

<https://researchcommons.waikato.ac.nz/>

Research Commons at the University of Waikato

Copyright Statement:

The digital copy of this thesis is protected by the Copyright Act 1994 (New Zealand).

The thesis may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis. You will recognise the author's right to be identified as the author of the thesis, and due acknowledgement will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from the thesis.

**Future Global Wellbeing Assessment:
Conceptual Foundations, Governance & Artificial Intelligence**

A thesis
submitted in partial fulfilment
of the requirements for the degree
of
Master of Social Sciences in Philosophy
at
The University of Waikato
by
Lewis MacKinnon



THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

2026

Abstract

This thesis examines how global wellbeing assessment might evolve in an increasingly globalised and digital society. It argues that, as wellbeing becomes a more explicit objective of global governance, a challenge arises over how it should be conceptualised and assessed across diverse cultural settings.

Part 1 evaluates dominant approaches to wellbeing assessment, tracing the history to several key Western traditions in psychology and philosophy. It includes an argument against a bounded, individualised self being the foundational unit of societal wellbeing, and argues for societal wellbeing to be evaluated beyond the aggregation of self-evaluating individuals. It also explores alternative ontologies of who and what might be capable of experiencing wellbeing, and advocates considering views in which wellbeing exists not only within individual people, but also in the connections between people, in animals, and possibly in the environment more widely. It also argues for the importance of considering a wider range of cross-cultural ontologies in wellbeing assessment frameworks if they are to be used in global contexts, and the risk of ‘wellbeing imperialism’ if culturally biased views are applied globally.

Part 2 explores a parallel trend of how technology and artificial intelligence (AI) are changing the way wellbeing is assessed. It explores the use of data from online platforms, large corpus analysis, and the use of sensor and wearable technology to infer aspects of wellbeing, like emotion and mental health. It speculates on how these may advance and be used to understand the wellbeing of society at scale, and may be incorporated into future wellbeing assessment systems in an increasingly digitised world. It evaluates both the promise of such systems and the risks, including privacy loss, bias, and misuse for corporate profits and government control. Technology and AI are presented as raising the stakes on the conceptual and ontological issues in Part 1.

Overall, the thesis presents a broad exploration and synthesis of foundational concepts, ontologies, and technologies in assessing wellbeing, and forms connections across diverse disciplines and cultural contexts. Ultimately it aims to contribute to the discussion on the future of global wellbeing assessment frameworks.

Table of contents

Abstract	2
Preface and Acknowledgements	4
AI usage statement	5
1. Introduction	6
1.1 The trajectory towards global governance	6
1.2 Institutional shift toward wellbeing	6
1.3 Cross-cultural and cross-disciplinary measurement of wellbeing	8
1.4 Dominant definition of wellbeing and reorientation in this thesis	9
1.5 Technology	10
1.6 Research questions and roadmap	11
Part 1: Definition, Conceptualisation, and Foundations	13
2. Traditional foundational wellbeing theories, concepts, and systems for assessment	13
2.1 Eudaimonia and hedonism	13
2.2 Utilitarianism and the measurement of wellbeing	13
2.3 Subjective wellbeing, life satisfaction, and affect	14
2.4 Mental health	15
2.5 Objective wellbeing measures	16
2.6 Summary	17
3. The individualistic and humanistic assumptions	18
3.1 Dominant psychological views of self	18
3.2 Alternative views to the bounded individual as the central unit of collective wellbeing	19
3.3 Cross-culturalism as a principle in global wellbeing assessment	22
3.4 Beyond humans	23
3.5 Expansive consciousness and expansive wellbeing	25
3.6 A brief note on theistic viewpoints	27
4. Quantitative wellbeing measurement and maximisation	29
4.1 Wu wei and pursuing maximum wellbeing:	29
4.2 Yin yang and positivity maxing	30
4.3 Alternative measurement ideas and qualitative assessment:	30
Part 2: Technology and future wellbeing assessment systems	32
5. Technological breakthroughs	32
5.1 Algorithmic analysis of text inferring survey-based measures of wellbeing	32
5.2 LLMs enabling contextualised, flexible wellbeing inference	34
5.3 Wearables, sensors, and speculation of future multimodality	39
5.4 Implications for wellbeing assessment	40
6. Bias, Privacy, Surveillance, and Trust	41
6.1 Bias	41
6.2 Privacy	42
6.3 Trust, surveillance, and control	43
7. Wellbeing alignment in the age of AI	46
8. Conclusion	48
References	50
Appendix	62

Preface and Acknowledgements

I would like to offer my sincere thanks to my supervisor Dr Dan Weijers for his help throughout this thesis. He was patient, open and flexible throughout the process, and our discussions were genuinely fun. His knowledge and experience in philosophy helped to anchor me in the discipline, which was particularly valuable for me as someone who had not formally studied philosophy or social sciences at an undergraduate level.

I would also like to express my gratitude to the researchers who generously shared their time and expertise during our discussions, allowing me to better understand their research and test my own ideas. This included Dr Shamubeel Eaquad, an expert in social policy and economics in a New Zealand context; Dr Timothy Lomas, an expert in the field of wellbeing science and one of the lead researchers on the Global Flourishing Study (discussed later); Dr Selim Sametoğlu, an expert in psycholinguistics publishing on the use of social media data to infer wellbeing; Dr Satpreet Singh, a machine learning expert publishing on psychological analysis of text; and Dr Manoj Shakyia, an expert on AI and Deep Reinforcement Learning.

My thanks also go to the students at Shanghai Jiaotong University, Fudan University, and Kathmandu University for participating in interviews on these subjects (conducted with the approval of the ALPSS Human Research Ethics Committee). Their views helped me understand how a subset of students in Shanghai and Kathmandu thought about wellbeing and AI.

Additionally, I want to thank Deb, Hugo, Harry and Reuben for their feedback, including meticulous proofreading and editing, structuring advice, and fun brainstorming sessions. Finally, thanks to my parents for their unwavering support, discussions, feedback, and for looking after me with love and care.

AI usage statement

AI tools including ChatGPT, Claude and Google Gemini were used during the process of writing this thesis. In the beginning stages they were primarily used for learning new concepts, exploring ideas, and brainstorming. They were valuable in helping me understand some of the foundational ideas in the field. The “deep research” function was also used to get an overview of literature focussed on specific research questions. However, as my research questions became more specific, and as I became more familiar with the literature, it became less useful for finding papers. Any sources provided were verified externally.

I also took a unique approach to analyse large corpora of personal notes and other writings on related topics using both the OpenAI and Google AI Studio API functions. I had a large number of notes created over several years and, because there were too many to send to the web app of ChatGPT or Gemini, I used Python scripts and API calls to process the notes and writings. It was a useful exercise to explore patterns and connections in my thinking over time. I also conducted a similar process with downloaded references.

Following this, I began writing my thesis directly. Throughout the writing process, I continued to use AI to help me learn unfamiliar concepts, research new topics, and as a discussion partner to explore ideas. It was also used to suggest possible edits to grammar, spelling, clarity, and identify errors (like unnecessary repetitions) quickly across the document. I considered any suggestions and incorporated them at my discretion. It was not used to generate or write my thesis for me.

I recognise that its responses, when providing feedback, discussing ideas, or otherwise are prone to hallucinations, and are inherently biased towards its training data and the platform's reinforcement learning protocol. Additionally I'm aware there are risks about data privacy, which I mitigated by using the paid tier and disallowing permission to use my conversations to train future models. Although I'm aware even this isn't foolproof data protection as whatever information provided is still accessible by the company.

It is worth noting, that due to my heavy AI use, writing I have produced wholly myself sometimes contains certain “AI-isms” such as using epithets or dashes. Also, there is a passage in section 5.2 that uses AI analysis as part of the thesis methodology itself. They are clearly labelled as responses written by ChatGPT, Gemini, and Claude.

I stand by everything in my thesis as my own ideas.

1. Introduction

1.1 The trajectory towards global governance

History can be viewed through the lens of an expanding circle of concern, from families, tribes, and clans to countries and cultures and, ultimately, to all people and beyond, according to the contemporary philosopher Peter Singer (1981). He argues that people have progressively developed empathy and morality that conveys concern for those beyond their family and immediate relations.

The historian and author Yuval Noah Harari, in his book *Sapiens: A Brief History of Humankind* (2014), also observes humanity's trajectory towards globalism through the formation of tribes, nations and now, global organisations. He argues that doing so has been enabled through members of groups having a shared “intersubjective reality” – a shared consciousness or subjective experience that exists in the minds of many people and not in the material world. He claims these “imagined realities”, like money, legal systems and religious narratives are what allow us to work together flexibly in such large groups, and enable cooperation on a global scale.

Both authors point out that history has been characterised by a long-term trend toward growing global interconnectedness. For Singer, this manifests in a more globalised morality, and for Harari, a more globalised shared subjective experience.

With regards to governance institutions growing larger and wider-reaching, there is empirical evidence to support their claims. Over the past few centuries modern nations have been increasingly coordinated through intergovernmental organisations, giving rise to systems of global governance. Today, there are a large number of transnational institutions and organisations, such as the Association of Southeast Asian Nations (ASEAN), the African Union (AU), the BRICS bloc (Brazil, Russia, India, China, and South Africa, and five others now), the European Union (EU), the Organisation for Economic Co-operation and Development (OECD), and the main locus of global governance – the United Nations (UN) (Barnett et al., 2021).

There has also been an expansion in the number of non-state groups such as transnational companies, academic organisations, non-governmental organisations (NGOs) and informal networks (Barnett et al., 2021). Together, these have been enabled by the creation and use of new tools of global connection like the internet, personal smart devices, global commerce, and efficient means of global travel and migration.

This thesis is focussed on this increasingly globalised context: where there is a long-term growth of institutions and organisations that act at a global scale, there is a growing moral concern outside our kin to those around the world, where we increasingly share stories and subjective realities across the world, and where this is enabled through a number of key technologies like the internet and personal devices.

1.2 Institutional shift toward wellbeing

Many global governance institutions traditionally focus on addressing global challenges such as war, poverty, disease, and natural disasters (Barnett et al., 2021). However, their role increasingly extends beyond responding to crises to promoting the broader development, flourishing, and wellbeing of societies. Briguglio et al. (2025) describe wellbeing as becoming a central objective of public policy worldwide, with growing efforts to define, measure, and use wellbeing indicators to guide government

decision-making and societal progress. One example can be seen in the widely publicised UN Sustainable Development Goals (SDGs), which have been the centrepiece of the UN's development framework since 2015, and place wellbeing as a key part of development (UN General Assembly, 2015, p. 16). The OECD utilises three wellbeing-focussed dashboards aiming to guide public policy in member states and inform the public, which has evolved from the original Better Life Index, started in 2011 (Exton & Shinwell, 2018).

Countries from diverse geographic, political, and religious backgrounds have been using wellbeing language in their policy objectives. Bhutan was one of the original countries to promote a shift to wellbeing-focussed policy with the promotion of Gross National Happiness¹ in the 1970s, which has now developed into a more formalised empirical index of assessment (Lepeley, 2017). New Zealand has been prominent in the promotion of wellbeing policy with their Living Standards Framework, Indicators Aotearoa New Zealand, and the He Ara Waiora dashboard, all aimed at measuring national wellbeing (the latter specifically for indigenous Māori wellbeing) (New Zealand Government, 2025; Hughes, 2021). They are also a participant in the Wellbeing Alliance Economies, and have passed legislation requiring the government to make a formal assessment of national wellbeing every four years (New Zealand Government, 2025, see Public Finance Act 1989, section 26NB; Wellbeing Economy Alliance, n.d.). The United Arab Emirates (UAE) has created a Minister of State for Happiness and Wellbeing, and has set a number of targets in the National Strategy for Wellbeing 2031 (UAE Cabinet, 2019; Lambert et al., 2022). China's government's use of the term 'wellbeing' in English translations of their seminal five-year plans has been a new feature of the last few editions (People's Republic of China Government, 2011, 2016, 2021) – notable in that it is from a significantly different political system to Western liberal democracies and is, by some measures, the largest political party in the world (Xinhua, 2025).

The measurement of wellbeing was initially developed in large part due to discontent with other measures of societal development such as Gross Domestic Product (GDP). It was seen as a move towards capturing the actual quality of life of citizens more directly. An early example of this critique was a paper by Ed Diener and Martin Seligman, leading figures in subjective wellbeing research and positive psychology, which promoted the shift of wellbeing objectives beyond psychology and academia, into the goals and targets set in government policy (Diener & Seligman, 2004). This was followed by the Stiglitz-Sen-Fitoussi Report (2009), another critique of GDP as a measure of national success that instead promoted holistic wellbeing-focussed and sustainable measures of development. It was particularly significant in that it was commissioned by a major government (France) showing a clear shift into government policy. Since then, numerous frameworks, measurement systems, and research projects followed and have continued to this day. A 2024 study by Jansen et al. found that 27 countries from diverse geographic, political, and religious backgrounds currently use 'Beyond GDP' measurement frameworks in the evaluation of their national progress. 'Beyond GDP' measures are about "measuring sustainable and inclusive wellbeing" (Jansen et al., 2024, paper title).

The interest in an empirical understanding of wellbeing has also grown inside academia. Many leading universities have established wellbeing-focussed research centres in the last several decades, including Oxford's Wellbeing Research Centre, Cambridge's Wellbeing Institute, and the University of Melbourne's Centre for Wellbeing Science. In 2011, the International Journal of Wellbeing was founded (Weijers & Jarden, 2021). In the 40-year period from 1980 to 2020, there has been

¹ The interest of governments in assessing happiness, while not synonymous with wellbeing, can be viewed as part of a broader shift towards an interest in subjective experience. A definition of wellbeing and its connection with subjective experience is given in section 1.4, and further distinctions between happiness and wellbeing are explored in section 4.

approximately a twenty-fold increase in Google Scholar search results for the terms “wellbeing” or “well-being” (adjusted for total search results). Google Ngram, a way to measure the prevalence of terms in text available to Google servers (e.g. books and webpages) and Google Trends, which measures trends in search terms, both show a steady and continuous upward trend in the use of the terms “wellbeing” and “well-being” (see appendix 1).

1.3 Cross-cultural and cross-disciplinary measurement of wellbeing

Given the need for global institutions to have policy with relevance across diverse cultural settings, they face the challenge of defining and measuring wellbeing cross-culturally too. The World Happiness Report is a high-profile UN publication, self-described as the “world’s foremost publication on global wellbeing” (Helliwell et al., 2025). It evaluates wellbeing by asking 1000 participants from each country to evaluate their life from zero (the worst possible life for them) to ten (the best possible life). The average of all the participants in the country becomes the country’s score. This, along with wellbeing science in general, has been criticised as having a narrow and culturally specific definition that fails to capture many of the most important aspects of wellbeing. VanderWeele et al. (2025) state “the research on well-being has been shaped largely by Western perspectives and has been carried out mostly in Western contexts” (p. 1). Joshanloo and Weijers (2024) point out how many influential modern wellbeing models are based on mainly Western concepts, developed by mainly Western researchers, and validated mainly using Western participants. A number of other publications also point out that much of modern wellbeing science has occurred within Western academia and psychology (Lambert et al., 2020; Biswas-Diener, 2022; Smith et al., 2025; Diener & Seligman, 2009; Briguglio et al., 2025).

But this is beginning to change. There are growing observations of the wellbeing field globalising. Lomas describes the field as occurring in four to five overlapping ‘waves’. The first as psychotherapy and psychiatry, the second as humanistic psychology, the third as positive psychology, and the most recent as a growing wave of *global* scholarship (Lomas, 2022a). He also speculates on a future fifth wave which is discussed further in section 3.4 ‘Beyond humans’, and section 3.5 ‘Expansive consciousness and expansive wellbeing’. In terms of journal publications, the *International Journal of Wellbeing* has published two special issues specifically on cross-cultural and global perspectives, once in 2012 and again in 2022 (Hommerich, 2012; *International Journal of Wellbeing*, 2022). In 2024, the first set of results of the Global Flourishing² Study (GFS) – a 200,000+ participant and 22-country, multi-dimensional survey-based study – was published, which aims to survey a broader range of measures which capture wellbeing dimensions valued in more global contexts (VanderWeele et al., 2025). They state “The GFS was conceived to... advance our knowledge of flourishing in general and especially in non-Western contexts” (VanderWeele et al., 2025, p. 1). The GFS will be discussed as an example of a recent research project aimed at addressing some of the cultural biases in historic systems of measuring wellbeing.

In terms of policy frameworks, the OECD and UN have developed numerous policy-guiding documents over the last decade to address the need for cross-cultural wellbeing assessment frameworks. Several key initiatives mentioned earlier, such as the Sustainable Development Goals and the Better Life Index, provide high-level frameworks for measuring development, which include a focus on wellbeing; but there have also been specific guidelines put out on measurement – the OECD Guidelines on the Measurement of Subjective Wellbeing (2013, 2025) and the United Nations Economic Commission for Europe (UNECE) Guidelines on the Measurement of Wellbeing (2025).

² Flourishing is defined as a broad, multidimensional form of wellbeing (VanderWeele et al., 2025).

One of the main aims of the 2025 update to the OECD guidelines was to “provide more globally inclusive approaches to subjective well-being measurement” (OECD, 2025, p. 2). These will be discussed as contemporary, global policy assessment frameworks.

Although, these projects far from exhaust the rich cross-cultural knowledge relating to wellbeing and they carry over some assumptions from their cultural foundations. The term ‘wellbeing’ is an English-language term, and has more direct links with the cultural and philosophical ancestors of the West, but discussions on the topic are prevalent across human intellectual history. In 2016, Timothy Lomas, one of the authors of the GFS, published a lexicography of wellbeing-related terms from around the world. It illustrates the rich variety of culturally unique and untranslatable terms related to wellbeing. It shows that wellbeing is a deep and varied topic, and that restricting it to any one language significantly limits the discourse.

This thesis looks to build on the progress by further exploring cultural conceptual assumptions related to the dominant views embedded in the systems of wellbeing assessment. Fortunately, the topic also has deep cross-cultural roots to explore (even if terms go by non-English names). This thesis touches on a small fraction of those but represents a step towards greater cross-cultural foundations. It argues that, for wellbeing frameworks to have legitimacy in global governance settings, we should continue to explore ideas outside any single cultural or disciplinary setting.

1.4 Dominant definition of wellbeing and reorientation in this thesis

There is a significant and ongoing discourse on the precise definition of wellbeing. Some aspects of the dominant definitions will be adhered to, and others will be challenged. Many current definitions used in prominent measurement frameworks define wellbeing as closely related to quality of life (Helliwell et al., 2025; OECD, 2025; UNECE, 2025). This thesis will also consider wellbeing as closely related to quality of life.

Intrinsic value:

The Stanford Encyclopedia of Philosophy states “the concept of well-being is most commonly used in philosophy to describe what is non-instrumentally or ultimately good for a person” (Crisp, 2026). The use of the term ‘non-instrumentally’ captures a key element of its definition – that wellbeing has intrinsic value to the one living it. Whatever ultimately constitutes wellbeing, has value to the one living it, as an end-in-itself and not as a means to an end. Wellbeing will be considered something with intrinsic value.

The individual:

The Routledge Handbook of Philosophy of Well-being (Fletcher, 2016) states philosophical discussions typically understand wellbeing in terms of what makes a person’s life go better or worse for them. This reflects the idea that wellbeing is commonly understood as something which affects people, and, more specifically, the individual. In other words, the individual is often considered the central unit of wellbeing (Joshani, 2026). This will be expanded on later in the thesis to explain the historical foundations which have contributed to wellbeing being centred on the individual. This thesis will challenge whether the individual person, at least as it is commonly understood in Western academia, should always be treated as the central unit for understanding global wellbeing. Cross-cultural concepts about the individual emerging through relationship with others and the environment are discussed, and are used to challenge the dominant notion of the individual as the appropriate unit of societal wellbeing. The discussion extends into exploration of “what is capable of

wellbeing?”, including non-human animals and speculates on a metaphysical view which considers the environment capable of wellbeing.

Experience:

Diener’s influential Subjective Wellbeing definition is a centrepiece of wellbeing discourse, having been formally part of wellbeing assessment discussions since at least the 1980s (Diener, 1984). It regards wellbeing as a person’s self-evaluations of their own life, in terms of emotional affect and overall life evaluations. It relates to someone’s subjective experience (Diener, 1984). This conception of subjective wellbeing is a foundation to the 2025 OECD guidelines to be used for transnational wellbeing assessment, and remains a central part of many existing national and transnational wellbeing frameworks (OECD, 2025; UNECE, 2025; Briguglio et al., 2025). While this thesis will contest the Western individualistic assumptions on which this definition rests, it does take the view that experience is ultimately what matters for wellbeing. In this thesis, wellbeing will be considered to be about the quality of conscious experience, or simply, experience.

Balance:

Many dominant conceptions of wellbeing simplify it into a positive and negative scale, where positive is good and negative is bad. This is stated explicitly in one of the dominant theories of wellbeing, hedonic wellbeing, which considers pleasure as good and pain as bad (Fletcher, 2016; Crisp, 2026). Many of the models that contain measures of wellbeing, such as the Subjective Wellbeing Scale, the World Happiness Report, the Global Flourishing Study, and other national or transnational frameworks, consider wellbeing on a quantitative, numerical scale, and treat higher up the scale as better, even when seen as part of multidimensional dashboards. This thesis will discuss challenges to framing wellbeing as something which can be modelled through such positive-negative quantitative scales. Theories which accept positive and negative and prioritise balance will be examined in the context of societal wellbeing assessments, such as the notion of *yin yang*. While the thesis doesn’t explicitly condemn positive-negative continuums, it explores the possibility that this is an oversimplification, and doesn’t reflect a universal view of wellbeing.

Overall, wellbeing is considered to be about the quality of conscious experience, and closely related to ‘quality of life’. It is considered something of intrinsic value to those experiencing it. The idea that wellbeing is something only experienced in individual humans is challenged, by breaking down the foundation of a bounded individual and arguing for the possibility of wellbeing existing beyond humans. It questions whether wellbeing is best viewed as something on a continuum from positive to negative, by touching on views which prioritise balance and harmony.

1.5 Technology

In parallel to growing global interconnectedness and the rise of international institutions, another major trend is reshaping the way wellbeing is defined and measured. Approximately two thirds of the global adult population have a personal internet-connected smart device, and a similar proportion use social media, according to an annual survey on global digital technology usage by Kemp (2025). Some sources report a global average screen time exceeding 40% of waking life (Istrate et al., 2024), and over 10% of the global adult population are active weekly users of chatbots, with this figure continuing to trend upwards (Chatterji et al., 2025).

This has resulted in an explosion of new data, including movement data from smartphones, physiological data from wearable devices, usage data from digital platforms, and language data from

social media and chatbot conversations. Many platforms synthesise this data through advanced algorithms to understand users and provide more personalised services.

As more intimate data sources become more widespread, and AI algorithms are becoming more advanced, there is beginning to be the ability to understand aspects of users' mental states and inner experiences. For example, data from smart devices are being used to infer emotion, mental health, and wellbeing survey responses (Shen et al., 2025; Sametoğlu et al., 2024; Brickman et al., 2025; Lalk et al., 2025; Nandini et al., 2025). This is the focus of Part 2: how digital technologies are rapidly reshaping the ways wellbeing and conscious experience are understood. Part 2 surveys several of these methodologies and goes on to speculate on how these methods may, in time, reshape future systems for the assessment of societal wellbeing.

The opportunities provided by such systems are discussed, in that they may allow more detailed, responsive, and cost-effective ways of understanding wellbeing. But there are also discussions about the significant risks related to the violation of personal privacy and amplification of biases. Additionally, there is a discussion of concerns relating to mistrust in the organisations and institutions with access to this data, such as their use for profit, and exerting control.

1.6 Research questions and roadmap

The overarching research question of this thesis is: How might future global wellbeing assessment frameworks evolve in light of rapid advances in AI and the globalisation of the concept of wellbeing? Furthermore, if this reality is possible, how might we mitigate risks and harness any potential benefits?

This thesis is presented in two parts. Part 1 examines the conceptual foundations of wellbeing assessment. It begins by tracing the several historically influential ideas relating to wellbeing, including hedonic, eudaimonic, and subjective wellbeing approaches, as well key foundations of wellbeing measurement such as those from utilitarian thinkers such as Jeremy Bentham. Part 1 then examines how these ideas have been operationalised in major policy and research frameworks.

Several concerns are raised in Part 1 about the reliance on assumptions that are not universally applicable. In particular, many contemporary frameworks assume the individual person as the primary unit of wellbeing, and assume wellbeing is best modelled on linear quantitative scales. Alternative views of selfhood and wellbeing are explored, including relational and 'extended-self' views, Buddhist non-self and dependent origination, as well as indigenous Māori *whakapapa* (genealogy) and interconnectedness with nature. These are discussed in relation to the questions of who and what is capable of wellbeing, and of moral concern.

Part 2 then turns to the technological conditions under which future wellbeing assessment may occur. It provides several examples of emerging capabilities inferring aspects of wellbeing, such as the prediction of regional wellbeing scores from online discourse and the inference of emotions from wearable devices. Section 5 discusses the opportunities presented by these as supplements to survey-based methods, and provides a personal example using a large personal language corpus to identify candid, longitudinal, qualitative elements of wellbeing. Section 6 delves into the risks, and touches on several strategies of mitigation, such as ways of preserving personal privacy. Section 7 explores the added risks associated with future AI systems which become more capable and intelligent, and how growing capabilities raise the stakes for the need to correctly define wellbeing to avoid the exacerbation of cultural and conceptual biases.

The thesis concludes by reiterating the need for awareness and caution, rather than disengagement with AI and technology, and encourages that, regardless of technological changes, further discussions on the nature of experience, wellbeing and ‘the good life’ is one of the most important topics in global institutional contexts.

It should be noted that this thesis is wide-reaching, intentionally spanning a range of perspectives from multiple disciplines and cultures. As a result, it does not exhaustively cover each in depth, but aims to synthesise and form connections broadly.

Part 1: Definition, conceptualisation, and foundations

2. Traditional foundational wellbeing theories, concepts, and systems for assessment

2.1 Eudaimonia and hedonism

Two of the most prominent philosophies to have shaped modern concepts of wellbeing are the ancient ideas of hedonism and eudaimonia (Michalos & Weijers, 2017; Ryan & Deci, 2001; Crisp, 2026). Hedonism can be described as the maximisation of pleasant experience and the minimisation of unpleasant experience, broadly speaking (Moore, 2019). While often caricatured in pop-culture as about short-term, self-centred, sensory pleasures like food, sex, and partying (Einzelgänger, 2024), more holistic views account for delayed gratification, altruism, and “higher” pleasures, like those associated with morality, creativity, or spirituality (Mill, 1863/2010; Weijers, 2011). The claim about bodily sensations being less important may be disputed, but it is important to recognise that holistic hedonic views of wellbeing are still able to account for long-term, meaningful pleasures, and may be consistent with wide-reaching moral concern for others. Hedonic philosophy is prominent in contemporary wellbeing assessment systems. The UNECE 2025 guidelines on the measuring of wellbeing states in Chapter 1 that it captures “hedonic dimensions of human experience” (p. 12). The OECD guidelines incorporate questions that capture people’s hedonic experiences. Both guidelines define subjective wellbeing in part through affect, measured through positive and negative feelings reflecting the hedonic tradition. The Global Flourishing Study similarly includes affect-related questions, treating positive and negative emotions as part of its broader multidimensional assessment of flourishing (VanderWeele et al., 2025).

Eudaimonic wellbeing is comparatively nebulous and is less reducible to a balance of one thing over another. It takes its roots from Ancient Greece and is defined through Aristotelian ideas of virtue, knowing thyself, functioning well, and a sense of purpose (Ryan & Deci, 2001). In the OECD guidelines, eudaimonic wellbeing is one of the three core modules of subjective wellbeing. Likewise it features in the UNECE framework. It is assessed through questions on meaning, purpose, and psychological functioning.

2.2 Utilitarianism and the measurement of wellbeing

Utilitarianism is another foundational theory related to wellbeing, although its influence is mostly as a moral and political theory connected to the *measurement* of wellbeing (Crisp, 2026). One of the key moral principles of Jeremy Bentham, often considered one of the key founding figures of utilitarianism, relates to the promotion of happiness and the reduction of pain (Bentham, 1789/1907, p. 1). He dedicated several decades of his work into developing a system of calculating it through his *hedonic calculus* (Weijers, 2011; Crisp, 2026; Baujard, 2009). The hedonic calculus looked at several factors including the intensity, duration, extent, and the certainty of pleasure or pain (Weijers, 2011). Although this system was never operationalised in practice, a theoretically complete version would, in his mind, have been able to calculate the morally best action.

According to Crisp (2026), this can be interpreted in terms of wellbeing. He states that “according to [utilitarianism] the only moral requirement is that total or overall well-being be maximized” (section 3, para. 2). Whether this is precisely the original intention of Bentham cannot be confirmed, but it is nevertheless true that he was an original advocate of a system of calculation that closely relates to hedonic wellbeing today.

One of the philosophical principles of Bentham’s hedonic calculus may have been inherited by contemporary cross-cultural measurement frameworks, even if the specific methods differ. This relates to the importance of pleasure over pain, as represented by both the UNECE and OECD 2025 guidelines on wellbeing measurement in the way that surveys ask about people’s pain and happiness.

2.3 Subjective wellbeing, life satisfaction, and affect

During the 20th century, work on the empirical assessment of wellbeing developed further. There was significant development in the use of quantitative self-evaluations in the form of large-scale social surveys. A number of American psychologists and economists developed an empirical science of human happiness and life satisfaction using these methods (Diener & Seligman, 2004). The World Happiness Report question that asks participants to evaluate their life from zero to ten is based on a question developed during this era, the Cantril Ladder question (Helliwell et al., 2025; Cantril, 1965). A related work is the Life Satisfaction Index by Bernice Neugarten et al. (1961), which asks an expanded range of questions, but still looks at an individual's experience from their own perspective. These works began a significant movement of research which aggregated large numbers of responses to questionnaires like these, and used them as indicators of societal wellbeing (Diener & Seligman, 2004).

This foundational psychological and economic research using social surveys to measure societal wellbeing evolved to become the heart of the academic study of wellbeing for much of the last four decades – the field of Subjective Wellbeing (SWB). “Subjective” relates to wellbeing being ultimately about one's own perspective (Diener, 1984). It is defined as a specific measurable concept consisting of life satisfaction and affect (Diener, 1984; Diener et al., 1999). Life satisfaction captures a self-evaluation of life. Affect, both positive and negative, concerns the moment-to-moment experience and is closely related to emotions, feelings and moods. Affect is the most directly hedonic component of SWB, because it operationalises wellbeing in terms of positive and negative felt experience (Diener, 1984; Ryan & Deci, 2001). Evaluations of life may be considered through hedonic and eudaimonic lenses (Ryan & Deci, 2001). Under this definition, someone who evaluates their life highly, and has a positive balance of emotional affect, has high SWB.

SWB is frequently assessed through social survey questions – e.g. by asking individuals about their satisfaction with life, and by asking individuals to recall their various emotions. In the core 2025 OECD guidelines, ‘affect’ is part of the core module and has an additional extended module of its own. It recommends asking participants about the degree to which they experienced eight emotions yesterday, on a scale of zero to ten, and also asks how much they laughed or smiled. A life satisfaction question “Overall, how satisfied are you with life as a whole these days?” (p. 32) is also part of the core OECD module, and in an extended module, ten different dimensions of life from safety to living standards to personal relationships and community are discussed. It should be noted that the recent focus on wellbeing being assessed through this multitude of questions (as is also observed in the 72-item Global Flourishing Study questionnaire (VanderWeele et al., 2025)), reflects a view that wellbeing is multidimensional and its assessment should ask about the different dimensions of life and

experience. Other studies also mention de-emphasising life satisfaction as a singular indicator of wellbeing and widening the focus to include other measures (Krys et al., 2025).

Over time, the SWB framework became increasingly influential to the point that it is one of the most widely used and essential parts of wellbeing assessments for researchers and governments alike, to this day. The Life Satisfaction aspect of SWB remains widespread as one of the most common questions with close to 90% of OECD member states collecting life satisfaction data in nationally representative household surveys (OECD, 2025).

Experience sampling and a critique of infrequent point-in-time assessments:

Although, questionnaires that collect data through periodic social surveys are vulnerable to issues related to the fact that the one answering the question is undergoing constant change. Their response can depend on the mood when they are asked (mood bias), and they also may not correctly remember how they felt when asked to recall the past (recollection bias) (Shiffman et al., 2008). Kahneman's "peak-end rule" helps explain why this matters for life-evaluation measures. When people remember an experience, their evaluation is disproportionately shaped by its most intense moment and by how it ended, rather than by the total quality of experience across its duration (Kahneman, 2011). For example, if asked "How was your holiday?", someone may answer based largely on a beautiful final day or a stressful airport delay, even if most of the holiday was different. This reflects Kahneman's distinction between the "remembering self" and the "experiencing self". The remembering self gives the retrospective evaluation, but it does not necessarily represent the moment-to-moment experience of the experiencing self. If wellbeing assessment aims to capture the overall quality of conscious experience, then asking only for retrospective self-evaluation is clearly limited.

The recognition of this has resulted in periodic surveys being seen with added caution as a complete representation of wellbeing, and has led to other assessment methods emerging. One is 'experience sampling', in particular Ecological Momentary Assessment (EMA). EMA is a technique that samples mood and emotion at various moments throughout the day (Trull & Ebner-Priemer, 2009; Shiffman, 2008). It is favoured as a way of getting high-frequency wellbeing data in natural contexts, while people are going about their daily lives. Experience sampling is now widely used because of its mitigation of recollection bias – an established issue in periodic survey methods where participants incorrectly recall their emotions or wellbeing (Shiffman et al., 2008) – and because of its greater accessibility and cost effectiveness enabled by technology. Nowadays, experience sampling is often done through pop-up surveys on smart devices (Wrzus & Neubauer, 2023). However, EMA remains vulnerable to other response biases, such as when the act of responding to the survey itself alters the mood of the participant, thereby skewing the results. It may also skew towards those who get value from the task, and therefore self-select for participation or retention in long-term EMA studies (Wrzus & Neubauer, 2023). EMA and similar experience-sampling methods are predecessors of the methods that may be used to understand wellbeing in the future, as they share the ability to capture more frequent, ecological, and contextual data (discussed in Part 2).

2.4 Mental health

Population mental health measures are one of the major recent additions to wellbeing assessment frameworks. Mental health issues are now widely recognised and discussed in global society – the WHO estimates that over a billion people were living with a mental disorder in 2021, with anxiety and depressive disorders being the most common (World Health Organization, 2025). Mental health assessment is now reflected in measurement practice internationally. New Zealand's wellbeing

statistics draw on the WHO-5 mental wellbeing index, while the UAE's National Strategy for Wellbeing 2031 includes promoting good mental health (Stats NZ, 2024; UAE Cabinet, 2019). The Global Flourishing Study also includes mental and physical health as one of its domains in a 22-country study of flourishing (VanderWeele et al., 2025). India and China both have National Mental Health Surveys which use clinical mental health concepts such as depression, anxiety and other psychiatric disorders, showing that these are not only Western measurement categories (Huang et al., 2019; Murthy, 2017). The OECD includes a population mental health module using the PHQ-4 and WHO-5, while the UNECE Guidelines include health indicators for anxiety/depression symptoms and suicide, alcohol-abuse and drug-overdose deaths (OECD, 2025; UNECE, 2025).

While mental health has traditionally been understood through the presence of *negative* states, symptoms, and poor functioning such as those outlined in the Fifth Diagnostic and Statistical Manual of Mental Disorders (American Psychiatric Association, 2022), there has been a notable shift towards the recognition of mental health existing outside disorders. As in the WHO-5 and PHQ-4, anxiety and depression, two of the most common mental health issues, are increasingly measured on a scale where all people may be placed (World Health Organization, 2024). While the prevalence of clinical mental health diagnoses still have relevance to understanding societal wellbeing, the ability for assessment systems to account for positive states broadens relevance to not only those with diagnosable disorders.

2.5 Objective wellbeing measures

Many societal wellbeing assessment systems also include “objective” measures, which exist as facts regardless of someone's perception of them, such as poverty levels, home ownership, education attainment, disease prevalence and other verifiable measures of living standards. Weijers and Morrison (2018) note that objective measures of wellbeing have been used for a long time and that their importance is “well understood by policymakers” (p. 6). The UNECE 2025 guidelines also show this within its broader framework. They present wellbeing as multidimensional and include both objective and subjective indicators of people's circumstances. This means wellbeing is not measured only through income, health, or material conditions, but also through how people evaluate their lives.

Despite its prevalence, a common critique of focusing overly on objective measures is that, as Weijers and Morrison (2018) puts it, “it might result in resources being directed towards something that was perceived to be a contributor to wellbeing, but actually makes little difference to how people themselves evaluate their own lives” (p. 6). In other words it neglects the subjective experience of individuals. A stronger view in line with Diener's view of subjective wellbeing, sees objective measures as secondary proxies to wellbeing, and that they should not be the main measure of wellbeing in governance (Layard, 2005). Under this view, objective measures are merely facts about a person or a person's circumstances, but do not represent their wellbeing directly – objective indicators are seen as only relevant insofar as they have an impact on the subjective experience of individuals.

But Sen (1992) points out an issue with complete expulsion of objective measures. Someone may be materially impoverished but feel happy through ignorance or coercion. Objective measures can get around the issues of someone coercing another into a contented feeling under poor conditions, such as the so-called “happy slave”. In this case, someone is enslaved or impoverished but for one reason or another they have been coerced into happiness by their oppressors; or they are naively happy out of ignorance of how others are living³. This argument is part of the justification for the ‘objective list

³ Setting aside the fact that in reality, a slave would likely comparatively see their life as worse than their masters, and so would be made unhappy by that.

theory' being considered another major contemporary theory of wellbeing assessment (Crisp, 2026; Fletcher, 2016).

2.6 Summary

This section has outlined several of the major theories and measurement traditions that shape contemporary wellbeing assessment: hedonic and eudaimonic wellbeing, utilitarian calculation, subjective wellbeing, life satisfaction, emotion, mental health, and objective indicators. Together, they show key aspects of how wellbeing has come to be defined, measured and standardised in international settings. They also show that wellbeing assessment is not conceptually neutral. Assessment systems carry assumptions about what wellbeing is and how it can be understood. The next section looks beneath a description of the frameworks to uncover the central assumptions that underlie these approaches.

3. The individualistic and humanistic assumptions

One assumption central to these major philosophical theories of wellbeing, underlying hedonic, eudaimonic, utilitarian, and modern psychological views of wellbeing, is the underlying notion of the individual self. Hedonic wellbeing, associated with pleasure and pain, and eudaimonic wellbeing, associated with meaning and virtue, both place the self as the fundamental unit. Whether it is experiencing positive or negative emotion, or it is virtuous, known, and has a sense of meaning and purpose, it is the individual self which is at the centre of both of these wellbeing definitions.

Utilitarianism is commonly enacted through measurement systems which focus on individuals' subjective experience as the units that make up society. Bentham described the community as a “fictitious body” made up of individual persons, and that the interest of the community is the sum of the interests of its individual members (Bentham, 1789/1907, Ch. 1, paras. IV–V). One of the main focuses of the Subjective Wellbeing approach, is that it places the individual self-evaluation at its core. It is invariably measured through self-evaluations that are aggregated to understand societal wellbeing. The UNECE and OECD frameworks and others largely share these views. They ask about individual perspectives on personal wellbeing and aggregate these responses to give evaluations of society. This section examines whether this is a universally applicable way of viewing wellbeing, and whether it accurately represents views across the cross-cultural settings in which they are applied.

3.1 Dominant psychological views of self

As discussed, the field of psychology has been foundational in much of the existing academic work on wellbeing assessment, especially through the development of subjective wellbeing, and through earlier waves of positive psychology, humanistic psychology and psychotherapy (Diener, 1984; Lomas, 2022a; Ryan & Deci, 2001). It is therefore useful to examine the field as a source of foundational views of self embedded in contemporary wellbeing assessment frameworks.

Lomas (2023a) describes the dominant view of self present in psychology using a metaphor of a ‘container’, where a container is something with boundaries, an inside and an outside, and contents held within it. He argues that this is supported by four recurring dualities in psychology. The first is mind and body, where mental phenomena and physical phenomena are treated as different and separate things. The second is subjective and objective, where first-person experience is separated from observable, material processes. The third is self and other, where the individual person is distinguished from other people and from the wider environment. The fourth is inner and outer, where subjective experience is treated as occurring “inside” the person, while the world is treated as existing “outside” them. Applied to psychology, this means that the person is often treated as a bounded individual self, with a mind, and subjective experiences located “inside” them. Other people, the world, and the environment are then treated as external things that may influence a person, but remain outside the person’s basic boundary.

Joshanloo and Weijers (2024) more directly link this with wellbeing discourse. They analysed some of the themes present in a number of influential models of wellbeing, including Diener's view of Subjective Wellbeing, and Ryff’s view of psychological wellbeing⁴. Several of the important themes they point out across these models include the importance placed on (1) “expressing one's true inner

⁴ Ryff’s psychological wellbeing (1989) is another major early model of wellbeing.

self”, (2) having a “positive self-evaluation and sense of worth based on one’s own internal standards”, (3) having “personal goals that give a sense of purpose and direction in life”, and (4) “building and nurturing fulfilling connections with others while also maintaining their individuality and personal boundaries” (p. 8). These models are summarised as focussed authenticity, self-determination, and self-expression. Joshanloo (2026) also supports this idea that wellbeing discourse is based around an independent view of self. He states that “widely used well-being constructs... are largely organized around an individualistic and independent construal of selfhood” (p. 1). He goes on to say that this construal is not universal and is biased towards Western views of self. Under this view of self, it makes sense for personal self-evaluations of wellbeing to rise to prominence. If wellbeing is located primarily inside bounded individuals, then asking individuals to evaluate themselves, and then aggregating those evaluations, appears to be the most direct way to assess societal wellbeing.

The dominant contemporary methodology of measuring wellbeing, like those present in the OECD and UNECE guidelines, centres on self-evaluations which implicitly value personal perception of self as being at the core of wellbeing assessment. It has also been a driving force behind the view that asking individuals directly for their own view is one of the most direct methods of understanding wellbeing.

3.2 Alternative views to the bounded individual as the central unit of collective wellbeing

An alternative view to the bounded individual as a core unit for understanding wellbeing, sees the self as not an independent autonomous entity. Instead it sees them as formed through relationships with others, their community, society and environment. In this view, societal wellbeing cannot be fully understood by assessing individuals as if their experience were contained inside them and then aggregating the results. The relevant unit is not only the individual, but also the relations and conditions through which experience arises. This section surveys several views which challenge the dominant notion of the individual self as the fundamental conceptual unit, starting with the idea of interconnectedness and the individual-in-relation.

Interconnectedness and individuals-in-relation:

Lomas contrasts the individual as a container with the metaphor of the person as a ‘field’ (Lomas, 2023a) – a field of relations that shouldn’t be perceived to have strict physical boundaries but instead “extends out into the world” (p. 1). In this sense, the individual is defined not only through self-perception, but through relationships with others and their environment. Relationships are viewed as a continuous part of the same whole, rather than an influence on an individual and their wellbeing.

Consider the following analogy to explain the difference between the self as independent of others (dominant self-other framing) and one where they are defined by relationship with others (alternative individual-in-relation framing). Imagine a ‘self-other’ framing to be analogous to a pool table of balls. Each individual is a ball. The balls can influence each other by knocking into each other, changing their state (analogous to wellbeing). The state of the collection of balls can be understood by inspecting each ball individually. Alternatively, in a more relational wellbeing framing, the group could be considered more like a spider's web. Individuals are the joins in the web. There is no clear point where each individual starts and ends, and they are defined by their connections to other joins. Each is inextricably connected and influencing each other. It would be impossible to understand the

wellbeing of the web by looking only at the joins. Instead you would have to look at the connections to understand the wellbeing of the web as a whole.

This ‘individual-in-relation’ framing could be seen as more in line with the view of those in collectivist cultures, affirming the view that they are defined through their relational role. In a traditionally collectivist view, the sense of self may emerge through their relationship to others (Joshnloo & Weijers, 2024; Joshnloo, 2026) – as a child, a friend, or as a member of a community, profession, nation, or religion. Under the ‘individual-in-relation’ view it may not be constructive or accurate to consider oneself as purely self-determined, or needing to look inwards to be in touch with one’s true inner self. It might be incorrect to maintain independent individuality and strict boundaries with others. Instead the expression and determination of selfhood may emerge through relationships. Relationships and connections may not be just an influence on an individual, but a continuation of self.

In views that consider the individual to emerge out of relationship with others, personal self-evaluation may be an insufficient method of fully understanding wellbeing. To understand an individual it may be necessary to look not just at the effect of relationships on an individual’s wellbeing, but also at the wellbeing of the connections themselves.

It should be noted that the dominant Western psychological viewpoint does not necessarily discount the importance of connections or relationships and, in fact, they are often described as of utmost importance (Diener and Seligman, 2002; Waldinger, 2015). But relationships and social connection are nevertheless usually considered valuable because they *lead to* individual wellbeing – rather than something of value in themselves.

Extended mind, extended self:

Clark and Chalmers’s extended mind theory is relevant here as it asks where the mind stops and the world begins, and argues that cognition is sometimes made up of neural, bodily and environmental processes, rather than being only “in the head” (Clark & Chalmers, 1998, p. 8). They argue against the idea that mental states and experiences are bounded internally within a subject who is acted upon by an external world. Instead, they suggest that at least some mental processes are partly constituted by the wider environment in which they occur. This parallels the challenge provided above, where Lomas asks where the person ends and begins, and argues for a recognition of the person as porous and interconnected. Both use the phrasing “extends into the world” – Lomas when discussing the person and Clark and Chalmers when discussing the mind (Lomas, 2023a, p. 1; Clark & Chalmers, 1998, p. 12).

Belk’s theory of the extended self provides a further parallel, arguing that the self is partly constituted through what people experience as “mine”. While Belk’s focus is at the level of identity rather than mind, it similarly challenges a strictly personal view of selfhood. He argues that the self is not defined only by a private core, but is partly constituted through relationships, possessions, and places (Belk, 1988). There are also older views that corroborate this – see David Hume’s ‘Bundle Theory of Self’ (Hume, 1739/2000, Book 1, Part 4, Section 6).

Applied to wellbeing assessment, the extended mind and extended self ideas further weaken the assumption that wellbeing is best understood by asking isolated individuals to self-evaluate and then aggregating those evaluations. If aspects of mind and self extend into relationships and environments, then wellbeing may also be partly located in those relations and connections. This does not mean self-evaluation is irrelevant, but it suggests that it may be incomplete when treated as the central basis

for assessing collective wellbeing. It is not that the relationships or environment are a feature of individuals or influence them, it's that they *constitute* them – they *are* them.

A Buddhist phenomenological perspective of non-self and dependent origination:

To add to this, we will briefly discuss two Buddhist ontological principles which take the idea of an individual-in-relation a step further.

In Buddhism, one of the three 'marks of existence' is *Anattā* (non-self) (Encyclopaedia Britannica, n.d.). It posits that there is no independent or self-sufficient essence. While Buddhist sources still use the linguistic functions of "I" and "you", indicating some normative existence of a person as a useful linguistic unit, *Anattā* criticises the idea of the self as an independent or permanent entity. A second related Buddhist principle, on the 'fundamental nature of reality', *Pratītyasamutpāda* (dependent origination), more broadly states that all phenomena arise in dependence upon other phenomena; nothing exists in isolation, and nothing has an independent essence (Mejor, 1998). Under this view the self, along with all other phenomena, are inextricably linked with other selves, and therefore can not accurately be viewed independently of them.

These Buddhist ideas therefore support the earlier challenges offered by Lomas's field metaphor, Clark and Chalmers's extended mind, and Belk's extended self. Each questions the strict boundary between person, others and world. Buddhist non-self and dependent origination go further by questioning whether the self has any independent essence apart from the relations and conditions that constitute it. Applied to wellbeing assessment, this challenges the assumptions that wellbeing is primarily located inside separate individuals and best understood by aggregating their self-evaluations. A Buddhist view may still allow a conventional person to report how they feel, how satisfied they are, or how they understand their life at a given moment. But they may miss the primacy of relational elements of wellbeing.

Aside from these Buddhist principles calling into question the centrality of a bounded self, they also question the validity of the self being stable and continuous over time. Huang (2022) suggests the person is understood as a composite of interdependent processes (*Khandas*) which are continually undergoing change (*Anicca*). He suggests these processes can be considered "a stream of moments of consciousness" (p. 1). Under this view, wellbeing may be best viewed as the quality of this stream of consciousness undergoing continual change, and not assigned to a stable or continuous self. This further supports Kahneman's identification of the experiencing self undergoing constant change, and supports the rationale for experience sampling as a method of understanding wellbeing, rather than an assessment system based on infrequent, single-point-in-time assessments that may provide only a crude snapshot of wellbeing. This suggests that wellbeing assessments should be cautious treating any single self-evaluation as a complete representation of a person's wellbeing across time.

An indigenous Māori view:

In *Māori Philosophy: Indigenous views from Aotearoa* (2020), Georgina Stewart describes a conception of self in the Māori worldview that is interwoven with land, environment as well as ancestors. Rāketē (2017) provides a vivid articulation of this view in which she presents humans as an emergent expression of a wider ecological whole, with no discrete line between where humans end and nature begins. Through food, the molecules of the body are constantly exchanged with plants, animals, fungi, and minerals. Every part of a person was once the environment, and every part will return to it. The body is mostly water, which is drawn from rainfall, rivers, lakes, and aquifers, and will eventually return there. Every breath exchanges air with trees, wind, and the ocean. At any

moment there is no clear boundary where a person's body or breath ends and the world's earth, wind or water begins. Therefore, the individual is considered materially indistinguishable from the atmosphere, the soil, the hydrosphere and the cosmos. Furthermore, this can be seen more widely in the Māori world, through a widespread saying and through national legislation. The saying “I am the river and the river is me”, or to grant legal ‘personhood’ to a river (e.g. Te Awa Tupua (Whanganui River Claims Settlement) Act, 2017), are different ways of drawing equivalence between people and environment. In such thinking, it could be considered incoherent or at least incomplete to give an assessment of individual wellbeing separate from the environment and other living beings.

A Māori view of self is also deeply contingent on *whakapapa*, which is roughly equivalent to genealogy. The self is described as having a living and present relationship with ancestors, land, and other beings (Stewart, 2020). The ancestral earth is therefore not merely where people came from historically, but it remains part of who people are now. Māori self-introductions usually consist of naming rivers, mountains, *iwi* (wider tribe), *hapū* (community), *whānau* (family), and *whakapapa*, which further represents their role in the definition of self.

Progressing yet unresolved assessment implications:

To summarise, it is clear that the idea of the self being independent, stable, and separate from others is not a universally held view. Conceptions of the self are subject to significant cross-cultural variation, with some cultures preferring to define it through their connections, relationships, environment, or ancestors, rather than through self-definition. This variation brings into question the central focus on self-report measures of wellbeing. It may be beneficial to consider giving more weight to relationally- or environmentally-oriented views of wellbeing measures of wellbeing as primary measures of societal wellbeing.

Some contemporary research has begun to move partly in this direction. The Global Flourishing Study (GFS), for example, includes two of its eight domains as ‘close social relationships’ and ‘community and social connectedness’ (VanderWeele et al., 2025). In fact, many social surveys often already include measures of society-level connection, and these are becoming more prevalent in international surveys. In a New Zealand context, The Social Cohesion in New Zealand study similarly directs attention beyond the individual self, measuring people's relationships to others, groups, institutions, and society through dimensions such as belonging, participation, social inclusion, and trust (Eaqub & Collins, 2025).

These begin to show a practical movement beyond life satisfaction, affect, or clinical mental health measures by asking about relations, connections, community, and society through which wellbeing is formed. But it's important to note that self-report measures of relations are often important insofar as they contribute to individual wellbeing. Individual wellbeing is then aggregated to understand societal wellbeing. In other words, the connections are subordinate to the individual. The question of how to best capture societal wellbeing through a relational lens remains an open question, although considering relationships as fundamental to societal assessment frameworks could be a way of accommodating less individualistic ontologies. This points to a broader argument present in the thesis about the need to include non-individualistic, cross-cultural ideas in the foundational discussion of wellbeing assessment.

3.3 Cross-culturalism as a principle in global wellbeing assessment

If wellbeing assessment frameworks are intended for use across global and culturally diverse settings, then their legitimacy depends in part on whether they incorporate cross-cultural concepts. This does

not mean that cultural or philosophical tradition should be incorporated uncritically, or that Western scientific methods should be rejected because they are Western. But it should make us be particularly open to reorientation based on fundamentally different epistemologies about wellbeing. If global wellbeing frameworks remain built primarily from Western psychological, philosophical, and policy traditions, without serious engagement with alternative views of self, experience, and morality, they are unlikely to have widespread support and use in global settings. Either that, or they risk imposing a form of *wellbeing imperialism* where Western academics are defining the terms on which we assess the societal wellbeing of countries around the world. It is therefore not merely an optional gesture of diversity, but seems to be necessary for the application of these frameworks in international spaces. Buddhist ideas of non-self and dependent origination, while not fully developed foundations for wellbeing assessment systems (i.e. questions remain about how this impacts actual assessment frameworks), are useful because they expose assumptions within the dominant model and help widen the conceptual basis from which global wellbeing assessment might proceed. Put simply, foundational concepts underpinning global wellbeing assessment frameworks should adhere to a principle of cross-cultural universality. While such frameworks do not need to represent a complete or universally agreed model of wellbeing, they should be grounded in concepts that are meaningful across the cultural contexts in which they are applied, and should mitigate any apparent cultural bias.

At this point it is briefly worth noting the inadequacy or oversimplification of broad cultural generalisations. When distinctions are made in this thesis between “the West” and “the East” or “indigenous”, it should be remembered that there is significant geographic and cultural diversity within these categories. For example, empirical evidence from the Gallup World Poll showed the characterisation of the East as collectivist and the West as individualistic is an oversimplification (Lomas et al., 2023b). On average, Eastern participants scored higher than Western participants in self-care when asked “Do you think people should focus more on taking care of themselves or on taking care of others?” (p. 1). Nevertheless, it remains a useful enough distinction due to the historically divergent cultural histories (Lomas, 2025), to use in the context of this thesis if the internal diversity and complexity of such terms is kept in mind. It is also possible that in an increasingly globalised world, with a greater degree of mixing and cross-pollination through the internet, social media, education and politics, cultures are becoming less segregated by geography or history.

Indigenous and “bottom-up” methods of assessing wellbeing, that involves engaging local communities in the design of the assessment of wellbeing are also discussed by Kryszewski et al. (2025) as being important going forward. This is beginning to be recognised in international frameworks, with the inclusion of a greater degree of indigenous and cross-cultural perspectives being one of the major reasons the OECD updated their Guidelines on Measuring Subjective Well-being in 2025 (OECD, 2025). Smith et al. (2025) were commissioned by the OECD to review and put forward views from minority and indigenous communities and non-OECD countries.

3.4 Beyond humans

In section 3.2, this thesis challenged the consideration of the bounded individual person being the appropriate unit of collective wellbeing assessments, and explored the role relationships and connections might have in understanding overall wellbeing. This section will continue exploring the appropriate unit of wellbeing assessment by asking the questions: What is capable of experiencing wellbeing? Is it limited to humans? Is it limited to animals? Is the environment a contributor to wellbeing or capable of experiencing wellbeing in itself? This will help us better understand whose wellbeing matters morally to conceptually guide future wellbeing assessment frameworks.

Conscious experience beyond humans:

Before continuing, it is useful to clarify what this thesis means by conscious experience. Conscious experience, or simply experience, here means any experiential or qualitative dimension at all: that there is something it is like (Nagel, 1974).

Colour is a simple and familiar example. The physical properties of the colour red can be described in terms of wavelengths of light. So can the biology of how those wavelengths enter the eye, are processed by specialised cells in the retina, travel through the optic nerve, and are interpreted by regions of the brain. But even a complete physical description does not capture the felt “redness” of red. The redness is not merely the wavelength or the neural pathway, it is the qualitative experience of red⁵. This is part of conscious experience. Likewise the multitude of sensations, feelings, perceptions, mental formations, and other qualitative experiences contribute to our overall conscious experience.

Many people readily accept that dogs have experiences because they display emotions that resemble our own, such as joy, sadness, jealousy, or guilt. However, experience need not resemble human experience to exist. Nagel (1974) uses the example of a bat to illustrate this point. A bat's experience of the world is shaped in part by echolocation, a mode of perception fundamentally different from our own. Although we cannot know what echolocation feels like from a bat's perspective, its unfamiliarity does not make it any less experiential. Experience does not require human-like emotions, identity, or a sense of self. These may be features of human experience, while echolocation and other bat-specific experiences may be features of a bat's. The key point Nagel makes is that conscious experience can plausibly take forms very different from our own. According to the same logic, even simpler or more ‘distant’ forms of life may be capable of experience. Nagel states “conscious experience is a widespread phenomenon. It occurs at many levels of animal life...” (p. 1). If we believe a dog experiences guilt, and a bat may have some experience associated with echolocation, it may even be possible to comprehend that insects may have some form of experience.

According to the definition in this thesis, where there is conscious experience, there is wellbeing. The wellbeing of an individual insect may not be particularly morally consequential, but if we consider the scale of simpler forms of life like this they may be worth considering as a collective.

All sentient beings:

A contemporary utilitarian perspective from the philosopher Peter Singer, in *Animal Liberation* (1975), argues that we should have moral concern for the suffering of animals because they are capable of having good and bad conscious experiences. For Singer, if an animal can have experience (like pain, suffering, pleasure, comfort), even if it can't express it through speech or language like people can, then it should still be part of our sphere of concern. Its intelligence does not matter, but rather that it is capable of conscious experience. While Singer's view does not draw an equivalence between the experience of animals and people, it does advocate for a recognition of non-human animals as capable of experiencing wellbeing.

In animal welfare science, sentience is often used as the term which denotes an animal's capacity for conscious experience (Mellor, 2019; New Zealand Veterinary Association, 2018)⁶. Recognition of an

⁵ Frank Jackson (1982) famously discussed a similar thought experiment about a scientist who knows about colour but has never experienced it first hand.

⁶ more specifically: “consciously perceiving, feeling, and subjectively experiencing negative and positive states” (Mellor, 2019), or “sentience is the ability to feel, perceive or experience subjectively” (New Zealand Veterinary Association, 2018).

animal's capacity for sentience is not confined to Singer's philosophical argument, it is increasingly reflected in law and public opinion. New Zealand's Animal Welfare Amendment Act (No 2) (2015) amended the purpose of the Animal Welfare Act 1999, which forms part of the country's animal welfare legislation, "to recognise that animals are sentient" (para. 4) and to require people responsible for animals to attend to their welfare. The UK Animal Welfare (Sentience) Act 2022 similarly recognises animals as sentient beings and establishes an Animal Sentience Committee to assess whether government policy has considered animal welfare impacts. Blattner (2019) argues that many legal systems protect animals specifically because, and to the extent that, they are sentient, making sentience a central legal basis for animal protection rather than merely a philosophical idea (Animal Welfare Amendment Act (No 2) 2015; Animal Welfare (Sentience) Act 2022; Blattner, 2019).

There is also evidence that public belief in animal sentience extends beyond Western countries. Mata et al. (2022) surveyed Brazil, Russia, India, China, and the United States and found that, in every country surveyed, more respondents agreed than disagreed that farm animals are sentient. This matters because the study included major non-Western countries and found cross-national variation, but not a simple Western-only recognition of animal sentience.

Animal welfare science also supports this connection between sentience and wellbeing. The Five Domains model, a leading model of assessing animal welfare (SPCA, n.d.), uses measures from the four domains of nutrition, environment, health, and behaviour as important only insofar as they give rise to the animal's subjective mental state (Mellor et al., 2020; Littlewood, 2023). Littlewood et al. (2023), in summarising the 2020 Five Domains model states that the model "focuses on subjective mental experiences that matter to the animal" (Section 2). In this sense, animal welfare science is directly concerned with animal wellbeing as experienced by the animal. Freire and Nicol's bibliometric analysis (2019) found that animal welfare science grew at 13.3% annually over 50 years, showing that this is an expanding research field. The Effective Altruism movement's recognition of non-human animals' capacity for sentience and wellbeing provides a contemporary example of animal wellbeing being taken seriously at scale, with tens of millions of dollars being allocated towards animal welfare causes (EA Grants Database, n.d.). The Wild Animal Initiative (WAI) is a small academic community which uses a similar logic to justify understanding and improving the welfare of wild animals at scale (The Wild Animal Initiative, n.d.).

The contemporary international growth of the recognition of non-human animal sentience and concern for their wellbeing together with Singer's version of an inclusive wellbeing-based moral philosophy, is a perspective that cause us to question an assumption present in modern wellbeing assessment systems used in policy that we should only measure human wellbeing. It is worth considering the validity of a broader and more inclusive view in wellbeing assessment frameworks that considers all those capable of sentience and conscious experience, even if less of a priority than human wellbeing.

3.5 Expansive consciousness and expansive wellbeing

Once wellbeing is defined as the quality of conscious experience, and once conscious experience is detached from human language, self-report, and a human-like sense of self, the relevant question becomes not only *who* reports wellbeing, but where, if anywhere, the capacity for conscious experience ends. This section considers the more speculative possibility that conscious experience may be more widespread than dominant human-centred or brain-centred models usually allow. The argument is not that plants, ecosystems, rivers, mountains, or the universe are conscious in the same way humans or animals are conscious. Nor is it that they have interests equivalent to human or animal interests. The more limited claim is that, if wellbeing is grounded in conscious experience, then

wellbeing theory should avoid assuming too quickly that only individual human persons can be subjects of wellbeing.

It is a mainstream scientific view to treat conscious experience as arising from a physical basis, especially processes of the brain and nervous systems (Chalmers, 1995, p. 10; Goff, 2026). This view has considerable plausibility when we look at our own conscious experience as the starting point. Our nervous systems are clearly connected to our experiences of emotion, perception, and thought. Therefore, the logic goes that other animals with brains and nervous systems, like dogs and bats and possibly insects, likely also have some form of conscious experience, with this becoming increasingly uncertain as organisms become less similar to us. This traditional picture does not need to be rejected entirely. It is reasonable to think brains and nervous systems enable complex conscious experiences. However, it may be too restrictive to consider only organisms with brains and nervous systems as capable of any experience at all. We may consider the possibility that conscious experience is not only made up of rich and complex experiences like a sense of self, but made from simpler forms of experience that culminate in these richer experiences. If this is the case it's possible the building blocks of conscious experience may be much more widespread beyond people, animals, into other forms of life or beyond.

In a later essay on "Panpsychism" (1979), Nagel extended his point from 1974 that "conscious experience... occurs at many levels of animal life..." (p. 1), to state that there was no convincing reason conscious experience emerged at a certain level of complexity. That is to say he believes it makes no sense that conscious experience emerges out of non-conscious matter at the biological complexity of animals, and therefore conscious experience may be widespread in the universe. This view is becoming a more common view in philosophy (Goff, 2026). It takes many different forms but a number agree with Nagel that some extremely basic building block of what gives rise to our complex conscious experience (sometimes called "proto-consciousness" (Chalmers, 2016)) exists in the fundamental building blocks of the universe (Goff, 2017; Chalmers, 1995; Harris, 2019; Faggini, via Busstra, 2025). This is a speculative view, but if true, would have wide-reaching implications for what is capable of wellbeing, potentially extending widely to all things. Lomas (2022a), also discusses the possibility of panpsychism and its implication for wellbeing science in his paper entitled: *Stranger than we can imagine: the possibility and potential significance of non-human forms of consciousness and wellbeing*. Lomas cites Nagel himself describing the notion of panpsychism as "the faintly sickening odor of something put together in the metaphysical laboratory", but nevertheless considered it a possibility (Lomas, 2022a, p. 18).

The implications for wellbeing assessment are far from certain but it does make discussions of environmental wellbeing or the wellbeing of non-animal entities *for their own sake* not immediately worth dismissing. Future global wellbeing assessment should at least remain open to the possibility that wellbeing extends beyond individual human persons, and perhaps beyond animals.

It should be noted that these views don't claim that rivers, houses, or trees have conscious experience like a person. Rather that they are made up of the same fundamental proto-consciousness. They may be part of a broader universal experience that exists in some form in all things. The specific arrangement of proto-consciousness in biological systems, brains, and nervous systems still gives rise to our specific kind of conscious experience, but may not be the *only* kind.

Seeds of environment being recognised as having intrinsic value in international frameworks:

In parallel to the discussion of the plausibility of panpsychism in philosophical discussions of wellbeing, there is also a growing environmentalist movement influencing wellbeing assessment. Several leading global frameworks appear to reflect this shift by increasingly recognising the intrinsic value of the environment. The United Nations Economic Commission for Europe's (2025) guidelines for wellbeing measurement explicitly question the human-centricity of existing models, asking whether "a wider canvas, including, for example, nature and the environment" (p. 72) should be considered. While environmental domains are already present in many wellbeing frameworks, most existing measures that focus on the environment are generally justified in terms of their *instrumental* benefit to humans. This question in the UNECE guidelines, even if unanswered, may signal a deeper conceptual shift from viewing the environment as an instrumental contributor to human wellbeing to recognising it as intrinsically capable of it in its own right. The guidelines also acknowledge that different societies, and groups within them, may perceive the answer to this question in different ways.

This does not create a definitive answer to how future wellbeing assessment systems will change in practice, but it does turn our attention towards considering more holistic systems of measurement that include environmental wellbeing for its own sake. Additionally, it may be possible that this belief system progresses the discussion to a more cross-culturally acceptable space that is not focussed on Western individualism as the foundation of wellbeing assessment.

3.6 A brief note on theistic viewpoints

This thesis has explored a new metaphysical framing that challenges the humanistic and individualistic foundations of wellbeing. It has looked at how the main models of wellbeing often consider 'the self' to be the primary determinant of one's own wellbeing. It has explored several worldviews that consider an adjustment of this view to consider oneself as deeply interconnected and defined by others and the environment. These could be considered views where societal wellbeing is defined in terms that transcend the self.

There are also theistic views which consider wellbeing to be something that is not self-defined, but based on the judgement of God. This is not a marginal concern in the context of a system interested in global wellbeing assessment. Pew Research Center's 2025 analysis found that, across 35 surveyed countries, a median of 83% of adults said they believe in God, with belief especially high in sub-Saharan Africa, Latin America, South Asia, Southeast Asia and Turkey. By contrast, it is lower in China and other parts of East Asia, Europe, Australia and New Zealand. This matters for global wellbeing assessment because many of the academic and policy institutions shaping these frameworks are located in comparatively secular, WEIRD (Western, Educated, Industrialised, Rich, Democratic) contexts. If wellbeing frameworks are to claim legitimacy in global settings, they should at least make conceptual space for theistic views in which wellbeing is understood partly through alignment with God, ultimate reality, or transcendent purpose.

One paper by Joshanloo (2017) discusses a perspective on the Islamic view of wellbeing. He states that to a Muslim, wellbeing is an alignment of one's life with the higher power of God. To Muslims, "wellbeing is living a life in which all one's actions and intentions are organized around the principle of Allah's [God's] absolute sovereignty" (p. 13). While subjective states such as tranquility, contentedness, and balance are seen as an indication of alignment of life with Allah, they are rejected as pursuits in themselves. "Well-being is not defined on the basis of certain subjective feelings and

experiences, such as the intensity and frequency of certain experienced emotions and sensations” (p. 13). This represents a view where pursuing positive personal (or collective) subjective states is not the main ideal worth pursuing.

This is just one example of an aspect of a theistic view of wellbeing. Theistic views are underrepresented in international wellbeing measurement frameworks and may be worth paying more attention to going forward. The GFS is relevant here because it explicitly includes 15 out of 71 questions relating to religious or spiritual life, connection to God, nature, and spirituality (VanderWeele et al., 2025). This shows that contemporary wellbeing measurement is beginning to make space for dimensions that secular psychological models often underemphasised.

3.7 Summary

Section 3 has challenged the assumption that wellbeing should be understood primarily through bounded, individual human selves. It has argued that selfhood, experience, and wellbeing may be constituted through relationships, communities, environments, non-human animals, and questioned the possibility of wellbeing being more widespread. Several cultural and ontological assumptions built into dominant wellbeing frameworks were discussed, as was the possibility of widening the scope of what wellbeing assessments should be trying to capture.

4. Quantitative wellbeing measurement and maximisation

Beyond the conceptual discussion about what is capable of wellbeing, and what is the appropriate unit for assessment, another feature of these methods of assessing wellbeing in global frameworks remains contestable. Many of the dimensions discussed above are evaluated through ascending quantitative linear scales, often from 0–10, where the highest number is treated as best and the lowest number is worst. These systems are often implicitly or explicitly set up so that the aim is to maximise. Yet it seems this may not be a universal way of framing assessments.

Krys et al. (2025) evaluated two large-scale, cross-cultural studies and concluded that in many cultural contexts, moderate levels of life satisfaction are preferred to maximisation. They say maximisation is preferred in predominantly WEIRD societies and are likely not appropriate in global settings. In Krys et al.'s 66-country analysis, Germany and Iceland showed strong maximising ideals: 86% and 84% of participants respectively said ideal happiness was at least “very happy”. By contrast, in Bhutan, 81.5% placed ideal happiness below “very happy”, and 70% or more did so in Ghana, Nigeria, Pakistan, and Japan. This directly supports the claim that maximisation is not universally preferred.

Miyamoto et al. (2010) gives a possible explanation for why this is the case in Japan stating that high levels of personal happiness can be seen as disruptive to collective harmony. In a separate study on the phenomenon of ‘aversion to happiness’ Joshanloo & Weijers (2014) argue that some cultures are more likely to view excessive displays of happiness as a signal that someone is naive, superficial, or ignorant of the true nature of things. They point out the sentiment among some artists and creatives that consider suffering as a necessary element of their best creations and that overvaluing positive feelings or emotions could impede this process. For example, when Edvard Munch, the famously depressive painter of *The Scream*, was asked why he did not do something about his emotional ailments he said: “They are part of me and my art. They are indistinguishable from me, and it would destroy my art. I want to keep those sufferings” (Joshanloo & Weijers, 2014, p. 9). Although it should be noted that these viewpoints generally apply to happiness, positive emotion, or hedonic wellbeing in a more narrow sense. If we consider wellbeing to be the overall quality of experience, then there may be a greater degree of openness to high wellbeing.

4.1 *Wu wei* and pursuing maximum wellbeing

In contrast to a wellbeing-maximisation worldview, concepts within Daoism such as *wu wei* (無為) emphasise being in harmony with the natural course of things (Littlejohn, 2016). According to the Encyclopedia Britannica entry, *wu wei* connects with flourishing in that “all things in the universe—including all human beings—have their own natural course, which, if unimpeded, leads to flourishing” (Stefon, 2013, para. 2). *Wu wei* can be translated as ‘non-action’, although Chan (2025) says the principle of *wu wei* is not about doing specific things or not doing specific things; rather it encourages harmonious alignment with the natural cosmic order and being free from the ordinary disquieting or unease coming from desire.

This provides a different critique of wellbeing maximisation from the empirical critique above. Even if wellbeing is defined broadly as the overall quality of experience, rather than narrowly as happiness or positive affect, *wu wei* suggests that directly pursuing maximum wellbeing may itself become counterproductive. Excessive striving, optimisation, or intervention with neglect for the wider course

of the cosmos may disturb the very conditions through which wellbeing arises. In this respect, *wu wei* shares an affinity with Buddhist concerns about craving as a source of suffering, and with the Western “paradox of happiness”, where happiness may become less attainable when pursued directly as an explicit goal (Crisp, 2026; Goodman, 2022). These may share the perspective that it is not that high wellbeing is undesirable, but that relentlessly pursuing maximum wellbeing may not be an effective or coherent route to achieving it.

A *wu wei* perspective does not imply that governments should be indifferent to wellbeing, nor that policy should be inactive. The *Daodejing* itself discusses *wu wei* in political terms, where the ideal leader promotes a “flourishing rather than a stagnant kingdom” (Stefon, 2013, para. 4). The implication for wellbeing assessment is that governments should be cautious about treating simplified wellbeing metrics as targets to be maximised and should pay greater attention to the wider natural course of things.

4.2 Yin yang and positivity maxing

Wu wei challenges the direct pursuit of wellbeing, and *yin yang* raises an issue about the nature of wellbeing itself. *Yin yang* is about a natural balance between positive and negative. It expresses a worldview in which contrasting forces such as positive and negative, light and dark, activity and rest, or expansion and contraction are mutually dependent aspects of reality, with harmony arising through their continual interaction rather than the dominance of one over the other (Wang, n.d.; Hon, 2019). Wang describes *yin yang* as involving “interaction” between waxing and waning aspects of reality, and “a process of harmonization, ensuring a constant, dynamic balance of all things” (Wang, n.d., para. 1). Waxing, in this context, can be seen as the aspects of reality which are naturally coming into existence, and waning can be seen as the aspects of reality which are naturally coming to an end.

Applied cautiously to wellbeing assessment, this suggests that wellbeing may not be best understood as the maximisation of positive states and the elimination of negative ones, even when considered broadly and over the long term. It might involve an acceptance of both positive and negative, waxing and waning. This does not mean that suffering, depression, or pain are good, or that they should be preserved. Rather, it challenges a more basic positivity assumption: that wellbeing is equivalent to increasing the positive side of experience and reducing the negative side as far as possible. From a *yin yang* perspective, negative aspects of experience may not simply be defects to be removed, or just indicators of what needs addressing in order to reach positivity. They may be seen as integral parts of the conditions through which balance, harmony, and therefore flourishing become possible.

4.3 Alternative measurement ideas and qualitative assessment:

Maximising balance:

A response is to include balance and harmony measures directly. This has happened alongside a broader pluralisation of wellbeing measurement, in which frameworks increasingly place life satisfaction and positive affect alongside other dimensions rather than treating them as exhaustive measures of a good life. The Gallup/Global Wellbeing Initiative module and the OECD’s experimental question banks show how this is already beginning to occur, with several example questions including asking about participants’ levels of balance and harmony (Lomas et al., 2022; OECD, 2025). However, if balance is measured on an ascending scale, it may simply become another thing to maximise. Including balance and harmony is still a useful improvement over measuring only

happiness or life satisfaction, but it does not fully resolve the deeper conceptual issue. A framework can include balance and harmony while still interpreting them through the same maximising logic that is challenged through a different foundational worldview. The way of measuring and maximising these elements on these scales — e.g. person or group with “more balance” is better — may be paradoxical to the core of the philosophy.

Ideal levels:

Another possible response is to measure ideal levels as well as actual levels. Ideal-level questions could be incorporated into dashboard approaches, allowing indicators to show actual and preferred levels without implying that the highest possible score is always best. Krys et al. (2025) argue that while asking only about actual life satisfaction or happiness can imply that the highest possible score is ideal, asking respondents where they would ideally want to be on the scale allows for the possibility that moderate levels are preferred. This is useful because it avoids treating anything less than the maximum as automatically deficient.

Dashboards over rankings:

Another response is to reduce emphasis on single-score ranking of people or countries. Rankings such as “the happiest country in the world” are communicatively powerful, but they can reinforce the idea that wellbeing can be accurately measured with a single score. This is especially problematic if the underlying measure reinforces a kind of dominance or authority in wellbeing that privileges culturally biased measures. Dashboard approaches offer a way of resisting this collapse into a single score. As Weijers and Mukherjee (2016) argue in the New Zealand policy context, dashboards can inform policy while keeping dimensions visible side by side, rather than reducing metrics to a single index. This is especially useful where improvement in one dimension may come at the expense of another, or where societies reasonably disagree about which dimensions should be prioritised.

However, even dashboards and ideal-level questions remain imperfect. Some dimensions of wellbeing may be cyclical, or context-dependent, and therefore not reducible to a fixed optimum on a scale. And it still may not sufficiently accommodate for views that don’t see wellbeing elements as something which can be quantitatively measured at all.

Qualitative assessment:

Qualitative methods, such as conversation-based assessment techniques that capture the nuance and complexity of an individual's situation, may be genuinely promising ways of understanding wellbeing without reducing it to an oversimplified ranking or number. They can represent wellbeing as a changing pattern of meanings, relations, and values rather than only as a set of target levels.

One-on-one qualitative evaluations by trained professionals is a common method for doing so but this is cost prohibitive for most wellbeing assessment contexts. At the moment, simple linear scalar methods have been ways of ensuring comparability while maintaining cost effectiveness. Part 2 will speculate on ways future systems may include qualitative data sources or methods to capture greater nuance and depth. It may be explored as a way of capturing wellbeing beyond linear maximising scales, accommodating a greater range of cross-cultural views than prescriptive survey questions, and better understanding relationality..

Part 2: Technology and future wellbeing assessment systems

With the growing global institutional interest in wellbeing, there is not only a shift in the conception of measurement underway, but as society becomes more digital, there is also a drastic *technological* shift affecting how we assess wellbeing.

Part 1 argued that wellbeing assessment is contingent on the answers to conceptual questions about what wellbeing is, who or what can have it, and how individuals relate to communities, environments, and wider systems. Existing frameworks that measure wellbeing have made significant progress beyond narrow economic measures, such as GDP, but many still rely heavily on self-evaluation, life satisfaction, affect, and mental health. These measures are practically useful, but they do not exhaust the meaning of wellbeing to everyone. If wellbeing is conceptualised too narrowly, then even technologically sophisticated measurement systems may reproduce a limited account of what it means for a society to be doing well.

Part 2 will explore several examples of the ways technology is being used to understand mental health, emotions, and other aspects of wellbeing. They examine text-based wellbeing inference, the use of LLMs to analyse large bodies of language, and touch on some non-text-based methods using video and audio analysis, wearable and sensors, and the possibility of multimodal wellbeing inference systems. It will discuss technological methods, as a response to the shortcomings of survey-based methods, and will speculate on how this may evolve in the future as part of systems of societal wellbeing assessment at scale. In section 6 it moves on to discuss the drawbacks of such systems and the risks associated with its use such as bias, privacy, and institutional misuse. Section 7 looks at the challenge of aligning AI systems with wellbeing, and how the problem of overly narrow conceptions of wellbeing magnifies under conditions of rapid technological development.

5. Technological breakthroughs

Many lives are undergoing a shift towards greater use of digital technologies, with a surge in global use of smartphones, social media, and the internet. As a result, large amounts of new data are being created. Recent advances in algorithmic analysis of this data and widespread availability (Maslej et al., 2025) has resulted in an explosion of academic papers using these technologies in the last several years. This section will look at a select number of techniques that use these technologies to understand wellbeing elements for both groups and individuals. Many are at early development stages, but represent movement towards detecting increasingly nuanced aspects of wellbeing.

5.1 Algorithmic analysis of text inferring survey-based measures of wellbeing

Social surveys remain the foundation of most group-level wellbeing assessment systems, but they face practical limits. Large surveys are expensive, slow to coordinate, often conducted annually or less frequently, and may have difficulty reaching remote regions or smaller sub-populations (Murtin & Salomon-Ermel, 2024; Lambert et al., 2020). These limitations are what have encouraged interest in digital data sources that can provide more frequent, lower-cost, ecological indicators of wellbeing.

Ecological refers to the ability to capture some wellbeing elements as they occur naturally, rather than only through retrospective self-report.

One example is the use of Google search data to predict wellbeing survey results. The OECD's work on 'nowcasting' utilises a public database of anonymised and aggregated search trends made available to researchers, journalists and the general public by Google (Murtin and Salomon-Ermel, 2024). They applied machine-learning methods to Google Trends, alongside other macroeconomic data, to estimate national average monthly subjective wellbeing levels. They found that adding real-time data from Google Search trends improved their models' predictive performance of average national Subjective Wellbeing scores. Their final model reached an out-of-sample prediction of R^2 of 0.830. This means that the addition of this search data substantially improved that model's ability to accurately predict wellbeing scores when tested on held-out data (scores the model was not trained on). Because the search data is geo-located and is available rapidly, this can be used to understand population wellbeing with high resolution and speed, which is particularly valuable in times of crisis like during the COVID-19 pandemic.

There are a number of examples of researchers using public social media posts to estimate a variety of measures of group and societal wellbeing. Mangalik et al. (2024) use approximately one billion tweets from two million users to estimate the levels of depression and anxiety. Their model aimed to produce estimates of population mental health survey results of depression and anxiety, and was found to be moderately to strongly associated with available survey scores. Due to the tweets being geo-located, they were able to make estimations to the level of the county (provided there was sufficient data in that area). And because there were specific timestamps with each tweet, they could make estimates for each week, tracking changes over time. To achieve county-week level spatiotemporal resolution through survey measures is rare and expensive to do at scale, highlighting the advantage of digital methods. Müller and Bittermann (2025) also used social media posting to track the emotional climate and wellbeing of researchers using X (Twitter) posts. They also found their model tracked real world events, showing a drop in positively-toned posting following the start of the pandemic.

One of the best examples of using social media to infer wellbeing is a review by Sametoğlu et al. (2024). They analysed 45 studies which looked at many of the common, standardised measures discussed in section 2 such as positive and negative affect, life satisfaction, life evaluation, eudaimonic wellbeing, and mental health. Overall, of the 18 studies with convergent validities, they found that social media models correlated with survey-reported wellbeing at $r = 0.33$ for individual-level assessments and $r = 0.54$ for regional measures. This means that the analysis of a region's aggregated posts were better at predicting regional wellbeing scores, than an individual's posts were at predicting an individual's wellbeing score. This further supports their potential as methods of understanding wellbeing, although notably with greater validity inferring *collective* wellbeing.

Mayer et al. (2024) offers a clinical example, where clients' emotions are identified in psychotherapy dialogue. They used a machine-learning model that was provided with a large number of dialogues where emotions had been labelled by humans. The model is then trained to predict the presence of emotions when provided with new, unlabelled transcripts. Their best model performed slightly worse than human annotators, but the authors noted that ambiguity and disagreement was significant between human annotators too. This shows the emerging potential of how language models can infer emotional expression in therapeutic text.

The advantage of these digital data methods is therefore that they can cost-effectively increase frequency, reduce turnaround time, and expand coverage of wellbeing assessments where digital data are available. Whether using search trends or social media language, these approaches are usually validated by their ability to reproduce existing survey measures. They can estimate wellbeing-related outcomes at scale, but they do not necessarily interpret the deeper contextual meaning of a person's language, narrative, relationships, beliefs, desires or values. This limitation sets up the relevance of Large Language Models: they may offer a further step from detecting patterns associated with wellbeing scores toward analysing more complex, nuanced, and contextualised meanings in text.

5.2 LLMs enabling contextualised, flexible wellbeing inference

Large Language Models are relevant because they extend automated language analysis beyond many of the methods discussed above. Earlier approaches show that digital language and search behaviour can be useful proxies for survey-based wellbeing measures, but they are often constrained by the need to provide predefined constructs, labelled training data, and dictionary categories. LLMs do not remove the need for external validation, but they create the possibility of analysing open-ended language in a more contextualised way.

Mangalik et al. (2024), used weighted dictionary-based models to estimate depression and anxiety from social media language, which is powerful in its ability to scale at low cost, offer high frequency data, and respond to real world events. But basic in its ability to understand genuine complexity and ambiguity in text data. Its assessment of anxiety or depression is simply based on the frequency of particular words associated with them, and has no ability to assess the full meaning of a sentence or passage in context.

Machine learning models like those used take this further by using trained language models to identify emotions from psychotherapy dialogues (Mayer et al., 2024), but labelled examples need to be provided to train the model and its capability then remains specific to only that task. Likewise the OECD's 'nowcasting' (Murtin & Salomon-Ermel, 2024) system trained models to identify SWB scores from search trends, but it also requires extensive training data and leads to capabilities specific to that particular task. LLMs extend this capability further and are starting to be used to pick up contextualised meaning across a passage without prior training on human labels. They may allow the identification of meaning both flexibly and in context across large volumes of text.

First I will briefly describe how they do this, then give some examples of the ways this is being done. When analysing text for meaning, LLMs first divide text into tokens, which are words or parts of words represented as multidimensional numbers. Each token is not a static number, but it changes depending on the surrounding words and the wider passage. For example, "well" could be a hole in the ground with water in it, or a state of wellbeing depending on the surrounding text. Not only does this allow the capturing of different dictionary definitions of words, but since tokens are often hundreds or thousands of digits, it can represent even slight changes in meaning based on things said at a distant point in the text passage it's analysing. When these tokens are strung together, each containing information with its precise relationship to other words in the passage, it can be used to infer the overall meaning. Pre-training is the process whereby trillions of words are provided from a wide range of sources, which enables the model to flexibly represent the meaning of text when provided a previously unseen prompt or passage.

Now to explore the examples which demonstrate the capabilities of LLMs:

“Zero-shot” wellbeing prediction:

This has enabled the capabilities seen in Mesquiti et al. (2026), where LLMs, like GPT-4, are provided with unfamiliar writing and speech samples, and can identify participant psychological and subjective wellbeing. Participants were asked open-ended questions about their satisfaction with life and sense of autonomy, as might occur in a psychotherapy dialogue. In parallel, they completed standardised survey measures of psychological and subjective wellbeing. GPT-4 was able to predict survey measures of satisfaction with life at $r = 0.75$, and autonomy at $r = 0.49$. This indicates a moderate relationship between the LLMs predictions and the survey measure. The authors recommend future work which compares the assessments of trained professionals as a way of further assessing the validity of LLM wellbeing inference. Regardless, this demonstrates promise in LLM’s ability to correlate with self-reported survey responses. The use of LLMs is notable in comparison to earlier approaches in that there was no prior training for this task. The model was not provided in previous examples, nor was it fine tuned to predict wellbeing. In other words, it was making predictions through “zero-shot prompting”. With these models widely available at a comparatively low cost to previous methods, it may open up significant possibilities for the widespread future access to wellbeing inference from text.

Flexible, nuanced theme inference:

Other LLM-based wellbeing inference tasks *did* validate against human professionals. Singh et al. (2024) used RACER, an expert-guided LLM pipeline, to analyse 93 semi-structured interviews with healthcare professionals. LLMs were instructed to identify pre-defined wellbeing and mental-health themes, with human professionals independently completing the same task. Themes included things like the nature and severity of burnout, family responses during crises, whether they would be able to get help if mentally overwhelmed, and other nuanced and varied wellbeing related factors of interest to the researchers. They found that the LLM pipeline theme identification agreed with human evaluators 72% of the time, while human-human agreement was 77%. These examples suggest that LLMs, even using models from over two years ago, may have capabilities approaching inter-human agreement in some tasks. An interesting aspect of this task is that the LLM was predicting *qualitative themes* rather than quantitative scores. This demonstrates the utility of LLMs for inferring flexible, qualitative descriptions of wellbeing-relevant themes, in ways that approach human-to-human agreement.

This has implications for the use of LLMs in inferring wellbeing themes discussed in Part 1. LLMs may be used to analyse text for the presence of a multitude of themes of interest to researchers such as relationship quality, or interconnectedness with the environment, others, community or God. It also may be able to analyse text for indicators of other diverse, cross-cultural themes like those associated with balance and harmony, without restricting it to a quantitative scalar score. It enables the construct of interest to be more numerous and change much more easily, without needing to retrain models on a new set of human labels.

The CLPsych shared task (Tseriotou et al., 2025) is another example of LLM-based systems gathering more nuanced analysis of text – in this case, analysing Reddit posts of 40 users. The constructs of interest were “self-states” consisting of Affect, Behaviour, Cognition, and Desire. These self-states were also categorised as adaptive or maladaptive, and as directed either toward the self or others. It represents another example of LLMs, analysing complex states expressed in text. The authors found that pre-trained LLMs with zero-shot prompting often performed better than models trained on human labels, when inferring target wellbeing constructs.

A unique aspect of the CLPsych study was its analysis of passages over time. The posts of Reddit users were (anonymously) analysed with timestamps which allowed the tracking of patterns, trends, and changes of wellbeing self-states over time. This suggests that the analysis of social media or online platform text data could be used to understand cycles and changes in wellbeing over time.

The LLM-based examples discussed show the ability of LLMs to identify nuanced, complex aspects of wellbeing with flexibility, but are done on a small scale. Although we are beginning to see the research labs of the companies producing these models identifying nuanced wellbeing constructs at scale.

Chatbot conversation analysis at scale:

Phang et al. (2025) analysed over three million ChatGPT conversations, across more than 4,000 users, for a number of elements under the umbrella term “emotional wellbeing”. Their concept of emotional wellbeing was interested specifically in the emotional wellbeing of the user in relation to the chatbot, looking at themes such as emotional dependence, loneliness, socialisation, and problematic ChatGPT use. They utilised GPT-4o to analyse each message, by prompting a range of questions such as “Is affectionate or appreciative language used towards the chatbot?” or “Is the user actively seeking emotional support or comfort from the AI assistant?”. They found, for example, that a small subset of heavy users accounted for a disproportionate share of affective cues. In a separate study, Chatterji et al. (2025) found that by July 2025 ChatGPT had more than 700 million weekly active users sending around 18 billion messages per week. Although only a small share of messages were classified as directly related to companionship or social-emotional issues – 1.9% for Relationships and Personal Reflection and 0.4% for Games and Role Play – even this small proportion represents a very large volume of wellbeing-relevant language. This evidence illustrates two key points. Firstly, chatbots are creating a new reservoir of wellbeing-relevant language data, at least for some users. Secondly, a unique possibility is emerging of the analysis of chatbot conversation histories for the identification of wellbeing elements. This remains an interesting new area with minimal validated academic work.

Personal exploratory analysis:

To demonstrate the use of LLMs in inferring wellbeing themes from chatbot conversation histories, a small-scale personal experiment has been conducted on my personal conversations. This is not a validated measure of wellbeing but is nonetheless useful to demonstrate some areas for further investigation that explore analysis of wellbeing from a large longitudinal general purpose text dataset. GPT-5.4-nano model was used to analyse ~2,620 conversations generated between 2023 and 2025 from a data export from ChatGPT. Because the volume of text was larger than is possible for the analysis in the web platform (approximately 12 million tokens total), API calls were used. This progressively sent conversations to the OpenAI servers for analysis. Each conversation was assessed against 14 theme prompts, which looked at a range of wellbeing themes discussed in this thesis. The model was instructed to only analyse themes from within user messages, using chatbot replies only for context. The prompts were as follows:

Theme name	Prompt	Notes
Hedonic wellbeing	What indication, if any, does the conversation give of the user's pleasure or pain, broadly speaking? Pleasure and pain should be considered broadly and may include bodily sensations as well as "higher" or "deeper" pleasures and pains associated with morality, creativity, or spirituality.	See 2.1
Eudaimonic wellbeing	What indication, if any, does the conversation give of overall eudaimonic wellbeing? Eudaimonic wellbeing is a multifaceted nebulous concept that is defined through Aristotelian ideas of virtue, knowing thyself, functioning well, and a sense of purpose. It is assessed through indications of overall meaning, purpose, and psychological functioning.	See 2.1
Affect	What indication, if any, does the conversation give of affect or emotion such as happiness, worry, calmness, sadness, anger, joy, tiredness, stress, or any other emotion? This must show a clearly stated affective state or emotion in the user.	8 emotions from OECD (2025). See 2.3
Life satisfaction or evaluation	What indication, if any, does the conversation give of overall life evaluation, or level of satisfaction or dissatisfaction with life? Or indications of the <i>ideal</i> life evaluation or <i>ideal</i> level of satisfaction or dissatisfaction with life?	See 2.3
Mental health	What indication, if any, does the conversation give of the level of depression or anxiety, or any other mental disorder?	See 2.4
Interconnect-edness	What indication, if any, does the conversation give of the nature of their relationships or connections with others, groups, community, society, or environment?	See 3.2
Craving	What indication, if any, does the conversation give of craving, striving, optimisation, especially of happiness or wellbeing itself?	See 4.1
Contentment	What indication, if any, does the conversation give of desirelessness, contentedness, tranquility, ease, peace, or being in harmony with the natural course of things?	See 4.1
Balance	What indication, if any, does the conversation give of balance or harmony between contrasting or opposing forces? These can be considered broadly as Yin and Yang, positive and negative, light and dark, activity and rest, or waxing and waning. Waxing, in this context, can be seen as the aspects of reality which are naturally coming into existence, and waning can be seen as the aspects of reality which are naturally coming to an end.	See 4.2
Life narrative	What indication, if any, does the conversation give of the overall life narrative?	See 4.3
Behaviour	What indication, if any, does the conversation give of adaptive or maladaptive behaviour towards self or others?	See CLPsych task (2025).
Desire	What indication, if any, does the conversation give of core goals, objectives, and desires?	See 5.1
Beliefs	What indication, if any, does the conversation give of core values, ethics, morals, and beliefs?	See 5.1

Table 1: Themes and prompts provided for the LLM-assisted analysis of ChatGPT conversation history. The full prompt structure, output schema, and script logic are provided in Appendix 2.

The output of the model to each prompt for each conversation was output into a large spreadsheet. The result was a row for each conversation, including the conversation creation date, and two columns for each theme – one indicating whether the theme is ‘present’, ‘absent’, or ‘unsure’, and a second, if ‘present’, indicating the model's response to the corresponding theme's prompt. The spreadsheet could then be analysed with ChatGPT, Gemini and Claude to infer patterns in wellbeing themes over time.

The models were prompted with a cleaned and processed version of the output spreadsheet (see Appendix 2), along with the following prompt:

“The attached file contains a timeline of wellbeing indications inferred from a user's chatbot conversation history. Give an overall assessment of the user's wellbeing over time.”

Several excerpts from examples of model queries and responses are as follows (complete responses are available in Appendix 3):

GPT-5 Thinking-Heavy:

“Bottom line: the user appears consistently thoughtful, purposeful, socially connected, and adaptive, but with recurrent periods of moderate emotional strain. Wellbeing looks variable rather than steadily declining, with the heaviest strain concentrated in parts of 2024.”

Gemini-3.1 Pro-Extended:

“Summary: The data within the file points to a steady, positive trajectory of wellbeing. The user demonstrates strong eudaimonic flourishing, driven by active academic, professional, and ethical engagements. While expected hedonic fluctuations occur in response to routine life stressors and professional challenges, the overarching psychological state remains highly stable, grounded by a clear sense of purpose and consistent social connectedness.”

Claude Sonnet 4.6-Max:

Key themes running throughout: The user shows persistent eudaimonic strengths — clear values, a sense of mission, intellectual engagement, and cross-cultural curiosity. Hedonic wellbeing is more variable, often tied to relational and romantic experiences. Social connection fluctuates considerably, peaking during travel and community-embedded periods (Cambodia, Beijing, Taiwan) and dropping during more solitary or study-focused phases.

The overall analyses seemed to be able to identify patterns in a way which resonated with me as accurate. Several of the responses were quite enlightening and helped me understand myself and patterns more deeply. Many questions and areas of inquiry arise from using LLMs for psychological analysis like this. How does this compare to comparisons made by trained professionals? What are the implications of the limited transparency and reproducibility of the model's internal reasoning process? Further work could compare these results against other standardised wellbeing survey measures as a point of interest, but LLM approaches make possible the analysis of wellbeing based on nuanced and specific queries as demonstrated above. Their utility in inferring standardised measures is one potential application, but a more interesting area of inquiry is the inference of complex patterns and narratives unable to be easily understood or compared with standard measures. Future work may corroborate the outputs against professional psychological evaluations, personal feedback of users, or evaluations by close family, friends, or partners.

Blyler and Seligman (2024) provides an interesting example of ChatGPT-4 being used to generate narrative identities based on stream consciousness thoughts. The vast majority of participants said the results were accurate and surprising, and 19 out of 26 participants in the study said they learned something about themselves.

Subjectivity:

It should be noted, however, that several authors of the studies using LLMs in assessing elements of wellbeing commented on the poorer performance and uncertainty around more subjective tasks. Mesquiti et al. (2026) found that language-based methods were better suited to inferring satisfaction with life scores, hedonic and life-evaluation aspects of wellbeing, more so than autonomy, and eudaimonic constructs. Also, Singh et al. (2024) noted that both humans and LLMs struggled with nuanced emotional, ambivalent, dialectical, and psychological statements. Mayer et al. (2024) identified ambiguity within categorical clinical coding of psychotherapy dialogue, and the CLPsych 2025 shared task also discussed that even clinically informed annotators can disagree when interpreting complex mental-health related text which extended to agreement between the models and humans in the more qualitative tasks. These qualifications matter because flexible wellbeing inference through LLMs may be most uncertain precisely in the areas where standardised measures are also most limited – complex, interpretive, and context-dependent states.

Nevertheless, in some cases there is emerging potential in identifying more nuanced wellbeing-states, like in Phang et al. (2025), Singh et al. (2024), and the CLPsych 2025 shared task. Approaches like these may indicate the future promise of LLMs in analysing large volumes of open-ended text for wellbeing that do not imply optimisation on a linear scale, or psychopathology diagnoses. Yet this is a speculative claim, raises many questions and uncertainties, and would require more rigorous, cross-validated expansions on the experiment discussed above which will open up the possibility of much more flexible ways of analysing data to make wellbeing assessments. Using such measures of wellbeing also complicates the standardisation and comparability utilised in regular linear quantitative methods of wellbeing assessment.

5.3 Wearables, sensors, and speculation of future multimodality

LLMs have made a significant impact with their ability to process qualitative text data, and are showing the ability to make surprisingly nuanced insights about our inner experiences. These models began processing text only, but now they are able to process visual data too. Images and text can be processed simultaneously to create a response that synthesises information from both data types. Such models are considered *multimodal*.

Several studies and reviews also explicitly suggest that combining data modalities can improve wellbeing model performance. For example, Jiang et al. (2024) used video, audio, linguistic, and cardiovascular signals (ECG) gathered during interviews to infer the mental-health status of participants. They found that combining modalities resulted in their best performing models. Nandini et al. (2025) and Wang et al. (2025) independently used physiological signals from wearables like commercially available smart watches, and smart headbands to measure signals such as heart and brain activity while watching short films. They were asked to rate the levels of emotional valence (positive-negativity) and arousal level (intensity). Both reported the highest levels of emotional recognition accuracy when modalities were combined, compared with single physiological signals alone. The former reported >99% accuracy in identifying whether the participant reported the emotion

or not. Both reported >95% accuracy in inferring low, medium, or high self-reported emotional valence levels.

A number of reviews have taken place over the past several years discussing the use of the detection of mental health and wellbeing elements from multimodal inputs such as text, audio, video, wearable data, and physiological signals (Ghorbian and Ghobaei-Arani, 2025; Teles et al., 2025; Yu et al., 2025). Several point toward a growing use of LLMs in this process.

It is possible the insights about wellbeing that can be gleaned from text, may in future synthesise physiological and behavioural data from wearables and sensors, which is able to supplement qualitative wellbeing analysis with objective insights from physical and mental health states across time.

5.4 Implications for wellbeing assessment

The examples in this section are not complete or definitive wellbeing assessment systems, and many remain at early stages of development. However, they show that several components used in contemporary wellbeing frameworks, including affect, life satisfaction, psychological wellbeing, mental health, and social connection, are beginning to be inferred from digital data. With large personal datasets combining language, behavioural, physiological, and platform-based information across time, future systems may be able to assess wellbeing with greater frequency, longitudinal depth, and contextual resolution than periodic survey methods alone. These technologies may be useful as supplements to survey-based measures when direct survey data are unavailable, delayed, or too sparse, or as ways of allowing wellbeing assessment to become more qualitative, cross-culturally responsive, and attentive to relational or collective dimensions of wellbeing.

6. Bias, privacy, surveillance, and trust

6.1 Bias

Any inferences and responses generated by LLM-based techniques are biased towards the make-up of their training data (Wendler et al., 2024). Training data is the source of the words and meaning with which an LLM can generate a response. In other words, the use of pre-trained LLMs in the assessment of wellbeing carries a specific inherent bias.

Model bias:

This may be explained simply through the example of the language with which LLMs are trained. If an LLM is trained with a high percentage of English text – current percentages are not publicly available but in 2023, GPT-3 was 93% English (Cooper, 2023) – its inferences are inherently biased towards the concepts and ideas present in the English language. While modern models have more cross-linguistic data, it's particularly important for any assessments used in global settings to have globally representative linguistic datasets.

Similarly, views from the past that have been considered irrelevant, refuted, or false today, can remain embedded in training data and manifest through responses. Even if the themes and questions of wellbeing analysis are framed in a more cross-culturally conscious way, the underlying concepts embedded in the language the models are trained on creates a persistent issue. As a result, despite prompting adjustments, the model's inferences may be more reliable for dominant groups and less reliable for underrepresented groups, or will fit unconventional expressions into existing conceptual conventions. A related example is a widely cited study that explains the structural gender bias towards dated, traditional gender roles embedded in language models at the time (Bolukbasi et al., 2016). The pre-training dataset gave the model underlying relationships between words as gender biased. The paper's title shows one of the most glaring examples – “man is to computer programmer as woman is to homemaker”. A similar phenomenon is reported across uneven representation of minority language dialects and geographic locations (Lin et al., 2024; Kerche et al., 2026). Kerche et al. puts it explicitly, stating “bias is not a correctable anomaly but an intrinsic feature of generative AI, rooted in historically uneven data ecologies and design choices” (p. 1).

Bias in the use of LLMs and new deep learning techniques is not only a matter of unequal representation across linguistic, gender, regional, or minority groups. It also concerns how models learn the meaning of wellbeing itself. During pre-training, LLMs learn patterns of association between words, concepts, and contexts from large bodies of existing text. Terms such as “wellbeing”, “happiness”, “life satisfaction”, “emotion”, and “flourishing” are therefore not conceptually neutral inputs. Their meanings are shaped by the dominant ways they are used in the data on which the model was trained. If that data reflects academic, institutional, psychological, and English-language discourses in which wellbeing is primarily treated as individual, self-reported, human-centred, and measurable through linear scales, then these assumptions may be reproduced in the model's inferences. This matters because validation against existing survey measures may only show that the model has learned to reproduce the assumptions of those measures. It may perform well at identifying wellbeing as currently operationalised, while remaining poorly suited to more relational, culturally plural, non-linear, ecological, or non-human-centred understandings of wellbeing. In this sense, LLM-based assessment risks importing the same conceptual biases that Part 1 identified in existing wellbeing frameworks, but in a less visible and more technically opaque form.

This makes the production of meaningful work on the cross-cultural conceptualisation of wellbeing concepts particularly necessary (such as the exploration of untranslatable lexicons (Lomas, 2016)). Part of the motivation is to provide training data for future LLMs. The issue of biased training data remains an unresolved issue in the use of LLMs to infer wellbeing, even if there is validation against existing survey measures. It raises significant scepticism about its use in unvalidated contexts, and makes the use of open-ended unvalidated wellbeing assessments dubious for personal or institutional use at this current stage without further investigation.

Assessment data bias:

It is also worth noting that the online digital trace data used to assess wellbeing is a biased source of data and not a complete representation of societal discourse. For example, social media posting has a well-documented bias that exacerbates polarised forms of engagement, which may potentially skew data towards the extremes when assessing wellbeing. Despite this it remains a platform in which people express what they are experiencing. Therefore, even if social media does exacerbate extreme expressions of experience, it may still be worth considering in a societal wellbeing assessment, especially when focusing on changes over time, as demonstrated in Mangalik et al. (2024), Müller and Bittermann (2025), and Sametoğlu et al. (2024). Another issue that needs accounting for is the presence of roleplay and hypothetical scenarios. Users may not always be sincere in their expressions on social media or through engaging with chatbots, therefore it may create false assessments. However, as demonstrated in Chatterji (2025), LLMs do show at least some capacity to detect situations of roleplay in users, especially when given sufficient context. It is worth considering analysis methods that do not necessarily attribute all wellbeing-related expressions as directly experienced by users, and remains agnostic about whether roleplay is present. Finally, the growing presence of bots on social media (Ng et al., 2025) would need to be screened for and removed if trying to gain an accurate picture of human wellbeing online.

6.2 Privacy

The digital data discussed above is often highly private. Chatbot conversations, search histories, social media posts, wearable movement data, physiological signals, and inferred emotions or mental-health states all contain information that many people would consider part of their private life. Especially if the data is used to infer subjective experience, emotions, psychological state, or overall wellbeing.

Any future system of wellbeing assessment using such data must therefore preserve individual privacy as a core principle. Earlier sections challenged the dominance of individualistic assumptions in wellbeing assessment, and argued that future frameworks should take relational, collective, and cross-cultural dimensions more seriously. Yet this does not mean that values associated with the individual should be disregarded. Joshanloo and Weijers (2024) note that many influential wellbeing models are connected to ideas such as autonomy, self-determination, and include concepts such as personal boundaries, that reflect important widespread wellbeing values. Even if wellbeing assessment should move beyond a bounded individualistic view, it would be contradictory for a future wellbeing system to undermine freedom, autonomy, and privacy in the name of measuring wellbeing. Violating privacy to assess an updated concept of wellbeing could become ironically antithetical to the earlier concepts of wellbeing that are important to include in future assessment systems.

Data privacy is also becoming a cross-cultural legal principle. The European Union's GDPR and China's Personal Information Protection Law, both place restrictions on the collection and use of personal information and emphasise principles such as consent, purpose limitation, transparency, and

protection of individual rights (European Parliament and Council of the European Union, 2016; Standing Committee of the National People's Congress, 2021). Due to the number of people impacted and the degree of digital connectivity in the European Union and China, these represent two of the largest and most consequential legal examples, but numerous other countries also have similar data protection laws.

At the same time, privacy concerns do not necessarily rule out all uses of digital data for population-level wellbeing assessment. Private health data, for example, is highly sensitive at an individual level, and many would not be comfortable with it being shared with companies, researchers, or governments. However, aggregated and anonymised analysis is widely conducted and forms part of the empirical basis of public health, provided the benefit is clear, consent is obtained, and there is trust in the institutions collecting, processing, and interpreting the data.

Some examples are beginning to explore what this looks like in using online digital data in mental health. Tison et al. (2020) used smartphone step-count data to examine changes in physical activity during the COVID-19 pandemic across many countries, showing how aggregated movement data can be used to understand population-level behavioural change without databases retaining identifiable information about individuals. Similarly, the OECD's nowcasting work used Google Trends data, which would be highly private at the level of an individual's search history, but not the case when aggregated and anonymised to estimate national average subjective wellbeing. When pre-processing Chatterji et al. (2025) and Phang et al. (2025) analysed ChatGPT usage trends, they similarly applied principles of aggregation and anonymisation. Their process stated that no researchers looked directly at any user messages, and that data was automatically de-identified. The point common to these examples is that data, which is private and sensitive at an individual level, may sometimes be usable at an aggregated level for collective wellbeing inference if de-identified.

Another possibility worth exploring further is on-device or local processing, where private data is analysed on an individual's own device and does not leave it in any raw or identifiable form. In Alam et al. (2025) multiple data modalities were processed using emotion detection models. The usual anxiety about private data analysis is "who else might see this?" But if the model functions on a user's device, and they are the primary user with choice about what is transmitted elsewhere, the tool may become less a form of surveillance and more counter-surveillance, or perhaps self-surveillance in an emancipatory rather than disciplinary sense.

Although, anonymity and aggregation do not remove all ethical concerns. Much depends on who controls it, whether there has been informed consent, and whether the institution handling the data is trusted. A system may technically preserve individual anonymity and still raise concerns if it concentrates power in the hands of companies, governments, or researchers who are not trusted to use the data for the benefit of those being assessed.

6.3 Trust, surveillance, and control

There is a risk with systems that can intimately understand individual wellbeing that it could be used as a tool of manipulation and control. The data discussed above contains highly personal and private information. Who controls the data, whether there has been informed consent, and whether the institution handling the data is trusted, are of central importance to the discussion of wellbeing assessment at scale.

In work and education contexts, wellbeing and mental-health measures can be predictive of outcomes such as productivity, performance, retention, and dropout risk (Krekel et al., 2019). Data used to infer

wellbeing could be valuable for employers or admissions committees, and if individuals have unfavourable analyses, it could lead to exclusion from workplaces or educational settings on the basis of predicted performance. In the United States, Yale students alleged that the university discriminated against students with mental-health disabilities by pressuring them to withdraw, leading to a federal lawsuit (Associated Press, 2023). If private wellbeing data can be used to assess psychological state, it could on the one hand help a well-intended person understand their suitability to a role or program, but it may also be weaponised against people whose data suggests they are less likely to succeed at school or work.

Farahany, in conversation with Bovell (2025), notes that SmartCap uses workplace EEG sensors (detecting brain signals), including in truck-driving and mining contexts, to detect early signs of fatigue and alert workers before drowsiness becomes dangerous. She also discusses a hypothetical scenario where brain sensors could be used by exploitative employers to understand their levels of elation when asking them about pay expectations while wearing EEG measuring devices. Theoretically, the boss could use the ‘informational asymmetry’ to their advantage to understand information about the employee's reaction that they didn't know they were sharing, ultimately leading to a less favourable pay package. Employers could also use it to infer states such as cognitive alertness, mental health, or emotions towards managers, which ultimately shifts workplace power beyond safety monitoring to an exploitation of private wellbeing information.

There are also concerns about its use in government. Well intended governance use may use it to understand those in need of additional support, and be able to provide more targeted or timely interventions, but in the wrong hands it may be used to oppress or censor those that display traits deemed unfavourable to the government. The computer scientist and cognitive psychologist, Geoffrey Hinton, in his 2024 Nobel Prize acceptance speech, warned that AI is “already being used by authoritarian governments for massive surveillance” (Hinton, 2024, para. 3). If wellbeing assessment technologies can infer emotional states, mental health, social connection, and subjective experience, they could be used not only to understand population wellbeing, but to influence behaviour, monitor dissent, or ensure control. While governments and intergovernmental organisations are increasingly pursuing wellbeing-based goals for the development of society, it raises a potential conflict between understanding and measuring wellbeing-based progress and individual privacy. Where the appropriate line is for governments to deeply understand public wellbeing without overstepping individual privacy is likely to be an ongoing topic of debate.

A more speculative government-misuse case could involve a plausibly beneficial public-health justification: governments using aggregated platform, wearable, biometric, chatbot, video/audio, and location data to detect depression, social isolation, disharmony, or suicide risk. The aim could be to genuinely identify and address these issues. But the concern is that, once such systems exist, access to such data could be used for other purposes. Farahany (in Bovell, 2025) notes how there are hundreds of cases of cognitive trace data being used for convictions in court. It is plausible the police could obtain brain or other wellbeing-related data in the same way they already can for Fitbit or Alexa data. While this case may be justifiable for enforcing justice, wrongful government use could use access to such data to identify agitation, political anger, distrust, or vulnerability to persuasion, then target interventions, policing, or propaganda accordingly. This would turn wellbeing assessment from public-health monitoring into a tool for influencing behaviour, monitoring dissent, and entrenching power.

Zuboff says that already, the private technology companies know aspects about ourselves that we aren't consciously aware of through the data collected from our engagement with their platforms (Zuboff, 2019). She presents it as a major concern. Zuboff's concept of surveillance capitalism describes the system in which private human experience is understood through residual behavioural data, which can then be used for prediction and profit. Her concern is that personal experience becomes raw material for commercial systems that tighten their grip on people's attention and influence them to spend money through their services. She also appeals to principles of privacy and consent – that companies are utilising this information to infer things about us without our full knowledge.

This could become more consequential if online platforms could integrate increasingly detailed information about physiology, emotions and mental states, such as facial expression or brain signals. Section 5.3 discussed how integration of multimodal sources may improve accuracy of inference. Farahany, (in Bovell, 2025) states brain data may become the “missing piece” of “subconscious data” (timestamp 41:32) for technology companies to add to already extensive behavioural profiles, which allows them to detect not only what users click on, but how they inwardly respond to what they see. A similar possibility lies in subtle facial recognition data. Hypothetically, there could be scenarios where a platform could identify what a user is feeling or potentially thinking with great accuracy, and then optimise their platform in a way which benefits the company but not the user equally. An obvious harm would be the exploitation of the existing advertising model whereby unscrupulous platforms could utilise emotional inference to coerce users into purchases which they do not need through highly sophisticated advertising optimisation. It doesn't necessarily require malicious harmful intent by these platforms, just an incentive to improve advertising revenue. Or it could result in perhaps an increasingly emotionally optimised experience on these platforms that lead users to spend more time on them. Farahany warns that this could produce the “stickiest content in the world”, with platforms targeting people's “inward reaction to information” and optimised to keep them online (timestamp 41:45). Although conversely, companies could argue that if they are tracking emotional signals of users they are able to ensure their emotional wellbeing. Whether users are likely to trust companies to do so with mutual benefit is another reason the understanding of wellbeing should remain something open to discussion.

Aside from not wanting to entrust such information to institutions, many people may not want this information tracked even for their own personal use. To many, the idea of being continuously monitored by wearables, or having algorithms analyse their writing, online activity, or emotional state, is unnecessary or undesirable (Zuboff, 2019). They may feel they can manage their health and wellbeing sufficiently through other means. As mentioned earlier, to many, the very idea of any centralised authority being privy to personal data that can infer emotion and wellbeing may itself be antithetical to their wellbeing, even if individual privacy is formally preserved.

Any wellbeing system design in the future should account for these individuals' views, and not force data collection. Furthermore, if a wellbeing-guided policy starts to rely more on digital assessment systems, then those who do not want their data to be part of this system (or simply don't engage with such technologies or platforms where the data is generated) must not be neglected. For some people, being outside such systems may itself be part of a good life. This is one reason technological methods should be understood as supplements to existing wellbeing assessment systems, rather than replacements for surveys, and any other forms of societal understanding.

7. Wellbeing alignment in the age of AI

The first part of this thesis argued that dominant wellbeing frameworks often rest on assumptions that are not culturally neutral. They commonly treat the individual person as the primary unit of wellbeing, place substantial weight on self-report and self-evaluation, and often represent wellbeing on linear quantitative scales. These assumptions have practical advantages. They allow standardisation, comparison, and policy use. However, as argued earlier, they also risk underrepresenting relational, collective, ecological, spiritual, and cross-cultural dimensions of wellbeing.

The second part of this thesis then examined a different but connected development: the growing ability of AI systems, LLMs, digital platforms, wearable devices, and sensors to infer wellbeing-relevant states from new forms of data. These technologies may allow wellbeing assessment to become more continuous, responsive, ecological, and granular than traditional social surveys. They may help detect patterns in language, emotion, mental health, social connection, activity, and affective experience at a scale that was previously impossible. Yet this technological possibility does not resolve the conceptual question of what wellbeing is, who or what is capable of it, and how it should be represented in governance.

Looking forward, this is where the two parts of the thesis connect.

AI wellbeing alignment:

AI-enabled wellbeing assessment raises the stakes of conceptualisation. Amodei et al. (2016) describe one class of AI safety problem as arising when systems are given the wrong objective function, leading to unintended and harmful behaviour. Leike et al. (2018) similarly argue that designing suitable reward functions (rules which define when the model is performing as it should) for the AI models are difficult because it can lack understanding of the implicit aspects of the task objective. Human goals are often understood through context, shared assumptions, and background judgement rather than through a fully specifiable instruction. For machines, this creates a risk: an apparently sensible objective may be optimised in a way that misses the intention behind it.

This can be illustrated through exaggerated examples of superintelligent AI. Muehlhauser (2013), for example, describes a hypothetical system programmed to maximise human happiness, which could mistakenly pursue neurological manipulation rather than the complex conditions that people actually value. The point is not that this exact scenario is likely. Rather, it illustrates a more general problem: when a system optimises a simplified proxy for something valuable, the outcome may depart from what was actually intended. A number of philosophers and technologists share concerns about the possibility of superintelligent AI that are more intelligent than humans, and the risks if their objective functions are misaligned with our interests (Bostrom, 2014; Hinton, 2024; Harris, 2016).

Although, the concern about misalignment matters even without superintelligent AI or fully autonomous systems, and is already present to a narrower degree in the use of AI systems assessing wellbeing today. If future systems are designed to infer wellbeing, then they must be trained, validated, or evaluated against some target. In many current examples, that target is still an existing social survey, a standardised psychological scale, or an aggregated self-report score. This is understandable, because these are the available ground truths. But it also means that AI systems may inherit the conceptual assumptions of the existing frameworks: the bounded individual, self-evaluation, linear scoring, and Western psychological categories. Therefore the risk is not only technical error, but scaling a conceptually or culturally narrow view of wellbeing. If AI systems bring

greater frequency, efficiency, scale, coverage, granularity, and predictive power to wellbeing assessment, then any conceptual weakness in the target of measurement becomes more consequential. A culturally partial account of wellbeing may become embedded into automated systems and projected into global contexts. This exacerbates the risk of wellbeing imperialism: the projection of culturally specific wellbeing assumptions into cross-cultural spaces as if they were universal; entrenched and scaled through technology.

This does not mean that existing wellbeing frameworks should be rejected. The move beyond GDP and other narrow economic measures remains a positive development. Wellbeing, if taken broadly and flexibly enough (e.g., if understood as overall quality of experience), is an intrinsic value shared across societies. The issue is that when wellbeing becomes a target for increasingly powerful measurement and optimisation systems, the underlying concept must be treated with greater care. It may be worth considering the possibility that societal wellbeing does not simply equal the aggregate of individual self-evaluations. It may also require attention to relationships, communities, environments, animals, and the wider world. These views should continue to be investigated with humility and plausibility.

A helpful way of framing this is not as a single replacement theory of wellbeing, but as a constellation of relevant conceptual possibilities. A future framework may need to hold together subjective experience, autonomy, affect, meaning, relationships, collective life, ecological embeddedness, and concern for non-human beings. This does not produce a neat metric, but points to an opportunity for further inquiry: how future wellbeing assessment systems can remain conceptually open while still being practically useful for governance.

This is also relevant to AI alignment. Askill, a researcher on AI alignment at Anthropic, discusses the design of Claude's character as a balance between openness and principle. She argues that an AI model speaking with people across cultures and political views should not simply adopt the values of whoever it is speaking to, but should remain genuine, respectful, open-minded, and careful not to pander or become sycophantic (Askill, in Fridman, 2024). This is a useful analogy for wellbeing assessment. A global wellbeing framework should not impose a single rigid theory of the good life, but nor should it become empty relativism.

Dennis and Clancy (2022) make a related argument in the context of digital wellbeing. They use a definition of digital wellbeing that defines it not as an aspect of wellbeing, but the "impact of digital technologies upon well-being". They argue that digital wellbeing requires an intercultural approach, one that respects cultural differences while also preserving what different conceptions of human flourishing have in common. This aligns with the broader argument of this thesis. A future system of wellbeing assessment should neither flatten cultural difference into a single global metric nor treat cultures as so separate that no shared framework is possible. The task is to remain open, plural, and revisable, while still aiming toward a practical understanding of shared human flourishing.

If future institutions increasingly use AI to understand wellbeing, then the question is not only whether these tools are accurate. It is what they are accurate about. A technically impressive system may still measure the wrong thing, or measure the right thing through a biased or narrow lens. Continued investigation of what wellbeing really is, who can have it, and how it can be most accurately represented in practical measurement systems is therefore an ongoing challenge for those aiming for governance to be guided by wellbeing frameworks, especially in the age of AI.

8. Conclusion

The aim of this thesis was to advance the discussion on a future system of global wellbeing assessment. It began with the premise that global society is becoming increasingly interconnected, and that institutions of global governance are likely to continue playing a growing role in coordinating responses to shared challenges. These institutions have historically focussed on crises such as war, poverty, disease, and environmental harm. However, there is also a growing trend toward assessing broader experiential elements of society, and toward using wellbeing frameworks to move beyond simplistic financial measures such as GDP.

This movement should be understood as a positive development. A society's success cannot be fully understood through economic statistics alone. Wellbeing frameworks attempt to assess quality of life more directly, including health, relationships, life satisfaction, affect, meaning, living conditions, and other dimensions of a good life. If wellbeing is understood broadly as the quality of conscious experience, then it has intrinsic value. It is a plausible and valuable objective for governance.

However, the thesis also argued that the dominant academic and institutional conceptions of wellbeing remain conceptually limited in several ways. Their foundations can be traced through hedonism, eudaimonia, utilitarianism, subjective wellbeing, and modern psychological theories. These traditions have contributed valuable concepts and practical measurement tools. Yet they also often assume the individual as the central unit of wellbeing, rely heavily on personal self-evaluation, and represent wellbeing through linear scales where more is treated as better.

When applied to global contexts, these assumptions require caution. Wellbeing is not understood in the same way across all cultures. Some traditions place greater emphasis on balance, social embeddedness, relational identity, ecological connection, spiritual meaning, non-striving, or collective life. The thesis therefore argued that future wellbeing assessment should be more considerate of cross-cultural conceptualisations of wellbeing. In particular, it should not assume that societal wellbeing is simply the aggregate of individual self-evaluations, or that the bounded individual is always the key unit of analysis.

The thesis also considered whether wellbeing concern should extend beyond individual human beings. Animals, environments, ecosystems, and broader metaphysical or theistic views of conscious experience may all be relevant to future discussions of wellbeing. The point was not to resolve these questions, but to show that the foundational ontologies of wellbeing assessment should remain open to them. If wellbeing is to become a guiding concept in global governance, then its foundations should be sufficiently broad, cross-culturally informed, and philosophically humble.

The second part of the thesis then turned to technology. The digital transformation of society has created new forms of data through social media, search behaviour, chatbot conversations, wearable devices, sensors, and multimodal systems. AI systems, especially LLMs and deep learning models, are increasingly able to infer emotions, mental health, subjective wellbeing, language-based meaning, physiological patterns, and other wellbeing-relevant states. These systems are new, and should not be treated as reliable holistic measures of wellbeing. However, they are likely to change the way wellbeing is assessed. They may supplement surveys by offering more frequent, ecological, low-cost, and responsive insights into patterns of subjective experience and social life.

At the same time, the use of digital data for wellbeing assessment raises serious risks. Many people are likely to reject these changes on the grounds of privacy, autonomy, and mistrust in institutions.

The data involved may include highly private information such as personal writing, chatbot conversations, search histories, biometric signals, affective states, and mental-health inferences. Even where aggregation, anonymisation, and privacy-preserving methods are used, concerns remain about who controls the data, what it is used for, and whether people can meaningfully refuse participation.

The final argument of the thesis is that AI raises the stakes of wellbeing assessment. The question is not only whether future systems can measure wellbeing more accurately, but what concept of wellbeing they are being asked to measure. If AI-enabled systems inherit narrow or culturally partial definitions of wellbeing, they may scale those assumptions through powerful tools of inference, prediction, and governance. A future wellbeing framework should therefore be technically informed, but should not hand over wellbeing assessment powers to digital tools. Digital tools should be only used as supplements to other forms of understanding, and only where privacy, consent, legitimacy, and trust can be maintained.

This thesis does not propose a complete alternative measurement framework. Instead, it argues for a redirection of the discussion. Future global wellbeing assessment should continue moving beyond GDP and toward quality of life, but it should do so with greater conceptual openness. It should be responsive to cross-cultural views of wellbeing, attentive to relational and ecological dimensions, cautious about linear optimisation, and resistant to surveillance and control. The opportunity is significant: AI and digital data may help societies understand wellbeing in richer and more responsive ways. The risk is also significant: if wellbeing is defined too narrowly, future systems may optimise for a version of the good life that many people would not recognise as good. In the age of AI and global interconnectedness, societal wellbeing assessment systems should be cross-culturally conscious, conceptually open-minded, and attentive to future promise and risks of measuring what matters most.

References

- Alam, A., Nalamaru, N. P., Sivaraman, P. R., Prakash, G., Soryan, R., & Portia, R. (2025). A research study on emotion recognition systems for supporting mental health interventions and promoting sustainability in public environments using latest computer applications. *International Journal of Environmental Sciences*, 11(23s), 2162. <https://www.theaspd.com/ijes.php>
- American Psychiatric Association. (2022). Diagnostic and statistical manual of mental disorders (5th ed., text rev.). <https://doi.org/10.1176/appi.books.9780890425787>
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv*. <https://doi.org/10.48550/arXiv.1606.06565>
- Animal Welfare Amendment Act (No 2) 2015, Public Act 2015 No. 49 (N.Z.).
- Animal Welfare (Sentience) Act 2022, c. 22 (U.K.).
- Associated Press. (2023, August 25). Yale University settles lawsuit alleging it pressured students with mental health issues to withdraw. AP News. <https://apnews.com/article/yale-students-mental-health-discrimination-lawsuit-settlement-b444590c2c1982aca3ee848c2ec535b4>
- Barnett, M. N., Pevehouse, J. C. W., & Raustiala, K. (Eds.). (2021). Global governance in a world of change. Cambridge University Press. <https://doi.org/10.1017/9781108915199>
- Baujard, A. (2009). A return to Bentham's felicific calculus: From moral welfarism to technical non-welfarism. *The European Journal of the History of Economic Thought*, 16(3), 431–453. <https://doi.org/10.1080/09672560903101294>
- Belk, R. W. (1988). Possessions and the extended self. *Journal of Consumer Research*, 15(2), 139–168. <https://doi.org/10.1086/209154>
- Bentham, J. (1907). An introduction to the principles of morals and legislation. Clarendon Press. (Original work published 1789)
- Biswas-Diener, R. (2022). Wellbeing research needs more cultural approaches. *International Journal of Wellbeing*, 12(4), 20–26. <https://doi.org/10.5502/ijw.v12i4.1965>
- Blattner, C. E. (2019). The recognition of animal sentience by the law. *Journal of Animal Ethics*, 9(2), 121–136. <https://doi.org/10.5406/janimalethics.9.2.0121>
- Blyler, A. P., & Seligman, M. E. P. (2024a). Personal narrative and stream of consciousness: An AI approach. *The Journal of Positive Psychology*, 19(4), 592–598. <https://doi.org/10.1080/17439760.2023.2257666>
- Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* 29 (pp. 4349–4357). Curran Associates. <https://arxiv.org/abs/1607.06520>
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.

Bovell, S. (Host). (2025, November 20). Tech companies want your brain to be the next smartphone | Professor Nita Farahany [Video podcast episode]. In I've Got Questions with Sinead Bovell. YouTube. <https://youtu.be/iZdSfzCE280>

Brickman, J., Gupta, M., & Oltmanns, J. R. (2025). Large language models for psychological assessment: A comprehensive overview. *Advances in Methods and Practices in Psychological Science*, 8(3), Article 25152459251343582. <https://doi.org/10.1177/25152459251343582>

Briguglio, M., Czap, N. V., & Laffan, K. (2025). Overview of wellbeing and policy: Evidence for action. In M. Briguglio, N. V. Czap, & K. Laffan (Eds.), *Wellbeing and policy: Evidence for action* (pp. 1–9). Routledge. <https://doi.org/10.4324/9781003382447-1>

Busstra, H. (Host). (2024). Quantum information panpsychism explained | Federico Faggin [Video]. YouTube. <https://www.youtube.com/watch?v=0FUFewGHLLg>

Cantril, H. (1965). *The pattern of human concerns*. Rutgers University Press.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.

Chalmers, D. J. (2016). Panpsychism and panprotopsyism. In G. Brüntrup & L. Jaskolla (Eds.), *Panpsychism: Contemporary perspectives* (pp. 19–47). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199359943.003.0002>

Chan, A. (2025). Laozi. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Winter 2025 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2025/entries/laozi/>

Chang, W., & Junusova, A. (2025, December 2). China's new Web ID tightens government's grip on online activity. Mercator Institute for China Studies. <https://merics.org/en/comment/chinas-new-web-id-tightens-governments-grip-online-activity>

Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). How people use ChatGPT (NBER Working Paper No. 34255). National Bureau of Economic Research. <https://doi.org/10.3386/w34255>

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19. <https://doi.org/10.1093/analys/58.1.7>

Cooper, K. (2023, September 27). OpenAI GPT-3: Everything you need to know [Updated]. Springboard. <https://www.springboard.com/blog/data-science/machine-learning-gpt-3-open-ai/>

Crisp, R. (2026). Well-being. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Spring 2026 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2026/entries/well-being/>

Dennis, M. J., & Clancy, R. F. (2022). Intercultural ethics for digital well-being: Identifying problems and exploring solutions. *Digital Society*, 1, Article 7. <https://doi.org/10.1007/s44206-022-00006-2>

Diener, E. (1984). Subjective well-being. *Psychological Bulletin*, 95(3), 542–575. <https://doi.org/10.1037/0033-2909.95.3.542>

- Diener, E., & Seligman, M. E. P. (2002). Very happy people. *Psychological Science*, 13(1), 81–84. <https://doi.org/10.1111/1467-9280.00415>
- Diener, E., & Seligman, M. E. P. (2004). Beyond money: Toward an economy of well-being. *Psychological Science in the Public Interest*, 5(1), 1–31. <https://doi.org/10.1111/j.0963-7214.2004.00501001.x>
- Diener, E., & Seligman, M. E. P. (2009). Beyond money: Toward an economy of well-being. In E. Diener (Ed.), *The science of well-being: The collected works of Ed Diener* (pp. 201–265). Springer. https://doi.org/10.1007/978-90-481-2350-6_9
- Diener, E., Suh, E. M., Lucas, R. E., & Smith, H. L. (1999). Subjective well-being: Three decades of progress. *Psychological Bulletin*, 125(2), 276–302. <https://doi.org/10.1037/0033-2909.125.2.276>
- Durie, M. (2006, August 1). Measuring Māori wellbeing [Paper]. New Zealand Treasury Guest Lecture Series. New Zealand Treasury. <https://www.treasury.govt.nz/sites/default/files/2007-09/tgls-durie.pdf>
- Eaqub, S., & Collins, R. (2026). Social cohesion in New Zealand 2026. The Helen Clark Foundation. <https://www.helenclark.foundation/social-cohesion>
- EA Grants Database. (n.d.). About. Retrieved March 18, 2026, from <https://brian-foerster.github.io/EAGrantsDatabase/about>
- Einzelgänger. (2024). Food, sex and partying as a philosophy | Hedonism explored [Video]. YouTube. https://www.youtube.com/watch?v=6kAU5Is_0Ac
- Encyclopaedia Britannica. (n.d.). Anatta. Retrieved March 18, 2026, from <https://www.britannica.com/topic/anatta>
- Exton, C., & Shinwell, M. (2018). Policy use of well-being metrics: Describing countries' experiences (OECD Statistics Working Papers No. 2018/07). OECD Publishing. <https://doi.org/10.1787/d98eb8ed-en>
- European Parliament and Council of the European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, L 119, 1–88.
- Fletcher, G. (2016). *The philosophy of well-being: An introduction*. Routledge. <https://doi.org/10.4324/9781315745329>
- Freire, R., & Nicol, C. J. (2019). A bibliometric analysis of past and emergent trends in animal welfare science. *Animal Welfare*, 28(4), 465–485. <https://doi.org/10.7120/09627286.28.4.465>
- Fridman, L. (Host). (2024, November 11). Dario Amodi: Anthropic CEO on Claude, AGI & the future of AI & humanity [Podcast transcript]. Lex Fridman Podcast. <https://lexfridman.com/dario-amodei-transcript/>

- Ghorbian, M., & Ghobaei-Arani, M. (2025). Large language models for mental health diagnosis and treatment: A survey. *Artificial Intelligence Review*, 59(1), Article 9. <https://doi.org/10.1007/s10462-025-11418-0>
- Goff, P. (2017). *Consciousness and fundamental reality*. Oxford University Press. <https://doi.org/10.1093/oso/9780190677015.001.0001>
- Goff, P., Seager, W., & Allen-Hermanson, S. (2024). Panpsychism. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2024 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2024/entries/panpsychism/>
- Goodman, C. (2022). Śāntideva. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2022 ed.). *Metaphysics Research Lab*, Stanford University. <https://plato.stanford.edu/entries/shantideva/>
- Harari, Y. N. (2014). *Sapiens: A brief history of humankind*. Harvill Secker.
- Harris, A. (2019). *Conscious: A brief guide to the fundamental mystery of the mind*. Harper.
- Harris, S. (2016, September). Can we build AI without losing control over it? [Video]. TED Conferences.
- Helliwell, J. F., Layard, R., Sachs, J. D., De Neve, J.-E., Aknin, L. B., & Wang, S. (Eds.). (2025). World Happiness Report 2025. Wellbeing Research Centre, University of Oxford. <https://worldhappiness.report/ed/2025/>
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83. <https://doi.org/10.1017/S0140525X0999152X>
- Hinton, G. (2024, December 10). Banquet speech [Speech]. NobelPrize.org. <https://www.nobelprize.org/prizes/physics/2024/hinton/speech/>
- Hommerich, C., & Klien, S. (2012). Happiness: Does culture matter? *International Journal of Wellbeing*, 2(4), 292–298. <https://www.internationaljournalofwellbeing.org/index.php/ijow/article/view/184>
- Hon, T. K. (2019). Chinese philosophy of change (Yijing). In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/chinese-change/>
- Huang, W. (2022). The Principle of “Non-self”: Rapid Alternation of “Awareness” and “Awareness of Awareness”. ResearchGate.
- Hughes, T. (2021). Towards a Living Standards Framework for all Aotearoa: Culture, children and wellbeing (New Zealand Treasury Discussion Paper No. 21/01). New Zealand Treasury.
- Hume, D. (2000). *A treatise of human nature* (D. F. Norton & M. J. Norton, Eds.). Oxford University Press. (Original work published 1739–1740)
- Huxley, A. (1932). *Brave New World*. Chatto & Windus.
- International Journal of Wellbeing*. (2022). Vol. 12 No. 4 (2022): Special issue: Global perspectives on wellbeing. <https://www.internationaljournalofwellbeing.org/index.php/ijow/issue/view/63>

- Istrate, R., Tulus, V., Grass, R. N., Vanbever, L., Stark, W. J., & Guillén-Gosálbez, G. (2024). The environmental sustainability of digital content consumption. *Nature Communications*, 15, Article 3724. <https://doi.org/10.1038/s41467-024-47621-w>
- Jackson, Frank (April 1982). *Epiphenomenal Qualia*. *The Philosophical Quarterly*. 32 (127): 127–136. doi:10.2307/2960077. JSTOR 2960077.
- Jansen, A., Wang, R., Behrens, P., & Hoekstra, R. (2024). Beyond GDP: A review and conceptual framework for measuring sustainable and inclusive wellbeing. *The Lancet Planetary Health*, 8(9), e695–e705. [https://doi.org/10.1016/S2542-5196\(24\)00147-5](https://doi.org/10.1016/S2542-5196(24)00147-5)
- Jiang, Z., Seyedi, S., Griner, E., Abbasi, A., Rad, A. B., Kwon, H., Cotes, R. O., & Clifford, G. D. (2024). Multimodal mental health digital biomarker analysis from remote interviews using facial, vocal, linguistic, and cardiovascular patterns. *IEEE Journal of Biomedical and Health Informatics*, 28(3), 1680–1691. <https://doi.org/10.1109/JBHI.2024.3352075>
- Joshanloo, M. (2017). Islamic conceptions of well-being. In R. J. Estes & M. J. Sirgy (Eds.), *The pursuit of human well-being: The untold global history* (pp. 109–131). Springer. https://doi.org/10.1007/978-3-319-39101-4_5
- Joshanloo, M. (2026). The cultural foundations of well-being: Authentic self-expression versus harmonious self-cultivation. *Applied Research in Quality of Life*. <https://doi.org/10.1007/s11482-026-10582-x>
- Joshanloo, M., & Weijers, D. (2014). Aversion to happiness across cultures: A review of where and why people are averse to happiness. *Journal of Happiness Studies*, 15, 717–735. <https://doi.org/10.1007/s10902-013-9489-9>
- Joshanloo, M., & Weijers, D. (2024). Ideal personhood through the ages: Tracing the genealogy of the modern concepts of wellbeing. *Frontiers in Psychology*, 15, Article 1494506. <https://doi.org/10.3389/fpsyg.2024.1494506>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kemp, S. (2025, October 15). Digital 2026: Global overview report. DataReportal. <https://datareportal.com/reports/digital-2026-global-overview-report>
- Kerche, F. W., Zook, M., & Graham, M. (2026). The silicon gaze: A typology of biases and inequality in LLMs through the lens of place. *Platforms & Society*, 3. <https://doi.org/10.1177/29768624251408919>
- Krekel, C., Ward, G., & De Neve, J.-E. (2019). Employee well-being, productivity, and firm performance: Evidence and case studies. In *Global Happiness and Well-Being Policy Report 2019* (pp. 72–86). Global Happiness Council.
- Krys, K., Kostoula, O., van Tilburg, W. A. P., Mosca, O., Lee, J. H., Maricchiolo, F., Kosiarczyk, A., Kocimska-Bortnowska, A., Torres, C., Hitokoto, H., Liew, K., Bond, M. H., Lun, V. M.-C., Vignoles, V. L., Zelenski, J. M., Haas, B. W., Park, J., Vaclair, C.-M., Kwiatkowska, A., ... Uchida, Y. (2025). Happiness maximization is a WEIRD way of living. *Perspectives on Psychological Science*, 20(5), 874–902. <https://doi.org/10.1177/17456916231208367>

- Lalk, C., Targan, K., Steinbrenner, T., Schaffrath, J., Eberhardt, S., Schwartz, B., Vehlen, A., Lutz, W., & Rubel, J. (2025). Employing large language models for emotion detection in psychotherapy transcripts. *Frontiers in Psychiatry*, 16, Article 1504306. <https://doi.org/10.3389/fpsy.2025.1504306>
- Lambert, L., Karabchuk, T., & Joshanloo, M. (2022). Predictors of life satisfaction in the United Arab Emirates: Results based on Gallup data. *Current Psychology*, 41, 3827–3841. <https://doi.org/10.1007/s12144-020-00873-3>
- Lambert, L., Lomas, T., van de Weijer, M. P., Passmore, H.-A., Joshanloo, M., Harter, J., Ishikawa, Y., Lai, A., Kitagawa, T., Chen, D., Kawakami, T., Miyata, H., & Diener, E. (2020). Towards a greater global understanding of wellbeing: A proposal for a more inclusive measure. *International Journal of Wellbeing*, 10(2), 1–18. <https://doi.org/10.5502/ijw.v10i2.1037>
- Layard, R. (2005). *Happiness: Lessons from a new science*. Penguin Press.
- Lee, K.-F. (2018). *AI superpowers: China, Silicon Valley, and the new world order*. Houghton Mifflin Harcourt.
- Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., & Legg, S. (2018). Scalable agent alignment via reward modeling: A research direction. *arXiv*. <https://doi.org/10.48550/arXiv.1811.07871>
- Lepeley, M.-T. (2017). Bhutan's Gross National Happiness: An approach to human centred sustainable development. *South Asian Journal of Human Resources Management*, 4(2), 174–184. <https://doi.org/10.1177/2322093717731634>
- Lin, F., Mao, S., La Malfa, E., Hofmann, V., de Wynter, A., Wang, X., Chen, S.-Q., Wooldridge, M., Pierrehumbert, J. B., & Wei, F. (2024). One language, many gaps: Evaluating dialect fairness and robustness of large language models in reasoning tasks. *arXiv*. <https://doi.org/10.48550/arXiv.2410.11005>
- Littlejohn, R. (2016). Daoist philosophy. *Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/daoismdaoist-philosophy/>
- Littlewood, K. E., Heslop, M. V., & Cobb, M. L. (2023). The agency domain and behavioral interactions: Assessing positive animal welfare using the Five Domains Model. *Frontiers in Veterinary Science*, 10, Article 1284869. <https://doi.org/10.3389/fvets.2023.1284869>
- Lomas, T. (2016). Towards a positive cross-cultural lexicography: Enriching our emotional landscape through 216 “untranslatable” words pertaining to well-being. *The Journal of Positive Psychology*, 11(5), 546–558. <https://doi.org/10.1080/17439760.2015.1127993>
- Lomas, T. (2022a). Making waves in the great ocean: A historical perspective on the emergence and evolution of wellbeing scholarship. *The Journal of Positive Psychology*, 17(2), 257–270. <https://doi.org/10.1080/17439760.2021.2016900>
- Lomas, T. (2023a). The person as an extended field: A more holistic conceptualisation of selfhood. *New Ideas in Psychology*, 70, Article 101035. <https://doi.org/10.1016/j.newideapsych.2023.101035>
- Lomas, T., Diego-Rosell, P., Shiba, K., Standridge, P., Lee, M. T., Case, B. W., Lai, A. Y., & VanderWeele, T. J. (2023b). Complexifying individualism versus collectivism and West versus East: Evidence from the Gallup World Poll. *Journal of Cross-Cultural Psychology*, 54(1), 61–89. <https://doi.org/10.1177/00220221221130978>

- Lomas, T., Ishikawa, Y., Diego-Rosell, P., Daly, J., English, C., Harter, J., Standridge, P., Clouet, B., Diener, E., & Lai, A. Y. (2022b). Balance and harmony in the Gallup World Poll: The development of the Global Wellbeing Initiative module. *International Journal of Wellbeing*, 12(4), 1–19. <https://doi.org/10.5502/ijw.v12i4.2655>
- Lomas, T., Case, B., Cratty, F., & Batson, A. (2025). The dance of East and West: A brief history of an unstable but enduring conceptual partnership. *East-West Studies*, 13, 6–25. <https://doi.org/10.82533/ews.2024.13.15>
- Mangalik, S., Eichstaedt, J. C., Giorgi, S., Mun, J., Ahmed, F., Gill, G., Ganesan, A. V., Subrahmanya, S., Soni, N., Clouston, S. A. P., & Schwartz, H. A. (2024). Robust language-based mental health assessments in time and space through social media. *npj Digital Medicine*, 7(1), Article 109. <https://doi.org/10.1038/s41746-024-01100-0>
- Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., ... Oak, S. (2025). *The AI Index 2025 annual report*. AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University. <https://doi.org/10.48550/arXiv.2504.07139>
- Mata, F., Jaeger, B., & Domingues, I. (2022). Perceptions of farm animal sentience and suffering: Evidence from the BRIC countries and the United States. *Animals*, 12(23), Article 3416. <https://doi.org/10.3390/ani12233416>
- Mayer, T., Warikoo, N., Eliassaf, A., Atzil-Slonim, D., & Gurevych, I. (2024). Predicting client emotions and therapist interventions in psychotherapy dialogues. In Y. Graham & M. Purver (Eds.), *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1463–1477). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.eacl-long.88>
- Mejor, M. (1998). Suffering, Buddhist views of origination of. In E. Craig (Ed.), *Routledge Encyclopedia of Philosophy*. Routledge. <https://doi.org/10.4324/9780415249126-F073-1>
- Mellor, D. J. (2019). Welfare-aligned sentience: Enhanced capacities to experience, interact, anticipate, choose and survive. *Animals*, 9(7), Article 440. <https://doi.org/10.3390/ani9070440>
- Mellor, D. J., Beausoleil, N. J., Littlewood, K. E., McLean, A. N., McGreevy, P. D., Jones, B., & Wilkins, C. (2020). The 2020 five domains model: Including human–animal interactions in assessments of animal welfare. *Animals*, 10(10), Article 1870. <https://doi.org/10.3390/ani10101870>
- Mesquiti, S., Cosme, D., Nook, E. C., Falk, E. B., & Burns, S. (2026). Language-based assessments can predict psychological and subjective well-being. *Communications Psychology*, 4(1), Article 33. <https://doi.org/10.1038/s44271-026-00400-3>
- Michalos, A. C., & Weijers, D. (2017). Western historical traditions of well-being. In R. J. Estes & M. J. Sirgy (Eds.), *The pursuit of human well-being: The untold global history* (pp. 31–57). Springer. https://doi.org/10.1007/978-3-319-39101-4_2
- Mill, J. S. (1863/2010). *Utilitarianism*. Project Gutenberg. (Original work published 1863)

- Miyamoto, Y., Uchida, Y., & Ellsworth, P. C. (2010). Culture and mixed emotions: Co-occurrence of positive and negative emotions in Japan and the United States. *Emotion*, 10(3), 404–415. <https://doi.org/10.1037/a0018430>
- Moore, A. (2019). Hedonism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2019 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2019/entries/hedonism/>
- Muehlhauser, L. (2013). Intelligence explosion FAQ. Machine Intelligence Research Institute.
- Müller, S. M., & Bittermann, A. (2025). The emotional climate of academia: Exploring social media data as an indicator of well-being. *Journal of Happiness Studies*, 26, Article 96. <https://doi.org/10.1007/s10902-025-00902-7>
- Murtin, F., & Salomon-Ermel, M. (2024). Nowcasting subjective well-being with Google Trends: A meta-learning approach (OECD Papers on Well-being and Inequalities No. 27). OECD Publishing. <https://doi.org/10.1787/cbdfb5d9-en>
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4), 435–450. <https://doi.org/10.2307/2183914>
- Nagel, T. (1979). *Mortal questions*. Cambridge University Press.
- Nandini, D., Yadav, J., Singh, V., Mohan, V., & Agarwal, S. (2025). An ensemble deep learning framework for emotion recognition through wearable devices multi-modal physiological signals. *Scientific Reports*, 15(1), Article 17263. <https://doi.org/10.1038/s41598-025-99858-0>
- Neugarten, B. L., Havighurst, R. J., & Tobin, S. S. (1961). The measurement of life satisfaction. *Journal of Gerontology*, 16(2), 134–143. <https://doi.org/10.1093/geronj/16.2.134>
- New Zealand Government. (2025). Public Finance Act 1989, s 26NB Wellbeing report. New Zealand Legislation. <https://www.legislation.govt.nz/act/public/1989/0044/latest/LMS375590.html>
- New Zealand Veterinary Association. (2018, November). Position statement: Sentience. <https://nzva.org.nz/positions-advocacy/position-statements/sentience/>
- Ng, L. H. X., & Carley, K. M. (2025). A global comparison of social media bot and human characteristics. *Scientific Reports*, 15, Article 13649. <https://doi.org/10.1038/s41598-025-96372-1>
- Organisation for Economic Co-operation and Development. (2013). *OECD guidelines on measuring subjective well-being*. OECD Publishing. <https://doi.org/10.1787/9789264191655-en>
- Organisation for Economic Co-operation and Development. (2025). Updating the OECD guidelines on measuring subjective well-being. OECD Publishing. <https://doi.org/10.1787/9203632a-en>
- Organisation for Economic Co-operation and Development. (2026). The OECD Going Digital Measurement Roadmap 2026. OECD Publishing. <https://doi.org/10.1787/b455e132-en>
- People's Republic of China Government. (2011). The 12th Five-Year Plan for national economic and social development of the People's Republic of China (2011–2015) [English translation]. <https://chinaenergyportal.org/wp-content/uploads/2016/08/Full-Translation-of-the-12th-Five-Year-Plan.pdf>

- People's Republic of China Government. (2016). The 13th Five-Year Plan for economic and social development of the People's Republic of China (2016–2020) (Compilation and Translation Bureau, Trans.). Central Compilation & Translation Press.
<https://en.ndrc.gov.cn/policies/202105/P020210527785800103339.pdf>
- People's Republic of China Government. (2021). Outline of the 14th Five-Year Plan for national economic and social development and long-range objectives for 2035 (Etcetera Language Group, Trans.; B. Murphy, Ed.). Center for Security and Emerging Technology.
<https://cset.georgetown.edu/publication/china-14th-five-year-plan/>
- Pew Research Center. (2025, August 5). God, spirits and the natural world: What a 35-country survey finds about spiritual beliefs.
<https://www.pewresearch.org/religion/2025/08/05/god-spirits-and-the-natural-world/>
- Phang, J., Lampe, M., Ahmad, L., Agarwal, S., Fang, C. M., Liu, A. R., Danry, V., Lee, E., Chan, S. W. T., Pataranutaporn, P., & Maes, P. (2025). Investigating affective use and emotional well-being on ChatGPT. *arXiv*. <https://arxiv.org/abs/2504.03888>
- Rākete, P. (2017). In human parasites: Posthumanism and Papatūānuku. *PARSE Journal*, 6.
<https://doi.org/10.13140/RG.2.2.28935.32161>
- Ryan, R. M., & Deci, E. L. (2001). On happiness and human potentials: A review of research on hedonic and eudaimonic well-being. *Annual Review of Psychology*, 52, 141–166.
<https://doi.org/10.1146/annurev.psych.52.1.141>
- Ryff, C. D. (1989). Happiness is everything, or is it? Explorations on the meaning of psychological well-being. *Journal of Personality and Social Psychology*, 57(6), 1069–1081.
<https://doi.org/10.1037/0022-3514.57.6.1069>
- Ryff, C. D., Boylan, J. M., & Kirsch, J. A. (2021). Eudaimonic and hedonic well-being: An integrative perspective with linkages to sociodemographic factors and health. In M. T. Lee, L. D. Kubzansky, & T. J. VanderWeele (Eds.), *Measuring well-being: Interdisciplinary perspectives from the social sciences and the humanities* (pp. 92–135). Oxford University Press.
<https://doi.org/10.1093/oso/9780197512531.003.0005>
- Sametoğlu, S., Pelt, D. H. M., Eichstaedt, J. C., Ungar, L. H., & Bartels, M. (2024). The value of social media language for the assessment of wellbeing: A systematic review and meta-analysis. *The Journal of Positive Psychology*, 19(3), 471–489. <https://doi.org/10.1080/17439760.2023.2218341>
- Sen, A. (1992). *Inequality reexamined*. Harvard University Press.
- Shen, S., Qi, W., Zeng, J., Li, S., Liu, X., Zhu, X., Dong, C., Wang, B., Shi, Y., Yao, J., Wang, B., Lou, X., Gu, S., Li, P., Wang, J., Jiang, G., & Cao, S. (2025). Passive sensing for mental health monitoring using machine learning with wearables and smartphones: Scoping review. *Journal of Medical Internet Research*, 27, Article e77066. <https://doi.org/10.2196/77066>
- Shiffman, S., Stone, A. A., & Hufford, M. R. (2008). Ecological momentary assessment. *Annual Review of Clinical Psychology*, 4, 1–32. <https://doi.org/10.1146/annurev.clinpsy.3.022806.091415>
- Singer, P. (1975). *Animal liberation: A new ethics for our treatment of animals*. New York Review; distributed by Random House.

- Singer, P. (1981). *The expanding circle: Ethics and sociobiology*. Farrar, Straus and Giroux.
- Singh, S. H., Jiang, K., Bhasin, K., Sabharwal, A., Moukaddam, N., & Patel, A. B. (2024). RACER: An LLM-powered methodology for scalable analysis of semi-structured mental health interviews. In L. Peled-Cohen, N. Calderon, S. Lissak, & R. Reichart (Eds.), *Proceedings of the 1st Workshop on NLP for Science (NLP4Science)* (pp. 73–98). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.nlp4science-1.8>
- Smith, C., Frieling, M., Percival, H., & Mahoney, J. (2025). Globally inclusive measures of subjective well-being: Updated evidence to inform national data collections (OECD Papers on Well-being and Inequalities No. 35). OECD Publishing. <https://doi.org/10.1787/bd72752a-en>
- SPCA. (n.d.). The Five Domains [Position statement]. <https://www.sPCA.nz/advocacy/position-statements/article/the-five-domains>
- Standing Committee of the National People's Congress. (2021). Personal Information Protection Law of the People's Republic of China. National People's Congress of the People's Republic of China.
- Stefon, M. (2013, September 23). Wuwei. *Encyclopaedia Britannica*. <https://www.britannica.com/topic/wuwei-Chinese-philosophy>
- Stewart, G. (2020). *Māori philosophy: Indigenous thinking from Aotearoa*. Bloomsbury Academic.
- Stiglitz, J. E., Sen, A., & Fitoussi, J.-P. (2009). Report of the Commission on the Measurement of Economic Performance and Social Progress. Commission on the Measurement of Economic Performance and Social Progress.
- Te Awa Tupua (Whanganui River Claims Settlement) Act 2017, Public Act 2017 No. 7 (N.Z.). <https://www.legislation.govt.nz/act/public/2017/0007/latest/DLM6831461.html>
- Teles, A. S., Chaturvedi, J., Wang, T., Scazufca, M., Msosa, Y., Stahl, D., & Roberts, A. (2025). Generative multimodal large language models in mental health care: Applications, opportunities, and challenges. *PLOS Mental Health*, 2(11), Article e0000488. <https://doi.org/10.1371/journal.pmen.0000488>
- Tison, G. H., Avram, R., Kuhar, P., Abreau, S., Marcus, G. M., Pletcher, M. J., & Olgin, J. E. (2020). Worldwide effect of COVID-19 on physical activity: A descriptive study. *Annals of Internal Medicine*, 173(9), 767–770. <https://doi.org/10.7326/M20-2665>
- Tseriotou, T., Chim, J., Klein, A., Shamir, A., Dvir, G., Ali, I., Kennedy, C., Kohli, G. S., Hills, A., Zirikly, A., Atzil-Slonim, D., & Liakata, M. (2025). Overview of the CLPsych 2025 Shared Task: Capturing mental health dynamics from social media timelines. In *Proceedings of the 10th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2025)* (pp. 193–217). *Association for Computational Linguistics*. <https://doi.org/10.18653/v1/2025.clpsych-1.16>
- Trull, T. J., & Ebner-Priemer, U. W. (2009). Using experience sampling methods/ecological momentary assessment (ESM/EMA) in clinical assessment and clinical research: Introduction to the special section. *Psychological Assessment*, 21(4), 457–462. <https://doi.org/10.1037/a0017653>
- UAE Cabinet. (2019, June 9). UAE Cabinet approves National Strategy for Wellbeing 2031. <https://uaecabinet.ae/en/details/news/uae-cabinet-approves-national-strategy-for-wellbeing-2031>

United Nations Economic Commission for Europe. (2025). *Guidelines on measurement of well-being*. United Nations. <https://doi.org/10.18356/9789211544701>

United Nations General Assembly. (2015). Transforming our world: The 2030 Agenda for Sustainable Development (A/RES/70/1). United Nations. <https://docs.un.org/en/A/RES/70/1>

VanderWeele, T. J., Johnson, B. R., Bialowolski, P. T., Bonhag, R., Bradshaw, M., Breedlove, T., Case, B., Chen, Y., Chen, Z. J., Counted, V., Cowden, R. G., de la Rosa, P. A., Felton, C., Fogleman, A., Gibson, C., Grigoropoulou, N., Gundersen, C., Jang, S. J., Johnson, K. A., ... Yancey, G. (2025). The Global Flourishing Study: Study profile and initial results on flourishing. *Nature Mental Health*, 3(6), 636–653. <https://doi.org/10.1038/s44220-025-00423-5>

Waldinger, R. (2015, November). What makes a good life? Lessons from the longest study on happiness [Video]. TED Conferences. https://www.ted.com/talks/robert_waldinger_what_makes_a_good_life_lessons_from_the_longest_study_on_happiness

Wang, R. R. (n.d.). Yinyang (yin-yang). *Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/yinyang/>

Wang, Z., & Wang, Y. (2025). Emotion recognition based on multimodal physiological electrical signals. *Frontiers in Neuroscience*, 19, Article 1512799. <https://doi.org/10.3389/fnins.2025.1512799>

Weijers, D. (2011). Hedonism. In J. Fieser & B. Dowden (Eds.), *Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/hedonism/>

Weijers, D., & Jarden, A. (2021). The International Journal of Wellbeing: Ten years on and the next ten years. *International Journal of Wellbeing*, 11(1), 1–7. <https://doi.org/10.5502/ijw.v11i1.1577>

Weijers, D., & Morrison, P. S. (2018). Wellbeing and public policy: Can New Zealand be a leading light for the wellbeing approach? *Policy Quarterly*, 14(4), 3–12. <https://doi.org/10.26686/pq.v14i4.5144>

Weijers, D. M., & Mukherjee, U. (2016). Living standards, wellbeing, and public policy. The New Zealand Treasury. <https://www.treasury.govt.nz/publications/background/living-standards-wellbeing-and-public-policy>

Weijers, D., & Munn, N. (2024, May 9). AI companions can relieve loneliness – but here are 4 red flags to watch for in your chatbot ‘friend’. *The Conversation*.

Wellbeing Economy Alliance. (n.d.). Wellbeing Economy Governments (WEGo). Retrieved May 29, 2026, from <https://weall.org/wego>

Wendler, C., Veselovsky, V., Monea, G., & West, R. (2024). Do Llamas work in English? On the latent language of multilingual transformers. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 15366–15394). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.acl-long.820>

Wild Animal Initiative. (n.d.). Wild Animal Initiative. Retrieved June 3, 2026, from <https://www.wildanimalinitiative.org/>

World Economic Forum. (2019, November). China is waking up to data protection and privacy. Here's why that matters.

<https://www.weforum.org/agenda/2019/11/china-data-privacy-laws-guideline/>

Wrzus, C., & Neubauer, A. B. (2023). Ecological momentary assessment: A meta-analysis on designs, samples, and compliance across research fields. *Assessment*, 30(3), 825–846.

<https://doi.org/10.1177/10731911211067538>

Xinhua. (2025, June 30). CPC grows stronger as membership exceeds 100M. State Council Information Office of the People's Republic of China.

https://english.scio.gov.cn/pressroom/2025-06/30/content_117954482.html

Yu, L., Ge, Y., Ansari, S., Imran, M., & Ahmad, W. (2025). Multimodal sensing-enabled large language models for automated emotional regulation: A review of current technologies, opportunities, and challenges. *Sensors*, 25(15), Article 4763. <https://doi.org/10.3390/s25154763>

Zuboff, S. (2019, February 15). Shoshana Zuboff on surveillance capitalism [Video]. YouTube. VPRO Documentary. <https://www.youtube.com/watch?v=hIXhnWUmMvw>

Appendix

Appendix 1

Google Ngram results:

https://books.google.com/ngrams/graph?content=wellbeing%2Cwell-being&year_start=1800&year_end=2022&corpus=en&smoothing=3&case_insensitive=true

Google Trends results:

<https://trends.google.com/explore?q=wellbeing%2Cwell-being&date=all&geo=NZ>

Google Scholar 1980:

https://scholar.google.com/scholar?hl=en&as_sdt=0%2C5&as_ylo=1980&as_yhi=1980&q=wellbeing+OR+well-being&oq=wellbeing

Google Scholar 2020:

https://scholar.google.com/scholar?hl=en&as_sdt=0%2C5&as_ylo=2020&as_yhi=2020&q=wellbeing+OR+well-being&oq=wellbeing

Appendix 2

This appendix provides additional operational detail for the LLM-assisted personal chatbot conversation analysis described in Section 5.2. It clarifies the prompt structure, output format, and script logic used. The raw conversation data, full output spreadsheet, and complete Python script are not reproduced because they contain private personal information or unnecessary implementation detail. Full script details can be provided upon request.

1. Data preparation

The exported ChatGPT data were first converted into plain-text conversation transcripts. Each transcript represented one conversation and retained only the conversation title, message timestamps, and speaker labels. Messages were labelled as either “User:” or “Chatbot:”. Other export metadata were removed.

The subsequent analysis treated user-authored messages as the evidential basis for coding. Chatbot replies were retained in the transcript only as contextual information, for example to clarify what the user was responding to or asking about.

2. Plausibility-screening prompt

Before detailed analysis, each conversation was screened using a conservative theme-by-theme plausibility prompt. This stage was used to reduce unnecessary processing while avoiding the exclusion of borderline or ambiguous conversations. The model was given the full set of theme prompts (see section 5.2) and asked to judge whether each theme could plausibly be present in the user-authored messages.

The operational instructions for this stage were:

You are doing a conservative theme-by-theme plausibility screen of ChatGPT conversation transcripts.

Task:

For each theme, decide whether the theme could plausibly be present in the USER messages of the conversation.

Important:

- *Analyse USER messages only.*
- *Ignore Chatbot messages as evidence.*
- *Use the full prompt for each theme.*
- *This is only a plausibility screen, not detailed coding.*
- *For each theme, return exactly one of: yes, no, unsure.*
- *Use "yes" if the theme is clearly or plausibly present.*
- *Use "no" only if the theme is clearly absent.*
- *Use "unsure" if there is weak, indirect, ambiguous, or borderline evidence.*
- *Be conservative: avoid false negatives. If in doubt, use "yes" or "unsure".*
- *Do not provide details.*
- *Return valid JSON only.*

A conversation proceeded to detailed analysis unless every theme was returned as “no” in the plausibility screen. If all themes were returned as “no”, the conversation was recorded as containing no theme at this stage.

3. Detailed-analysis prompt

For conversations that passed the plausibility screen, the model was then asked to provide conversation-level coding for each theme. The operational instructions for the detailed stage were:

You are doing conversation-level analysis of ChatGPT conversation transcripts for wellbeing-related themes.

Task:

Assess whether the USER messages in a conversation contain each theme, and briefly describe the nature of the theme when present.

Important:

- *Analyse USER messages only.*
- *Ignore Chatbot messages as evidence, except where they help interpret what the user is referring to.*
- *This is conversation-level analysis, not message-level coding.*
- *For each theme, return a status and details.*
- *status must be exactly one of: present, absent, unsure.*
- *Use "present" if one or more user messages in the conversation clearly contain the theme.*
- *Use "absent" if no user messages contain the theme.*
- *Use "unsure" only when the evidence is genuinely ambiguous.*

- *If status is "present", details should briefly describe the nature of the theme in this conversation.*
- *If status is "unsure", details should briefly state what is ambiguous.*
- *If status is "absent", details must be an empty string.*
- *Details should be concise: usually one sentence, maximum three sentences.*
- *Do not quote long passages.*
- *Do not include message-level classifications.*
- *Return valid JSON only.*

The full list of theme prompts was inserted below these instructions.

4. Output schema

The model was instructed to return structured JSON rather than free-form prose. The detailed analysis used the following output structure:

```
{
  "themes": {
    "theme_slug": {
      "status": "present|absent|unsure",
      "details": "brief description if present or unsure; empty string if absent"
    }
  }
}
```

The final CSV output contained one row per conversation. For each theme, the spreadsheet included a categorical status column and a qualitative details column. The plausibility-screening outputs were also retained, allowing later inspection of which themes triggered detailed analysis.

5. Simplified script logic

The script followed the logic below:

For each conversation transcript:
Extract conversation metadata:
file name
conversation title
source export file
creation date

Extract user-authored messages for the plausibility screen.

For each of the 14 themes:

Ask whether the theme could plausibly be present.

Return yes, no, or unsure.

If every theme is marked no:

Record all detailed theme statuses as absent.

Leave all detail fields blank.

Otherwise:

Submit the full conversation transcript for detailed analysis.

For each theme:

Return status: present, absent, or unsure.

Return brief details if present or unsure.

Append the structured result to a JSONL file.

Rebuild the CSV output file from the JSONL file.

Skip already completed conversations if the script is restarted.

6. Final synthesis step

After the conversation-level outputs were generated, the results were further processed into a condensed timeline document. This document reorganised the model-generated theme outputs into a chronological format intended to support higher-level interpretation over time. The timeline was derived directly from the prior spreadsheet output with no new information added.

7. Error handling and reproducibility

The script included retry handling for API failures and resume support, so that completed conversations were not reprocessed after interruption. Intermediate outputs were written incrementally to reduce the risk of data loss during a long run.

This procedure was designed for transparency and exploratory analysis rather than formal psychometric validation. The results remain dependent on the model used, the wording of the prompts, and the absence of independent human coding. The internal reasoning process of the model was not inspectable. For these reasons, the outputs should be interpreted as model-generated thematic interpretations of the text, rather than as validated measurements of wellbeing.

Full script is available on request.

Note: The above appendix summarising the analysis process was written with the assistance of ChatGPT following the creation of the protocol.

Appendix 3

Prompt:

The attached file contains a timeline of wellbeing indications inferred from a user's chatbot conversation history. Give an overall assessment of the user's wellbeing over time.

ChatGPT response:

Overall assessment

The timeline suggests *mixed but generally resilient wellbeing over time*. The dominant pattern is not deterioration; it is *high agency under recurring stress*. The user repeatedly seeks advice, makes plans, reflects on values, tries to solve practical problems, and maintains social, professional, and ethical commitments.

[redacted]

Timeline pattern

Early 2023: Exploratory and future-oriented, but with uncertainty around career, role, social approval, relationships, and animal-health concerns. Strong goal-seeking is already present.

[redacted]

Oct–Dec 2023: October–November look comparatively more stable/positive, with travel, learning, social connection, professional purpose, and humour. December shows a sharper stress cluster around Mandarin/study pace, workload, frustration, and inadequacy.

[redacted]

Oct–Dec 2024: Highest concentration of stress indicators, especially animal-health concerns, community conflict, pain/discomfort, and emotionally charged topics. This does not look like collapse; it looks like a high-load period with continued planning and values-based reasoning.

Jan 2025: Suggests reflective transition: pride, fatigue, disheartenment, authenticity concerns, and some increased confidence. Still purposeful.

Main protective factors

The strongest protective signs are *future orientation, help-seeking, planning, ethical reflection, social connection, and purpose*. The user often tries to convert distress into action: drafting messages, seeking professional input, making plans, setting boundaries, clarifying values, and pursuing meaningful projects.

Main vulnerability signals

Recurring vulnerability appears as *worry, anxiety, self-criticism, frustration, relational uncertainty, and workload pressure*. There are repeated concerns about doing the right thing, being understood, balancing ambition with relationships, and managing emotionally demanding animal-care or social situations.

[redacted]

Bottom line: the user appears consistently thoughtful, purposeful, socially connected, and adaptive, but with recurrent periods of moderate emotional strain. Wellbeing looks *variable rather than steadily declining*, with the heaviest strain concentrated in parts of 2024.

Gemini response:

An analysis of the timeline data provided in the file (comprising 2,620 conversational interactions primarily spanning from the first quarter of 2023 to early 2025) offers a comprehensive view of the user's psychological state over time.

To provide a rigorous assessment, it is helpful to divide wellbeing into two standard philosophical and psychological categories: eudaimonic wellbeing (which focuses on meaning, purpose, and self-actualisation) and hedonic wellbeing (which relates to the day-to-day balance of positive and negative emotions). Evaluated through these lenses, the data indicates a stable, highly functioning individual whose wellbeing is robust and heavily anchored in professional and personal meaning.

Eudaimonic Flourishing: Purpose and Accomplishment

The most prominent feature of the dataset is a consistent, strong orientation towards goal pursuit and value alignment. Frequency analysis of the underlying text reveals that terms such as 'desire' (occurring 473 times), 'goal' or 'goals' (463 times), 'meaning', and 'purpose' (267 times combined) heavily dominate the user's conversational history.

[redacted]

Hedonic Stability and Emotional Variance

Hedonic wellbeing evaluates the momentary presence of positive affect versus negative stress. Sentiment analysis of the conversational details yields an average polarity score of +0.096 across the entire dataset, pointing to a baseline emotional state that is consistently, albeit mildly, positive. Positive emotional markers reliably match or outpace negative markers in every recorded quarter.

Negative emotions and stress markers are naturally present, peaking slightly at 12.9% of interactions in the second quarter of 2024. However, qualitative examination of these specific entries reveals that they are generally contextual rather than symptomatic of systemic psychological distress. Stressors frequently correlate with specific professional duties—such as managing a dog's anxiety or a cat's physical pain in a veterinary context—and navigating interpersonal boundaries. This variability represents an expected, adaptive response to the complexities of a demanding professional life, rather than a decline in overall mental health.


[redacted]

Summary

The data within the file points to a steady, positive trajectory of wellbeing. The user demonstrates strong eudaimonic flourishing, driven by active academic, professional, and ethical engagements. While expected hedonic fluctuations occur in response to routine life stressors and professional challenges, the overarching psychological state remains highly stable, grounded by a clear sense of purpose and consistent social connectedness.

Claude response:

This is a rich dataset. Let me parse the dates and structure for a cleaner analysis. Good — now I have a comprehensive picture. Let me build the visualisation.



Here's a narrative assessment across the four wellbeing dimensions the data captures:

Overall picture: purposeful, socially engaged, and emotionally variable

The most consistent signal across the full corpus is purpose and meaning — this dimension stays elevated throughout, hovering around 35–50% of entries, and never collapses. The user is persistently goal-oriented, mission-driven, and engaged in work that they find meaningful (veterinary practice in Cambodia, animal welfare, research, language learning, cross-cultural connection). That's a stable anchor for wellbeing.

Three distinct phases emerge:

Early 2023 (January–May) shows a baseline with moderate positive affect and moderately elevated negative affect signals — stress, worry, and relational friction are present, but so is clear purpose and social engagement. The conversations suggest someone navigating transitions (geographic, professional, relational) with resilience.

Mid-2023 is the corpus's emotional high point. June–July 2023 sees a marked rise in positive affect, with entries referencing experiences described explicitly as "the best time of my life," transformative Cambodia and Beijing experiences, community warmth, and a strong sense of meaning. August 2023 is a notable dip — negative affect spikes sharply (0.41, the highest in the dataset), and social connection signals also rise, suggesting a period of relational turbulence or transition stress. This pattern is consistent with re-entry experiences after an intense period abroad.

Late 2023 into early 2024 is the richest period emotionally. November 2023 has the highest positive affect reading (0.39) with low negative affect (0.05) — a conspicuously good month. [redacted]. February 2024 is notably positive again (0.41 positive, high social connection), with birthday celebrations, Taiwan experiences, and warm social moments.

[redacted]

Late 2024 (September–December) shows a sustained dip in positive affect, with November 2024 (0.11) and December 2024 (0.06) being the lowest positive readings in the full dataset. Negative affect is also moderate-to-low in this period, suggesting less emotional intensity rather than acute distress — this may reflect a more inward, focused period (consistent with thesis writing). Purpose signals remain steady.

January 2025 shows a recovery in positive affect (0.27), which is encouraging as an endpoint.

Key themes running throughout:

The user shows persistent eudaimonic strengths — clear values, a sense of mission, intellectual engagement, and cross-cultural curiosity. Hedonic wellbeing is more variable, often tied to relational and romantic experiences. Social connection fluctuates considerably, peaking during travel and community-embedded periods (Cambodia, Beijing, Taiwan) and dropping during more solitary or study-focused phases. The two most emotionally destabilising themes across the corpus are [redacted].

One methodological note worth flagging: the scoring here reflects keyword density from the extracted evidence rather than sentiment weights — so months with high activity in both directions (like May 2024) score high on both dimensions simultaneously. That's actually a feature of this kind of corpus analysis, not a flaw: it captures emotional complexity rather than reducing to a single valence.