_\(public\).pdf](http://files.givewell.org/files/conversations/Christian_Lengeler_08-20-2015_(public).pdf).
- . "Against Malaria Foundation (AMF)." 2015. Accessed December 5, 2015. <http://www.givewell.org/international/top-charities/AMF>.
- . "Donor Coordination and the 'Giver's Dilemma'." 2014. Accessed January 1, 2016. <http://blog.givewell.org/2014/12/02/donor-coordination-and-the-givers-dilemma/>.
- . "Errors in DCP2 Cost-Effectiveness Estimate for Deworming." 2011. Accessed December 9, 2015. <http://blog.givewell.org/2011/09/29/errors-in-dcp2-cost-effectiveness-estimate-for-deworming/>.
- . "GiveWell's 'Excess Assets' Policy." 2015. Accessed January 31, 2016. <http://www.givewell.org/about/official-records/Excess-Assets-Policy>.
- . "Process for Identifying Top Charities." 2013. Accessed September 6, 2015. <http://www.givewell.org/international/process>.
- . "The Money for Good Study." 2010. Accessed January 1, 2016. <http://blog.givewell.org/2010/07/20/the-money-for-good-study/>.
- Giving USA Foundation. *Giving USA 2015: The Annual Report on Philanthropy for the Year 2014*. 2015.
- Giving What We Can. "How Rich Am I?" 2015. Accessed December 7, 2015. <https://web.archive.org/web/20150218010009/https://www.givingwhatwecan.org/get>

involved/how-rich-am-i.

- Goldsmith, Connie. *Battling Malaria: On the Front Lines against a Global Killer*. Nonfiction - Young Adult. Ebsco Publishing, 2010. ISBN: 978-0-76-136363-7. https://books.google.de/books?id=u07%5C_UJ78vYsC.
- Graham, Jesse, Jonathan Haidt, Sena Koleva, Matt Motyl, Ravi Iyer, Sean P Wojcik, and Peter H Ditto. "Moral Foundations Theory: The Pragmatic Validity of Moral Pluralism." *Advances in Experimental Social Psychology, Forthcoming* (2012).
- Gray, Kurt, Liane Young, and Adam Waytz. "Mind Perception Is the Essence of Morality." PMID: 22754268, *Psychological Inquiry* 23, no. 2 (2012): 101–124. doi:10.1080/1047840X.2012.651387. eprint: <http://dx.doi.org/10.1080/1047840X.2012.651387>. <http://dx.doi.org/10.1080/1047840X.2012.651387>.
- Hauck, Katharina, and Peter C. Smith. "The Politics of Priority Setting in Health: A Political Economy Perspective." *Center for Global Development Working Paper*, no. 414 (2015).
- Haushofer, Johannes, James Reisinger, and Jeremy Shapiro. "Your Gain Is My Pain: Negative Psychological Externalities of Cash Transfers" (2015).
- Hillebrandt, Hauke. "Median GDP per capita: how much does the typical person earn in different countries?: A look at Global Inequality." 2016. Accessed July 8, 2016. <https://www.givingwhatwecan.org/post/2016/05/giving-and-global-inequality/>.
- Hope Consulting. "Money For Good II" (2011). http://www.hopeconsulting.us/wordpress/wp-content/uploads/2013/04/MoneyForGood_II_Full1.pdf.
- Hudson, James M., and Amy S. Bruckman. "The Bystander Effect: A Lens for Understanding Patterns of Participation." *The Journal of the Learning Sciences* 13, no. 2 (2004): 165–195.
- Impact Purchase. "The Impact Purchase: 2015 Impact Purchase." 2015. Accessed January 1, 2016. <http://impactpurchase.org/>.
- Institute for Health Metrics and Evaluation. "Global Burden of Disease Compare – IHME Viz Hub." 2013. Accessed December 9, 2015. <http://vizhub.healthdata.org/gbd-compare/>.
- IR Mapper. "IR Mapper: Mapping Insecticide Resistance in Malaria Vectors." 2016. Accessed January 31, 2016. <http://www.irmapper.com/>.
- Jamison, Dean T., Joel G. Breman, Anthony R. Measham, George Alleyne, Mariam Clae-son, David B. Evans, Prabhat Jha, Anne Mills, and Philip Musgrove. *Disease Control Priorities in Developing Countries*. 2nd ed. World Bank Publications, 2006.
- Jamison, Dean T., W. Henry Mosley, Anthony R. Measham, and Jose Luis Bobadilla. *Disease Control Priorities in Developing Countries*. 1st ed. Oxford University Press, 2006. <http://documents.worldbank.org/curated/en/1993/10/698818/disease-control-priorities-developing-countries>.
- Jenni, Karen, and George Loewenstein. "Explaining the Identifiable Victim Effect." *Journal of Risk and Uncertainty* 14, no. 3 (1997): 235–257.
- Kahneman, Daniel, Ilana Ritov, David Schkade, Steven J. Sherman, and Hal R. Varian. "Economic Preferences or Attitude Expressions?: An Analysis of Dollar Responses to Public Issues." In *Elicitation of Preferences*, 203–242. Springer, 2000.
- Katz, Yehuda. "Handlebars." 2016. Accessed March 12, 2016. <http://handlebarsjs.com/>.

- Kessler, Johannes, and Michael Tschuggnall. "Malaria Atlas Project: The Spatial Distribution of Plasmodium falciparum Malaria Endemicity Map in 2010 in Zambia." 2010. Accessed March 1, 2016. https://web.archive.org/web/20120610235223/http://www.map.ox.ac.uk/browse-resources/endemicity/Pf_mean/ZMB/.
- . "Malaria Atlas Project." 2010. Accessed March 1, 2016. <http://www.map.ox.ac.uk/>.
- . "RelaX: Relational Algebra Calculator." 2015. Accessed January 30, 2016. <http://dbis-w6666.uibk.ac.at/ra/help.htm>.
- Kogut, Tehila, and Ilana Ritov. "The 'Identified Victim' Effect: An Identified Group, or Just a Single Individual?" *Journal of Behavioral Decision Making* 18, no. 3 (2005): 157.
- Kugler, Matthew, John T. Jost, and Sharareh Noorbaloochi. "Another Look at Moral Foundations Theory: Do Authoritarianism and Social Dominance Orientation Explain Liberal-Conservative Differences in 'Moral' Intuitions?" [In English]. *Social Justice Research* 27, no. 4 (2014): 413–431. ISSN: 0885-7466. doi:10.1007/s11211-014-0223-5. <http://dx.doi.org/10.1007/s11211-014-0223-5>.
- Lopez, Alan D. *Global Burden of Disease and Risk Factors*. Disease Control Priorities Project. World Bank Publications, 2006. ISBN: 978-0-82-136263-1. <https://books.google.de/books?id=F8Abr-of0wIC>.
- MacAskill, William. *Doing Good Better: How Effective Altruism Can Help You Make a Difference*. Penguin Publishing Group, 2015. ISBN: 978-0-69-819110-5. <https://books.google.de/books?id=sAgbBQAAQBAJ>.
- Make-A-Wish® America. "Frequently Asked Questions | Make-A-Wish® America." 2016. Accessed January 1, 2016. <http://wish.org/content/faq>.
- Marcus, Bernard. *Malaria*. Deadly Diseases and Epidemics Series. Facts On File, Incorporated, 2009. ISBN: 978-1-43-810156-9. <https://books.google.de/books?id=UgBpkGsg0ukC>.
- Masum, Hassan, Ronak Shah, Karl Schroeder, Abdallah S. Daar, and Peter A. Singer. "Africa's Largest Long-Lasting Insecticide-Treated Net Producer: Lessons from A to Z Textiles." *BMC International Health and Human Rights* 10, no. 1 (2010): 1–6. ISSN: 1472-698X. doi:10.1186/1472-698X-10-S1-S6. <http://dx.doi.org/10.1186/1472-698X-10-S1-S6>.
- Matthews, Dylan. "You have \$8 Billion. You Want to Do As Much Good As Possible. What Do You Do?" 2015. Accessed January 19, 2016. <http://www.vox.com/2015/4/24/8457895/givewell-open-philanthropy-charity>.
- Mercurial community. "Mercurial SCM." 2016. Accessed January 12, 2016. <https://www.mercurial-scm.org/>.
- Monbiot, George. "There's a Population Crisis All Right. But Probably Not the One You Think." 2015. Accessed March 5, 2016. <http://www.theguardian.com/commentisfree/2015/nov/19/population-crisis-farm-animals-laying-waste-to-planet>.
- MoralFoundations.org. "Moral Foundations Theory." 2015. Accessed December 26, 2015. <http://www.moralfoundations.org/>.
- Moyo, Dambisa. *Dead Aid: Why Aid is Not Working and How There is Another Way for Africa*. Penguin Books Limited, 2011. ISBN: 978-0-24-195996-1. <https://books.google.de/books?id=154Z1x2Nj-4C>.

- Nassar, Salma. "Why We Do Not Measure Impact At Scale." 2015. Accessed January 18, 2016. <http://www.evidenceaction.org/blog-full/why-we-do-not-measure-impact-at-scale-and-are-unapologetic-about-it>.
- Norwood, F. Bailey, and Jayson L. Lusk. *Compassion, by the Pound: The Economics of Farm Animal Welfare*. Titolo collana. OUP Oxford, 2011. ISBN: 978-0-19-955116-3. <https://books.google.de/books?id=kxUA9r4If8C>.
- Open Philanthropy Project. "Giving Now vs. Later." 2015. Accessed January 19, 2016. <http://www.goodventures.org/research-and-ideas/blog/giving-now-vs-later>.
- . "Introduction To Risk Aversion." 2015. Accessed January 19, 2016. <http://www.slideshare.net/txslr/introduction-to-risk-aversion>.
- . "Our Process." 2016. Accessed January 19, 2016. <http://www.openphilanthropy.org/research/our-process>.
- OpenLayers community. "OpenLayers 3." 2015. Accessed January 12, 2016. <http://openlayers.org/>.
- OpenStreetMap. "OpenStreetMap." 2016. Accessed March 7, 2016. <http://www.openstreetmap.org/>.
- Oracle Corporation. "MySQL: The World's Most Popular Open Source Database." 2016. Accessed March 7, 2016. <https://www.mysql.com/>.
- Ord, Toby. "Moral Trade." *Ethics, forthcoming* (2015). <http://www.amirrorclear.net/academic/papers/moral-trade.pdf>.
- Organization, World Health. "Conditions for Use of Long-Lasting Insecticidal Nets Treated with a Pyrethroid and Piperonyl Butoxide" (2015). <http://www.who.int/malaria/publications/atoz/use-of-pbo-treated-llins.pdf>.
- Packard, Randall M. *The Making of a Tropical Disease: A Short History of Malaria*. Johns Hopkins Biographies of Disease. Johns Hopkins University Press, 2010. ISBN: 978-1-42-140175-1. <https://books.google.de/books?id=-ePZU9CX9L4C>.
- Parkin, Michael, Melanie Powell, and Kent Matthews. *Economics*. Addison-Wesley, 2008. ISBN: 978-0-13-204122-5. <https://books.google.de/books?id=Kp6Ls7j3tVIC>.
- Pivotal Labs. "Jasmine: Behavior-Driven JavaScript." 2016. Accessed March 12, 2016. <http://jasmine.github.io/>.
- PostGIS Project Steering Committee. "PostGIS: Spatial and Geographic Objects for PostgreSQL." 2016. Accessed March 7, 2016. <http://postgis.net/>.
- PostgreSQL Global Development Group. "PostgreSQL: The World's Most Advanced Open Source Database." 2016. Accessed March 7, 2016. <http://www.postgresql.org/>.
- Pritchett, Lant. "Using 'Random' Right: New Insights from IDinsight Team." 2015. Accessed January 18, 2016. <http://www.cgdev.org/blog/using-%C3%A2%C2%80%C2%9Crandom%C3%A2%C2%80%C2%9D-right-new-insights-idinsight-team>.
- ProPublica. "Nonprofit Explorer – Make-A-Wish Foundation of America." 2014. Accessed December 9, 2015. <https://projects.propublica.org/nonprofits/organizations/860481941>.
- Pylons Project. "Pylons Project: Pyramid: About." 2016. Accessed January 12, 2016. <http://www.pylonsproject.org/projects/pyramid/about>.

- Python Software Foundation. "Welcome to Python.org." 2016. Accessed January 12, 2016. <https://www.python.org/>.
- QGIS. "QGIS: A Free and Open Source Geographic Information System." 2016. Accessed March 7, 2016. <http://www.qgis.org/>.
- Rigg, Jonathan, Anthony Bebbington, Katherine V Gough, Deborah F Bryceson, Jytte Agergaard, Niels Fold, and Cecilia Tacoli. "The World Development Report 2009 'Reshapes Economic Geography': Geographical Reflections." *Transactions of the Institute of British Geographers* 34, no. 2 (2009): 128–136.
- Robert E. Hall, Marc Lieberman. *Microeconomics: Principles and Applications*. Cengage Learning, 2009. ISBN: 978-1-43-903897-0. <https://books.google.de/books?id=XR9uvGmChp8C>.
- Roll Back Malaria. "Malaria Programmatic Gap Analysis: Guidance Notes" (2013). Accessed March 2, 2016. http://www.rollbackmalaria.org/files/files/toolbox/docs/rbmttoolbox/GapAnalysis/tool_GapAnalysisGuide-en.pdf.
- Ronacher, Armin. "Jinja2: The Python Template Engine." 2014. Accessed January 12, 2016. <http://jinja.pocoo.org/>.
- Roodman, David. "The Impact of Life-Saving Interventions on Fertility." 2014. Accessed March 5, 2016. <http://davidroodman.com/blog/2014/04/16/the-mortality-fertility-link/>.
- Sassi, Franco. "Calculating QALYs, Comparing QALY and DALY Calculations." *Health policy and planning* 21, no. 5 (2006): 402–408.
- Schwitzgebel, Eric, and Fiery Cushman. "Expertise in Moral Reasoning?: Order Effects on Moral Judgment in Professional Philosophers and Nonphilosophers." *Mind & Language* 27, no. 2 (2012): 135–153.
- SDG International. "Incoming Program Officer: Lewis Bollard." 2011. Accessed January 26, 2016. <http://www.openphilanthropy.org/blog/incoming-program-officer-lewis-bollard>.
- Shah, Neil Buddy, Paul Wang, Andrew Fraker, and Daniel Gastfriend. "Evaluations with Impact: Decision-Focused Impact Evaluation as a Practical Policymaking Tool." *International Initiative for Impact Evaluation (3ie)* (2015). http://www.3ieimpact.org/media/filer_public/2015/10/01/wp25-evaluations_with_impact.pdf.
- Singer, Peter. "The Triviality of the Debate over 'Is-Ought' and the Definition of 'Moral'." *American Philosophical Quarterly* (1973): 51–56.
- Small, Deborah A, and George Loewenstein. "Helping a Victim or Helping the Victim: Altruism and Identifiability." *Journal of Risk and Uncertainty* 26, no. 1 (2003): 5–16.
- Software Freedom Conservancy. "Git." 2016. Accessed January 12, 2016. <https://git-scm.com/>.
- Stragies, Stephanie. "Wursthersteller erstmals fleischfrei." 2015. Accessed December 30, 2015. <https://vebu.de/news/2419-wursthersteller-fleischfrei>.
- Terra Daily. "Global Fund Faces Billion-Dollar Gap." 2011. Accessed January 22, 2016. http://www.terradaily.com/reports/Health_Global_Fund_faces_billion-dollar_gap_999.html.

- The Economist. "It's Not What You Spend: How to Make Aid to Poor Countries Work Better." 2015. Accessed January 1, 2016. <http://www.economist.com/news/international/21651814-how-make-aid-poor-countries-work-better-its-not-what-you-spend>.
- The PhilPapers Foundation. "Preliminary Survey Results." 2015. Accessed December 25, 2015. <http://philpapers.org/surveys/results.pl>.
- The Web Framework for Perfectionists with Deadlines. "Django." 2016. Accessed January 12, 2016. <https://www.djangoproject.com/>.
- The Week. "The Little Dog Lost at Sea." 2015. Accessed December 5, 2015. <http://theweek.com/articles/496623/little-dog-lost-sea>.
- The World Bank. "Sub-Saharan Africa (Developing Only) – Data." 2015. Accessed December 9, 2015. <http://data.worldbank.org/region/SSA>.
- Tomasik, Brian. "Gains from Trade through Compromise." 2015. Accessed December 28, 2015. <http://foundational-research.org/gains-from-trade-through-compromise/>.
- . "Humane Insecticides" (2007). <http://reducing-suffering.org/humane-insecticides/>.
- United Nations Sustainable Development Solutions Network. "World Happiness Report 2016 Update." 2016. Accessed July 8, 2016. <https://worldhappiness.report/>.
- Webb, James L.A. "The Long Shadow of Malaria Interventions in Tropical Africa." *The Lancet* 374, no. 9705 (2009): 1883–1884.
- . *The Long Struggle against Malaria in Tropical Africa*. Cambridge University Press, 2014. ISBN: 978-1-10-705257-4. <https://books.google.de/books?id=eV72AgAAQBAJ>.
- WHO. "Malaria." 2016. Accessed January 22, 2016. <http://www.who.int/ith/diseases/malaria/en/>.
- World Bank. "GINI index (World Bank estimate)." 2016. Accessed July 9, 2016. <http://data.worldbank.org/indicator/SI.POV.GINI>.
- World Health Organization. "Global Burden of Disease 2004 Update: Disability Weights for Diseases and Conditions" (2004).
- . "Insecticide-Treated Mosquito Nets: A WHO Position Statement." 2007. Accessed January 3, 2016. <http://files.givewell.org/files/DWDA%202009/Interventions/Nets/itnspospaperfinal.pdf>.
- Zope Foundation. "Buildout: Software Build System Reloaded." 2015. Accessed March 7, 2016. <http://www.buildout.org/>.