



Research agenda for 2023-2024

Ryan Dwyer, Michael Plant

July 2023





Contents

Foreword	3
One-page summary	5
Research agenda	6
Applied research to maximise global wellbeing	6
1. Cause area explorations	6
1.1 Non-mood-related mental health issues (e.g., psychotic and trauma-related disorders)	6
1.2 Child development effects (e.g., abuse, trauma, nutrition)	7
1.3 Fistula repair surgery	7
1.4 Basic housing improvements (e.g., concrete floors)	7
Longlist of future cause areas to explore	7
2. Organisation evaluations	8
2.1 Updated evaluation of psychotherapy	8
2.1.1 StrongMinds	8
2.2 Mental health organisations	9
2.3 Non-mental health organisations	9
Longlist of future organisations to explore	9
Foundational research	9
3. The nature of wellbeing	10
4. The measurement of wellbeing	10
5. Moral weights	11
5.1 Using WELLBYs to compare the value of extending lives against improving lives	11
5.2 Moral uncertainty (what to do when you don't know what to do)	11
6. Longtermism and population ethics	12
Conclusion	12
Contributions	12



Foreword

The idea that the quality of a society should be judged by the happiness of its people is an old idea, stretching back at least to the Enlightenment, if not Aristotle. What's new is that now we have data: in the last few decades, researchers have developed valid and reliable ways to measure wellbeing through self-reports, and amassed a treasure trove of evidence. For the first time in human history, we can work out the best ways to improve the human condition in a scientifically rigorous way.

But what are these ways? That's what the Happier Lives Institute was set up to find out. We are the first organisation to combine the principles of effective altruism (that is, we should do the most good we can with the available resources) with the idea of taking happiness seriously (by prioritising using happiness data to figure out what does the most good).

We conduct *foundational research* on the nature and measurement of wellbeing, as well as *applied research* where we synthesise the existing data on subjective wellbeing to discover which interventions and organisations are the most cost-effective. Our research draws on and further develops previous work in the fields of philosophy, economics, and psychology. Ultimately, we measure impact in WELLBYs (wellbeing-adjusted life years), a method born in academia and now used by, for example, the [UK Treasury](#), that we are developing and deploying.

We aim to have impact by:

1. Identifying the most cost-effective ways to make lives happier, then communicating these to philanthropists and policymakers.
2. More broadly, engineering a paradigm shift towards a wellbeing approach among decision-makers: vast sums of resources are aimed at improving people's lives, and if those can be spent even a bit better, that would do enormous good. But we think they can be spent substantially better.

This is our third research agenda, which sets out our research goals through the end of 2024 (see our [first](#) and [second](#) agendas). When we started in 2019, we wanted to know if taking happiness seriously might change the priorities for donors seeking to do the most good. That is, if we evaluated – for the first time – different interventions in terms of their impact on subjective wellbeing, would we uncover different top charities?

To that end, our first major line of applied research explored the [impact of mental health interventions on wellbeing](#). We found that [psychotherapy](#) was a promising – yet overlooked – intervention in low-income countries. We compared the cost-effectiveness of psychotherapy to other well-evidenced interventions, including [cash transfers](#), [antimalarial bednets](#), and [deworming pills](#). From this work, we've found that psychotherapy is several times more cost-effective than cash transfers or deworming for improving happiness. We concluded that comparing psychotherapy to



antimalarial bednets, a life-saving intervention, [depends heavily on various philosophical assumptions](#): treating depression ranges from about as good as to several times better than antimalarial bednets, depending on the assumptions. This has [led us](#) – and [others](#)¹ – to conclude that evaluating interventions in terms of their impact on subjective wellbeing is an important approach to setting global priorities.

Along the way, we've also pursued foundational research to better understand the nature and measurement of wellbeing. This has included evaluating philosophical views of [wellbeing](#) and [life satisfaction](#), pioneering [methods to conduct cost-effectiveness analyses using wellbeing](#), and conducting [novel research on wellbeing measurement](#). We strive for maximum rigour in our work, so this fundamental research is key to ensuring that we are doing the most good possible by aiming for the right outcomes.

As a small team, we're proud of what we've accomplished in these four years, but there is still much to do. Looking ahead, we plan to:

- Assess a wider range of interventions and charities to find more opportunities for donors to improve the wellbeing of existing people.
- Conduct further theoretical work: although we have established the feasibility of the subjective wellbeing approach in general, we will continue to refine our approach as we push the boundaries of research on the nature and measurement of wellbeing.

Read on for specific details about what we have planned.

¹ See also the [discussion](#) between Elie Hassenfeld from GiveWell with Rob Wiblin on *The 80,000 Hours Podcast*.



One-page summary

Applied research to maximise global wellbeing

1. Cause area explorations

To find new promising solutions to the biggest problems, we will explore:

- Non-mood-related mental health issues (e.g., psychotic and trauma-related disorders)
- Child development effects (e.g., abuse, trauma, nutrition)
- Fistula repair surgery
- Basic housing improvements (e.g., concrete floors)

2. Organisation evaluations

To better inform donors of the best giving opportunities to promote wellbeing, we will evaluate the cost-effectiveness of several organisations (partially informed by our cause exploration work):

- Updated evaluation of psychotherapy and StrongMinds
- Mental health organisations (organisations TBD)
- Non-mental health organisations (organisations TBD)

Foundational research

3. The nature of wellbeing

We will explore the underlying theory of wellbeing and wellbeing measurement, including:

- An academic paper on life satisfaction theories of wellbeing

4. The measurement of wellbeing

We will conduct new research on how to measure and interpret subjective wellbeing measures:

- An academic paper setting out our method for measuring impact using wellbeing
- A “how-to” guide explaining how we undertake cost-effectiveness analyses
- A revised paper on the [theory and current evidence on scale cardinality](#) (e.g., is your 7/10 the same as my 7/10?)
- An experimental survey to test assumptions about subjective wellbeing measures, including comparability, linearity, and the neutral point

5. Moral weights

We will conduct new research about how to make tradeoffs between different units of value:

- A working paper exploring a bargaining-based approach to moral uncertainty

6. Longtermism and population ethics

We will explore whether we should improve the wellbeing of people alive now or in future generations:

- An academic journal book review of Will MacAskill's [What We Owe The Future](#)



Research agenda

Applied research to maximise global wellbeing

An ultimate goal of our work is to identify the most cost-effective opportunities to improve global wellbeing. To do this, we follow three stages to identify where additional resources can do the most good:

1. We conduct broad analyses of different causes. This starts with an educated guess about which problems have the right combination of being large, solvable, and unduly neglected.
2. We assess different possible interventions for each promising cause, looking for the most cost-effective ones.
3. Where relevant, we hone in and compare the top organisations implementing those interventions.

We can then compare the impact of the new interventions and organisations we examined to interventions and organisations we've previously evaluated. We expect to run this process in cycles – picking new cause areas to investigate, then narrowing down to the specific organisations – which will enable us to look broadly and deeply at the same time.

1. Cause area explorations

A big portion of our work to date has focused on interventions that address [mental health](#), such as psychotherapy. But we think there are other neglected problems – unrelated to mental health – that may have large impacts on wellbeing as well. So far we have completed shallow reviews on [pain](#), [lead exposure](#), and [immigration](#).

Over the next two years, we will continue to explore new cause areas both within mental health and outside of it. Our goal is to review a range of topics to determine which areas are likely to have promising evidence-based and cost-effective solutions. To our knowledge, we are the first organisation to look for the top global interventions and organisations working on improving wellbeing, so this is uncharted territory. We expect that mapping out the evidence across different areas will help us develop a framework for understanding what types of interventions are most effective for promoting wellbeing.

1.1 Non-mood-related mental health issues (e.g., psychotic and trauma-related disorders)

Our previous cause area report on [mental health](#) focused specifically on depression and anxiety, which are the most common mental health disorders. However, there are other disorders that can severely impact SWB. Our upcoming review will briefly examine the most promising mental,



neurological, and substance use disorders beyond mood disorders. This includes problematic alcohol and drug use, schizophrenia and other psychotic disorders, and epilepsy. We can't expect to examine all issues in detail; our more modest objective is to identify the non-mood-disorder topic that seems most promising.

1.2 Child development effects (e.g., abuse, trauma, nutrition)

There are several critical periods of development in childhood that can fundamentally alter wellbeing over the lifetime. Our shallow review will explore the key events that can impact development and promising interventions to improve long-term SWB. Because childhood interventions can improve SWB in both the present and the future, the benefits for these interventions could be particularly large. This is a broad topic, so we plan to conduct a very shallow review first, and then follow up with more detailed reviews on the most promising topics.

1.3 Fistula repair surgery

Obstetric fistula is an abnormal opening between a woman's genital tract and her urinary tract or rectum. As a consequence of an unrepaired fistula, women suffer life-long incontinence, which results in shame and social exclusion. We suspect that using a SWB approach will more accurately capture the negative impacts of this condition than would be measurable on a conventional economic or health framework (e.g., [DALYs](#)). Our upcoming review will focus on finding high-quality evidence on SWB outcomes, and cost-effective organisations that deliver fistula repair surgery.

1.4 Basic housing improvements (e.g., concrete floors)

Inadequate housing is a large problem in many areas of the world and has a wide range of adverse consequences on health and wellbeing. Dirt floors are an important subproblem of poor housing, as they can be difficult to clean adequately and can therefore provide a means for parasitic infection. This is associated with intestinal parasitic infestations, especially in children as they are more likely to ingest contaminated substances from dirt floors. Concrete floors may improve wellbeing by preventing these infections and providing a healthier home. Our review will focus on finding high-quality evidence on SWB outcomes, and cost-effective organisations providing concrete floors or similar interventions to address this concern.

Longlist of future cause areas to explore

We also have a longlist of [cause areas](#) we would like to explore when we have more time, including psychedelics, opioids, poverty, loneliness, sleep, and air pollution.



2. Organisation evaluations

As mentioned, HLI's starting mission was to evaluate the world's top organisations in terms of subjective wellbeing (SWB). So far, we have evaluated the cost-effectiveness of organisations that provide [psychotherapy](#), [cash transfers](#), [antimalarial bednets](#), and [deworming pills](#). From this work, we've found that psychotherapy for depression is several times more cost-effective than cash transfers for improving happiness, deworming has an unclear long-term effect, and treating mental health compares well to providing antimalarial bednets – although comparing quality and quantity and life is an under-appreciated and complicated issue.

We think there are other cost-effective organisations doing impactful work, so we plan to evaluate three to six more interventions in the next two years. This will include updating our evaluation of psychotherapy, evaluating new mental health organisations, and evaluating non-mental health organisations. We hope to find two new organisations to recommend in time for giving season 2023, but we cannot guarantee what our research will find.

2.1 Updated evaluation of psychotherapy

We first published a [cost-effectiveness analysis of treating depression](#) in late 2021. This was a meta-analysis where we looked at nearly all the available studies of lay-delivered group psychotherapy for depression in low- and middle-income countries. As far as we are aware, it was the first meta-analysis on the impact of treating depression in terms of SWB. We produced a second analysis in 2022 in response to comments; the primary change was to include an estimate of [household spillovers](#).

We are now planning a further update in response to additional comments (e.g., from [James Snowden](#) and [GiveWell](#)). We expect this will include updating our analysis with recently completed studies and refining some technical aspects of the analysis, including:

- Our systematic review, and the weight we place on different sources of evidence
- Estimated spillover benefits for household members
- Cost estimates
- Technical details, such as:
 - How long do the effects of psychotherapy last?
 - How important is the expertise of the deliverer or number of sessions?
 - Are the effects of psychotherapy affected by publication bias?

2.1.1 StrongMinds

StrongMinds is a nonprofit organisation that provides group interpersonal psychotherapy for women struggling with depression in Uganda and Zambia. We first recommended [StrongMinds as a top organisation](#) at the end of 2022. This recommendation is based on our evaluation of



psychotherapy, so we will also update our assessment of StrongMinds after we update our psychotherapy evaluation.

2.2 Mental health organisations

Based on our [cause area report on mental health](#) and our [cost-effectiveness analysis of psychotherapy](#), we think mental health is a promising area in which to find cost-effective interventions to improve wellbeing. StrongMinds was the first mental health-focused organisation we evaluated, but we know of others that may be similarly cost-effective. Therefore, we plan to evaluate other organisations focused on mental health. While some of these organisations may deliver versions of psychotherapy, we expect others will provide different types of mental health interventions, such as social-emotional learning. We expect to examine [Friendship Bench](#), [Sangath](#), and [CorStone](#) unless we find something more promising.

2.3 Non-mental health organisations

While our research indicates treating mental health is highly cost-effective, we are not confident it is the best way to improve global wellbeing, and we want to explore other possible options. Above, we mentioned our shallow reports on [lead exposure](#) and [pain relief](#), which found that these are both promising areas to improve wellbeing. We plan to scope new cause areas (as discussed above in Section 1) to find promising interventions and cost-effective organisations. Our ultimate goal is to conduct full cost-effectiveness analyses for the organisations that seem most impactful, and to compare them to our other recommended charities.

Longlist of future organisations to explore

We also have a long list of [organisations](#) we would like to explore, including the [Shamiri Institute](#), [Action for Happiness](#), and [Koko](#).

Foundational research

We put most of our effort into our applied research. However, we can't do this applied work in isolation, because doing it involves running into a whole host of other questions, often in moral philosophy or the philosophy of science. The notable ones are: What is wellbeing? How well can we measure it? How can we compare saving lives to improving lives? How should we compare the importance of affecting the long term to doing good now? What should we do when we are unsure, morally, what we ought to do?²

Attempting to sidestep these questions often means answering them implicitly via assumptions; where possible, we prefer to engage with the assumptions we're making explicitly and try to make

² Some of these have been discussed in Michael Plant's DPhil thesis, *Doing Good Badly? Philosophical Issues Related to Effective Altruism* (2019).



more sensible ones. In our foundational work, we aim to think about these questions, focusing on those that seem most decision-relevant.

We aim to eventually publish much of this work in academic journals, which means we have a number of ongoing projects. For this section, we distinguish three statuses of this work: existing work (where a public document is available), current work (where research is ongoing), and future work (where we may investigate the issue but haven't yet started doing so). Some work is both existing and current (where we have extant research we are updating).

3. The nature of wellbeing

What *is* wellbeing – that is, the thing that ultimately makes our lives go well for us?

Our existing work on this topic includes a summary of the [philosophy of wellbeing](#), which explains the three main rival accounts of what wellbeing is, as well as a paper critically evaluating life satisfaction theories of wellbeing (one of the main philosophical accounts of wellbeing) — [Life satisfaction and its discontents](#) — which we plan to revise.

In the future, we plan to:

- Investigate what a 'meaningful life' is and how much it matters
- Write a book on Taking Happiness Seriously for the general reader, which would combine many of these topics in an accessible format

4. The measurement of wellbeing

Whatever wellbeing is, can we measure it in a scientifically reliable way?

Our article [To WELLBY or not to WELLBY?](#) sets out the WELLBY method, its strengths, weaknesses, and areas for future work. To expand on this, we are:

- Developing an updated version to be published in an academic journal
- Writing a 'how-to' guide for other researchers building on this method that details how we undertake WELLBY cost-effectiveness analyses

Our working paper [A Happy Possibility about Happiness \(and other\) Scales](#), a working paper attempts to provide the first overview of both the theory and evidence of the comparability of subjective wellbeing scales (e.g., is your 7/10 the same as my 7/10?). We plan to revise this for publication in an academic journal.



To fill information gaps related to this, in [Can we trust wellbeing surveys?](#), we conducted a pilot survey on new empirical tests of the comparability of happiness data, among other things. We plan to run a larger version of the survey and publish the results in an academic journal.

In the future, we plan to:

- Assess the social desirability bias and other self-reporting biases in SWB data (for example: Do people give answers surveyors want? Is it a problem? If so, can anything be done?)
- Explore whether the measure of SWB matters (for example, if the key outcome is [happiness rather than life satisfaction](#), do we get different priorities?)
- Although unlikely, we may also do some work relating to animal welfare; a challenge is that we prefer to rely on self-reports, which animals can't give.

5. Moral weights

Our work in 'moral weights' considers how to make tradeoffs where we have different units of value. We've addressed this topic in two primary ways:

5.1 Using WELLBYs to compare the value of extending lives against improving lives

We have explored this topic first in our report [Estimating moral weights](#) and then more extensively in [The elephant in the bednet](#). The latter emphasised the role of key philosophical issues (such as accounts of the badness of death, the 'neutral point' equivalent to non-existence on a 0-10 life satisfaction scale) in calculating the cost-effectiveness of top interventions. We also published a [basic application](#) that allows users to input their own views to see the implications of them.

In the future, we may look further into the plausibility of the different accounts of the badness of death.

5.2 Moral uncertainty (what to do when you don't know what to do)

We've published two working papers on moral uncertainty: [The property rights approach to moral uncertainty](#) and [Wheeling and dealing: An internal bargaining approach to moral uncertainty](#), which both explore a novel, bargaining-based approach to acting when you're uncertain what's morally right. (This is very roughly akin to the '[moral parliament](#)' approach.) We're currently working with two external co-authors on a new paper that combines these ideas, which we plan to publish in an academic journal.

In the future, we plan to explore some sort of 'moral algorithm' that, when combined with empirical data, would indicate how you should distribute your resources given different accounts of morality and moral uncertainty.



6. Longtermism and population ethics

In HLI, we primarily strive to improve the wellbeing of people alive now. But should we? *Longtermists* would say we do more good by focusing on future generations.

To explore this claim, we're working on an academic journal book review of Will MacAskill's [What We Owe The Future](#), which recently and prominently made the case for longtermism. Our review raises some challenges to the presentation of and claims in the book.

In the future, we plan to explore:

- Whether making people happier now is a credible cause for longtermists. This may seem like a suspicious convergence, but longtermists often claim priorities such as AI alignment and preventing pandemics are important, even if we solely consider present wellbeing, so we shouldn't dismiss the possibility. There is potentially a case to be made that if people were happier, perhaps they would be kinder and more cooperative, start fewer wars, and so on.
- Work in [population ethics](#), particularly on the plausibility of '[person-affecting](#)' views: in slogan form, those "in favour of making people happy, but indifferent about making happy people". If such views are true, that would count against longtermism; how much they count also depends on one's view of moral uncertainty.

Conclusion

Over the past four years, we have shown that it is not only possible to compare interventions and organisations in terms of their impact on wellbeing, but, by doing so, we've discovered new global priorities. Over the next two years, we plan to further develop our approach and to uncover more cost-effective opportunities to improve wellbeing around the world.

In addition to the work we have outlined here, we have a [full list of research ideas](#) that we hope to explore. Please let us know if you can suggest any topics that we should research. If you have any comments or questions on the research agenda, or you would like to work on any of these issues, we strongly encourage you to get in touch at hello@happierlivesinstitute.org. We look forward to hearing from you.

Contributions

The entire HLI team contributed to the development of our research agenda. Ryan Dwyer wrote the original draft, and reviewed and edited subsequent drafts. Michael Plant wrote the section on foundational research, and reviewed and edited drafts. Katy Moore copy-edited.