The Easterlin Paradox reconsidered: methodological challenges in answering whether money buys long-term happiness.

Master Thesis
July 2024

Bram van den Beld
M Economics
Tilburg University
2044059

Supervisor
prof. dr. Johan Graafland

Second reader
em. prof. dr. Lans Bovenberg

Word count: +/- 19800
(excluding footnotes)

Abstract

This thesis critically examines the Easterlin Paradox, which states that, in the long run, trends in happiness and economic growth are unrelated. It explores this subject by deconstructing the paradox into three stylized statements and discusses the current scientific state of affairs regarding each. The mechanisms proposed to explain the paradox are reflected upon, and attention is given to underexplored mechanisms. In an effort to advance the debate, economic problems are discussed that can explain why there is a lack of consensus on the existence and importance of the paradox, including the issue of causality and the measurement of income and happiness. The thesis considers the relevance for policy and argues against the use of the paradox in propagating degrowth. Finally, it provides a concise agenda with unresolved matters to be addressed by future research.

Contents

Introduction

Suppose that you would be offered **a million dollars** to read this thesis. Would you do it? Why?

The offer would most probably lead to suspicion and some frowned eyebrows, why would you receive a million dollars to perform a task that - since you began to read the first sentence - you were planning on doing anyway?
Assume that the suspicion could be overcome and that you could be convinced that the offer is in fact extremely serious: there are no hidden traps or false promises. I suspect that people could not wait to make the deal.

That is because people care deeply about money. We go to great lengths to attain it. We move to a different city to be able to get a better paying job and play the lottery in the hopes of winning a big prize that would change our lives forever. Businesses compete fiercely to be able to make a profit, and spend a huge amount on research and development to make sure this will also be the case in the future.

Unsurprisingly, money seems highly important in day to day life. Money makes us able to consume things that we feel we need and achieve things that we deem important. Money therefore seems to be closely related to happiness. It seems thus a reasonable statement to assume that *more* money makes us *more* happy. This *monotonicity assumption* - that more is better - lies at the basis of the work of utility theory, which is an important part of economics. The hypothesis that raising the income of all people would raise the happiness of all seems thus well grounded.

Interestingly enough, there are economists that disagree with that idea. The most prominent economist to do so is Richard Easterlin, who is known for the paradox that is named after him: The Easterlin Paradox. This paradox states that over the long run, countries whose economies grow faster do not seem to get happier faster. This thesis aims to dive deeper into the matter and see what the Easterlin Paradox can reveal about the relationship between happiness and income.

Outline

This master thesis aims to answer the question:

> *"What is the current state of research into the Easterlin Paradox and what are the remaining economic issues to tackle to help the debate move one step further?"*

It does so by dealing with the fundamentals in chapter 1. Here, the focus will lie on what the paradox does and does not entail. In chapter 2 then, there will be zoomed in on the three statements Easterlin developed to find out what the most recent scientific evidence on these matters is. Chapter 3 will discuss the proposed mechanisms behind the paradox and adds a discussion about underexposed mechanisms. Chapter 4 contains economic problems that remain around the paradox, including how to measure happiness and income, the issues around causality and the 'happiness' scales Easterlin uses. Chapter 5 applies the insights that the thesis has found on policy. It does so by reflecting on whether happiness is worth maximising and by looking at the case for degrowth. Chapter 6 constructs a research agenda with problems to be resolved in future research. At last, in chapter 7, concluding thoughts will be shared.

1: The Easterlin paradox: the fundamentals

This chapter will set the foundation for what is coming. It aims to clear up what the paradox does *(and does not)* entail. This is especially important since Easterlin himself seems to remark that the paradox gets misrepresented often, both in non-scientific and scientific literature, even in recent times. This misconception might partly explain why the debate is still going on fiercely. Let's clear up the ambiguity and make sure that we are all on the same page.

In 1974, Richard Easterlin published the influential paper 'Does Economic Growth Improve the Human Lot? Some Empirical Evidence.' In this paper, Easterlin was one of the first ones to let go of the Pigouvian presumption that *"changes in economic welfare indicate changes in social welfare in the same direction"* (Abramovitz, 1959). This presumption was still a dominant view at the time, for example in Nordhaus and Tobin (1972, p.532): *"Although GNP and other national income aggregates are imperfect measures of welfare, the broad picture of secular progress which they convey remains…"* Instead, Easterlin tried to look at the empirical evidence on whether economic growth is positively associated with social welfare, i.e. human happiness.

The original question - does economic growth improve the human lot - is, according to Easterlin, societal in nature. The answer to the question indicates whether a government who aims to maximise collective well-being should seek to increase GDP. (Easterlin & O'Connor, 2020, p.3-4).

Exactly fifty years have passed since the original paper was published, and there has been put extensive attention, effort and research into the subject. In these fifty years, the exact content and demarcation of the paradox have changed slightly. I do not aim to give a complete historical overview of this change, but do go into some detail of this in chapter 2.

A correct **contemporary** representation of the Easterlin paradox consists of the following three statements:

1. Within countries at a given point in time, there is a noticeable positive association between income and happiness at the micro *(individual)* level

2. Between countries at a given point in time, there is a positive association between income and happiness at the macro *(national)* level

3. Long-term trends in average happiness and average income are not significantly related.

The most important statement is statement 3, this is the new insight that Easterlin brought to the table and that got the most attention. What is crucial to note is that statement 3 on its own is **not** a paradox. The Easterlin paradox is a paradox only because the combination of statement 1 + 2 *(cross-sectional data)* seems to conflict with statement 3 *(time-series data)*.

Or, as Easterlin and O'Connor (2020) write:

> *"The Easterlin Paradox states that at a point in time happiness varies directly with income, both among and within nations, but over time happiness does not trend upward in correspondence with income growth." (p.2)*
>
> *"The paradox is the contradiction between observations on the relation of happiness to income at a point in time … and evidence on happiness and income over time …." (p.2)*

Using non-scientific language: the paradoxical element is that richer people are happier than poorer people and richer countries are happier than poorer countries, but over the course of time, countries which grow faster do not seem to get happier faster. (Plant, 2022, p.3)

Other 'basic' remarks are that when dealing with the paradox, 'income' always refers to **real** income, so a reflection of what money will buy; the quantity of goods and services. This makes 'layman explanations' for the paradox like *"rising income is not associated with more happiness since prices have risen as well"* invalid. Real income takes inflation into account. Following from this, when comparing countries, proper analysis uses PPP *(purchasing power parity)* data. This adjusts for differences in the cost of living between countries.

Furthermore, to depict the relationship between income and happiness, the literature mainly uses a measure of happiness and the **logarithm** of income. This has to do with how we perceive changes in income and how intense these changes feel. The effect an extra dollar has on happiness goes down, the more income one has. For more info on this, one could read Easterlin (2001, p.468) or Stevenson and Wolfers (2008, p.5) Though, thorough discussion of the correct functional form is quite rare.
Using the logarithm of income makes 'layman explanations' for the paradox like *"rising income is not associated with more happiness since the extra income one gets if one is already fairly rich does not add much"* invalid. This - on itself not incorrect - statement on the decreasing marginal benefit of income is already taken into account when using the logarithm of income.

Covering these more basic aspects can prevent a significant part of debate in *(popular)* science literature. Easterlin does **not** claim that income does not matter for happiness. Finding evidence that people do get happier from more income or that richer nations are happier cannot serve as evidence for the non-existence of the paradox. In order to credibly state that the paradox does not exist, one has to tackle statement 3 ('the Easterlin statement', if you will). I find support for this view in Kaiser and Vendrik (2019), who conclude that reliable tests of the paradox must allow for trends in happiness and growth.

These two authors also point to another misconception, namely that the paradox claims that happiness is constant over the long run. This need not be the case, and finding that happiness in fact is not constant over time therefore cannot serve as sufficient counter-evidence against the paradox. Countries can have rising, constant, or falling trends in happiness. The paradox only claims that since trends in average happiness and average income are unrelated, steeper uptrends in a country's income are not accompanied by greater growth in happiness. (Easterlin and O'Connor, 2020, p.5)

A final critical remark that is put to the table in Kamilcelebi and Veenhoven (2024, p.1) is that over time, the content of the Easterlin Paradox has shifted somewhat. For example, Easterlin's recent focus on the *long run* was not present to that extent in earlier work. To take the *contemporary* version of the paradox and then criticising the *older* papers that shined their light on the paradox need not be fair, since these papers might have commented on the somewhat different paradox *(both in content as in focus)* as existed at the time. An impartial analysis of both scientific 'camps' – paradox believers and paradox deniers -, needs to look at the most recent works. I aim to do so, especially in the last part of section 2.3.

Now that we have handled the basics and made sure that some misconceptions will not muddy the waters in the further discussion, it is time to dive deeper into the three statements.

2: Research on the three key statements

This chapter will provide a detailed overview of the three statements and aims to show the scientific state of affairs. A complete historical overview of research into the three statements will not be given, though I deem some historical context of how the consensus view on the topic came to be and changed over time helpful. The focus of this chapter will thus lie on the most recent findings.

2.1: Within countries, micro level

The focus in this section will lie on statement 1: "Within countries at a given point in time, there is a noticeable positive association between income and happiness at the micro *(individual)* level."

To research this statement, one basically asks whether greater happiness goes with higher income. At the time that Easterlin wrote his paper, quite some attention had already been put on this question and there existed numerous studies that contained empirical evidence to answer the question *(for example: Robinson and Shaver (1969), Gurin et al. (1960), Wilson (1967), Bradburn (1969))*. A big chunk of empirical work was focused on the United States, though, plenty of countries in Europe, Asia, Africa and Latin America were also represented.

Looking at this evidence, together with 'new' data from 1970, resulted in Easterlin stating: "There is clear indication that income and happiness are positively associated." (1974, p.99).

This indication has most definitely survived the test of time. In the fifty years that have passed, both data quality and quantity have steadily improved. The critique that the available data was not diverse or detailed enough might have been valid in the 1970's, but has since then been rendered ungrounded. The evidence is clear and practically unequivocally so: at a given point in time, within a country, there is a noticeable positive association between personal income and happiness.

This visualisation by Our World in Data shows the evidence of 106 countries. It does so by dividing every country into five groups of equal population based on their income *(income quintiles)*, these incomes are transformed to PPP values to enhance comparability. The lines then show, for every country, at a given point in time, the mean self-reported life satisfaction and mean income of people at a given income quintile.

Figure 1: *"Self-reported life satisfaction across the income distribution"* [a]

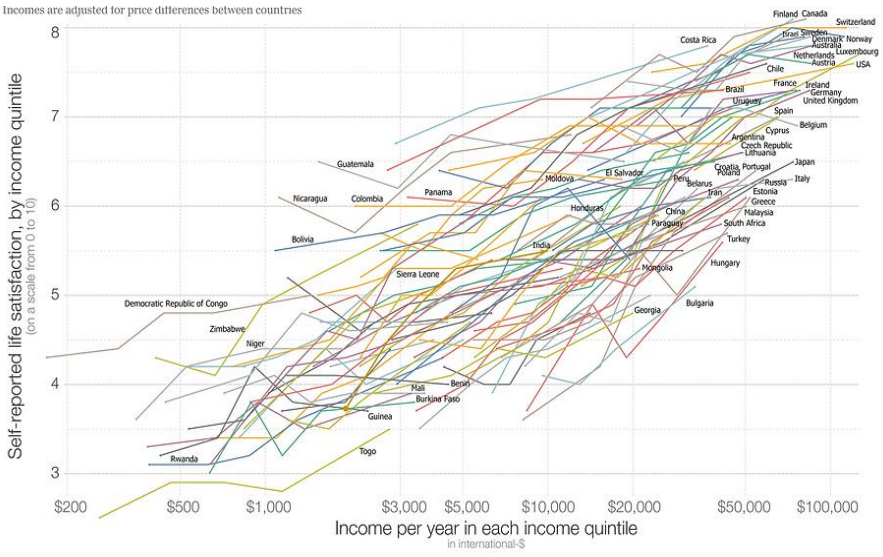

[a] **Source:** *(Our World in Data, Max Roser)* Self-reported life satisfaction across the income distribution. For each country, incomes have been split into five groups with the same number of people *(income quintiles).* Lines show, country by country, the average self-reported life satisfaction on a scale from 1 to 10 for people at a given income quintile, at a given point in time. Data is from 2008 to 2014, depending on the country. Data sources: World Bank, Gallup World Poll

Due to the fact that it contains so much data, readability may be a bit hampered. To be able to get the gist faster, this second visualisation by the Economist might be preferable to some:

Figure 2: *"Self-reported life satisfaction across the income distribution"* [b]

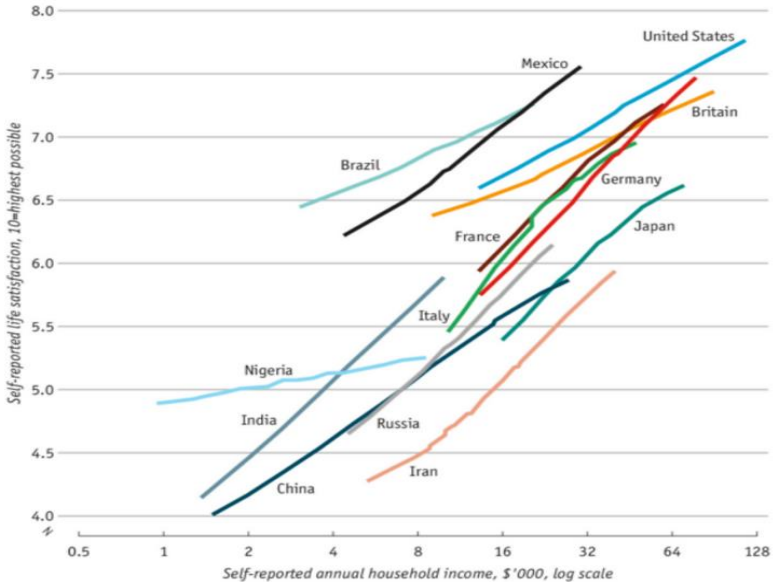

[b] **Source:** *(The Economist, 2013)* Lines show, country by country, the average self-reported life satisfaction on a scale from 1 to 10 for people at a given self-reported household income at a given point in time. Data from Stevenson & Wolfers (2013). Data is from 2010 to 2012, depending on the country.

The take-away from these figures is that as income goes up by a given percentage *(log plot),* life satisfaction scores seem to go up by a more or less fixed number of points. So, within a country, people who earn more money report a higher life satisfaction, but this happens on a log-linear scale: two people earning $4,000 and $12,000, would be expected to show the same difference in well-being as two people earning $12,000 and $36,000.

What is important to note is that these graphs show a simple bivariate correlation, and need not represent a causal link from income to happiness. If anything, one could defend the case that happier people are more likely to be successful in different aspects in life, including financial and career aspects, making them more likely to receive a higher income. Issues like reverse causality and omitted variable bias and their importance to the Easterlin paradox will be discussed in section 4.1.

Readers who have paid some attention to figure 1 & 2 may have noted that both figures depict life satisfaction on their y-axis instead of 'happiness'. The difference between these two concepts and the different ways of capturing and measuring happiness will come back in section 4.3. Though, for now, it is important to note that using life satisfaction, as measured by the Cantril ladder *(or comparable methods)* is just one way of measuring happiness and need not be the best or most accurate way of capturing happiness, depending on what one precisely aims to measure. In these types of methods, people are being asked to reflect on their lives and think about how they feel like things are going. The answer that results from this method is also known as ***evaluative well-being****,* named after the evaluative process that people go through when being asked such questions.

A different way to measure happiness is to try to capture ***affect*** *(also: emotional well-being)*. This can be roughly seen as how good or bad we feel from moment to moment, or how *emotionally* happy we feel. Affect measures try to distinguish between the different emotions that we may feel on a given day. Does one feel content, cheerful, exuberant or rather stressed, sad, angry and miserable? This focus on our day to day emotions is the reason that these methods are also sometimes referred to as 'experienced happiness'.

In 2010, Kahneman and Deaton performed a study where they tried to test whether the positive association between income and life satisfaction *(evaluative well-being)* (fig. 1 & 2) could also be replicated when looking at the relationship between income and affect.

They do so by looking at the Gallup-Healthways Well-Being Index data. According to the authors, this survey provides *"an unusually detailed measurement of well-being."*. That is because the GHWBI survey tries to capture multiple aspects of happiness. Both that of evaluative well-being, and that of affect. The latter is measured by questions about the presence of different emotions - enjoyment, happiness, worry, stress - in the evaluation of **yesterday.**

Kahneman and Deaton construct three categories; 'positive affect' - the average of reports of emotions as happiness, enjoyment and frequent smiling and laughter -, 'blue affect' - the average of worry and sadness - and 'stress'. To enhance comparability and to check for a positive relationship, they inverse 'blue affect' and 'stress' into 'not blue' and 'stress free'. When they plot these three categories *(measures of affect)* and the Cantril ladder *(a measure of evaluative well-being)* on income on a log scale, the results look like this:

Figure 3: *"Self-reported well-being measures across the income distribution"* [c]

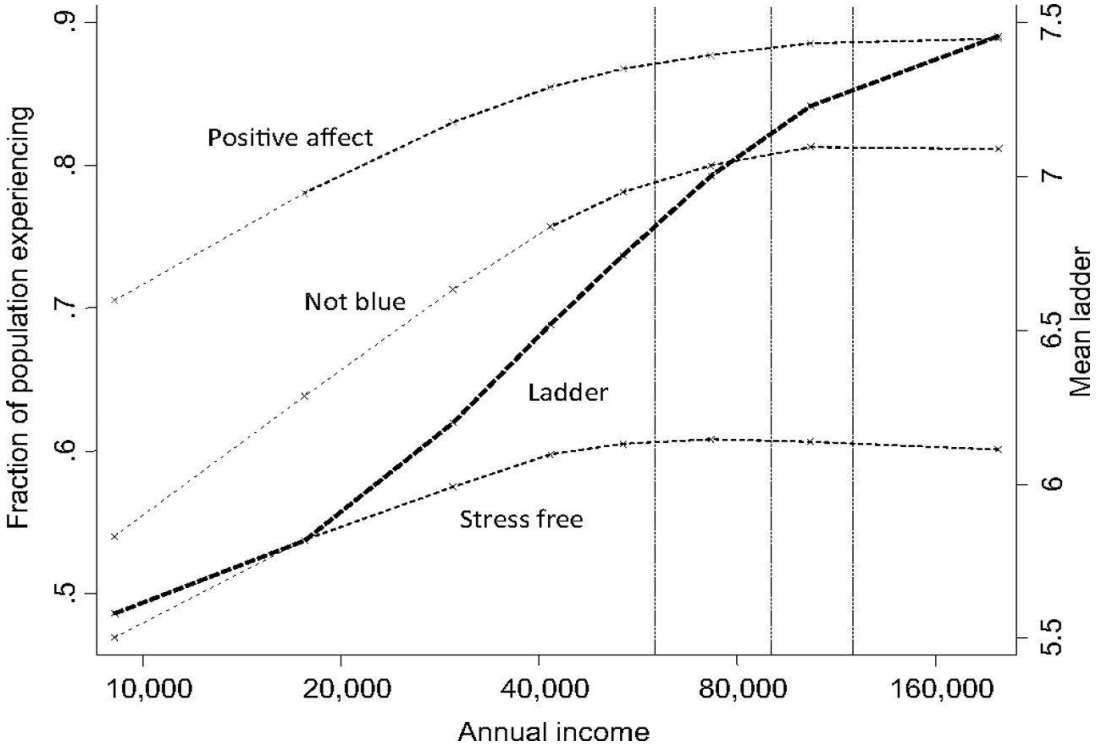

[c] **Source:** *(Kahneman and Deaton, 2010)* Three measures of affect *(positive affect, not blue, stress free)* and one measure of evaluative well-being *(ladder).* Positive affect is the average of the fractions of the population reporting happiness, smiling, and enjoyment. "Not blue" is 1 minus the average of the fractions of the population reporting worry and sadness. "Stress free" is the fraction of the population who did not report stress for the previous day. These three measures are marked on the left-hand scale. The ladder is the average reported number on a scale of 0–10, marked on the right-hand scale. These are visualised for eight different income groups in the USA and plotted relative to the median income in these groups. Data from 2008 and 2009.

For lower incomes, a doubling in a person's income seems to still be associated with a fixed increase in the measures for affect. Interestingly however, unlike evaluative well-being *(ladder)* the study seems to show that increases in affect *(positive affect, not blue, stress free)* plateau at around $75.000 *(in 2008 dollars)*. Therefore, the conclusion that the authors draw is:

> *"Emotional well-being [= affect] … rises with log income, but there is no further progress beyond an annual income of ~$75.000. … high income buys life satisfaction but not happiness …"* (p. 16489) [1]

Their conclusion makes intuitive sense, and is also the one that seems most representative of what most people have brought up in casual conversation when I discussed the subject of this thesis with them. Money goes only so far in making day to day life easier and more enjoyable, the argument seems to be. Some problems, like mental or physical illness, or trouble in your relationship are hard to solve with money and are therefore problems that even the richest of people still have to deal with. Maybe that around the threshold of $75.000 *(or ~$110.000, in 2024$),* all the aspects of life that can get better with money are 'solved' and one only remains with issues that hold no relationship whatsoever with income.

From this hypothesis follows that having a higher income than the threshold would only be of importance for happiness as represented by the evaluative well-being measure. An increase beyond the threshold would not matter for affect, because this extra income would not improve one's ability to do what matters most to their emotional well-being, such as enjoying hobbies and other forms of leisure and spending time with friends, family and loved ones.

However, this hypothesis that has grown out to be some sort of 'colloquial wisdom' has failed to replicate. In a study done in 2021 by Killingsworth, the author finds quite the contrary. Killingsworth starts by criticising the data used by Kahneman and Deaton, since the GHWBI asks people how they remember feeling during some period of the past *(the last day)*. *"This requires people to accurately remember how they felt across the various moments of the past and then accurately integrate those memories into a single estimate, an approach that is vulnerable to memory errors and biases in judgment."* (p. 1)

Furthermore, Killingsworth argues that Kahneman and Deaton's measure of emotional well-being *(affect)* actually partly represents a measure of evaluative well-being. He states: *"making a summary judgment of how one felt on average over some time period in the past [in K&D's case: a day ago] invokes a judgment process similar to evaluative well-being. … The true relationship between income and … [emotional] well-being could therefore be considerably stronger or considerably weaker than currently thought, and a plateau might exist at a different income level or not exist at all".*

To overcome this issue, Killingsworth used real-time smartphone data, where participants were asked multiple times a day 'how do you feel right now?'. This real-time character of the data makes it a closer representation of experienced *(vs. evaluative)* well-being than that of Kahneman and Deaton. Next to that, in Killingsworth's study, a considerable amount of participants with a high income were included. While the highest income group in the 2010 study has an median income of around $250K *(inflation adjusted)*, this new study also provides a data point for individuals with an median income of $625K. Overall, one can say that methodologically, this setup and data are superior to that of the one in 2010.

The results are the following:

[1] The authors imply here that affect *(or as they call it: emotional well-being)* is the best representation of happiness and evaluative well-being does not represent happiness. My reading of the literature is that Easterlin would probably disagree with them. More on this in section 4.3.

Figure 4: *"Mean reported levels of positive and negative feelings across the income distribution"* [d]

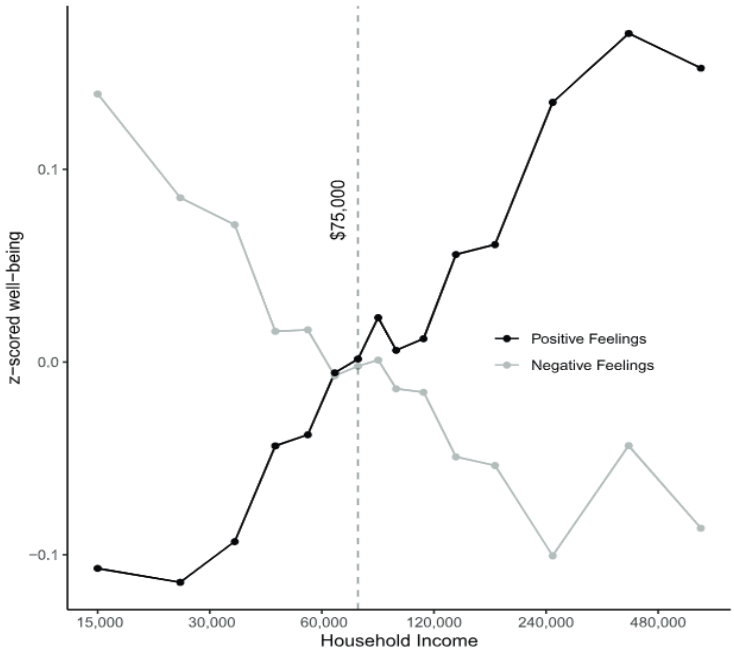

[d] **Source:** (Killingsworth, 2021) Mean levels of positive *feelings (Positive Feelings is the average of confident, good, inspired, interested, and proud)* and negative feelings *(Negative Feelings is the average of afraid, angry, bad, bored, sad, stressed and upset)* for each of the 15 income bands. The participants are from the USA and are aged 18-65.

This thus shows no evidence for a plateau, with both dimensions of affect - positive and negative feelings - showing a linear relationship with (log) income, even for high incomes. Killingsworth also replicates the finding of Kahneman and Deaton that there is no plateau for evaluative well-being *(life satisfaction)*, even for higher incomes, see:

Figure 5: *"Mean levels of affect and evaluative well-being across the income distribution"* [e]

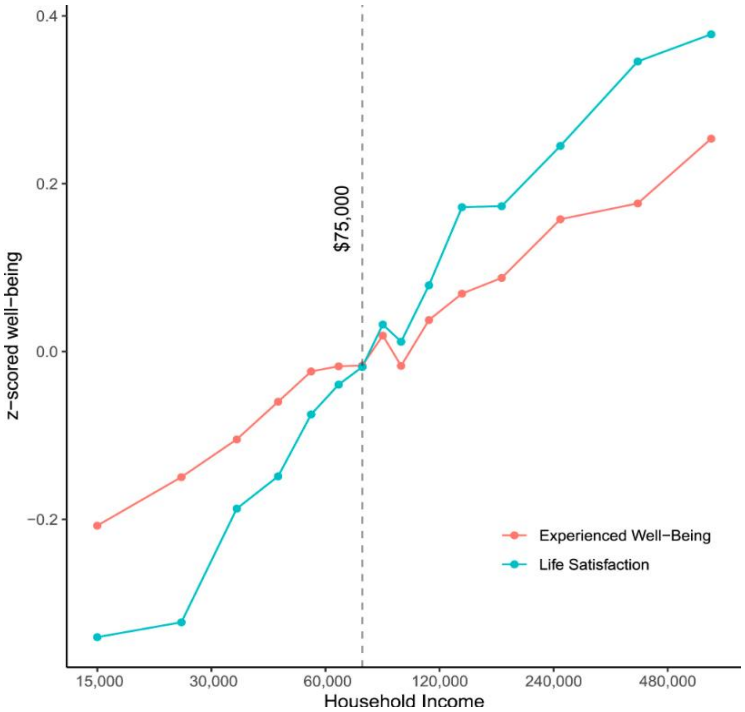

[e] **Source:** *(Killingsworth, 2021*) Mean levels of experienced well-being [affect] *(real-time feeling reports on a good–bad continuum)* and evaluative well-being *(overall life satisfaction)* for the 15 income bands. The participants are from the USA and are aged 18-65.

The result visualised in figure 5 also is compatible with the hypothesis that income matters more for evaluative well-being *(life satisfaction)* than for affect *(experienced well-being),* since the latter's slope is less steep. Furthermore, the underlying data indicates that differences in income below $80K were comparatively stronger in reducing negative feelings, while differences in income above $80K were comparatively stronger in increasing positive feelings (p. 3). This finding is interesting from a utilitarian perspective. If you regard emotional well-being *(affect)* as the sum of total positive feelings minus total negative feelings, it is interesting that the overall positive linear relationship between emotional well-being and (log)income exists not because of a constant relative importance of the two factors, but rather because the relative importance of the two changes as income rises.

The story does not end here, the unique part of this debate is that Killingsworth and Kahneman, after realizing that their results partly contradict each other, decided in 2023 to engage in an adversarial collaboration together with Mellers to search for a coherent interpretation of both studies. Such an adversarial collaboration, in which researchers with different views attempt to resolve their disagreement by doing joint research with the help of an impartial party, seems like an absolute victory for science, no matter the outcome.

The collaboration neatly explains the strengths and weaknesses of both papers, and does a re-analyzation of the data used in Killingsworth (2021). They find that the plateau does seem to exist, but only for the 15-20% of the population with the lowest emotional well-being. In other words, if you are unhappy to begin with, extra income over $100.000 *(that is more or less equal to the inflation corrected amount of $75.000 that Kahneman and Deaton found in 2010)* is not associated with an increase in emotional well-being. Though, if a person is not part of the 20% 'least happy' people, extra (log) income is associated with an increase in day-to-day positive emotions.

This insight is visualised in the following figure:

Figure 6: *"Emotional well-being across the income distribution and across the person-level happiness distribution"* [f]

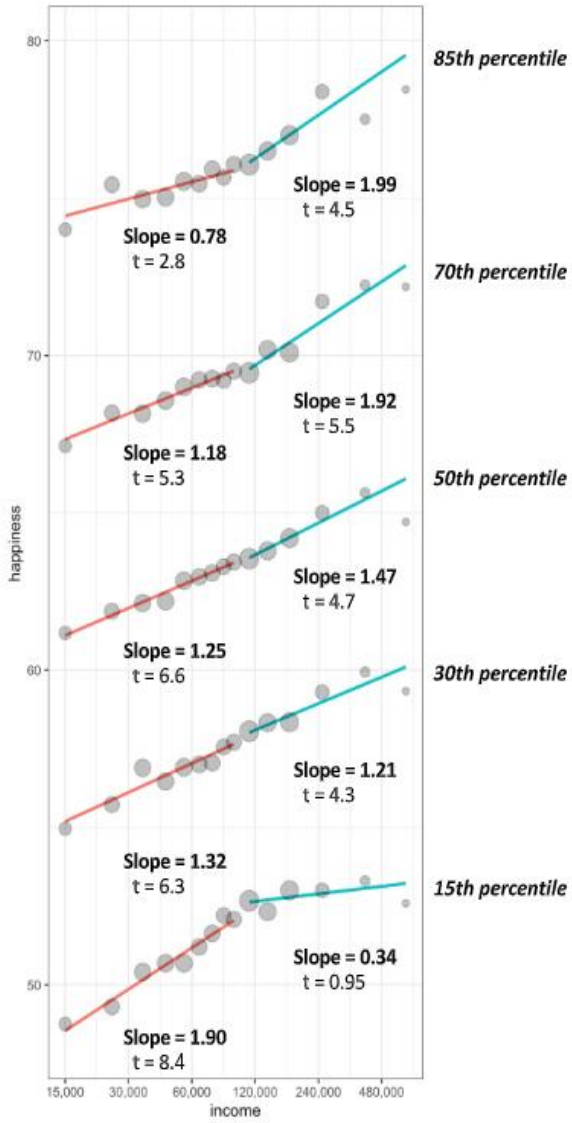

[f] **Source:** *(Killingsworth, Kahneman and Mellers, 2023)* Emotional well-being of the 15th, 30th, 50th, 70th, and 85th percentiles of the person-level happiness distribution, calculated within each income category. Slopes were calculated below and above 100k, using quantile regression

This figure might be a bit unintuitive at the first glance: it shows the result of the analysis where in each income group, the happiness score is divided into percentiles. Five of these percentiles have been plotted. The middle line *(50th percentile)* thus shows the median happiness rating for each income group.[2]
Their interpretation of this figure states that *"The suffering of the unhappy group diminishes as income increases up to $100K but very little beyond that. This income threshold may represent the point beyond which the miseries that remain are not alleviated by high income. Heartbreak, bereavement, and clinical depression may be examples of such miseries."* (p.3)

Another interesting finding is that for the most happy 30% in each income group, the happiness/(log)income relationship seems to accelerate beyond the threshold. This gives a surprising result, namely that the simple linear-log relationship *(of experienced well-being)* in figure 5 is the result of complementary nonlinearities. Though, to reinforce the claims following from figure 6, I deem a replication necessary, preferable in a different country or set of countries, only then can the aftertaste of HARK'ing fully be removed.[3]

Concluding this section, one can say that the association between (log) income and happiness within countries is both noticeable and positive. This result has been proven robust, and there seems no evidence of a certain general threshold after which this relationship weakens. Newer data of better quality is able to provide even more confidence in the relationship.

To check whether the relationship is not only statistically significant, but also economically meaningful, the interpretation of the data behind figure 5 is needed. For evaluative well-being *(life satisfaction)* a doubling in income is associated with a z-score increase of 0.2. This means that if one compares two groups of people within the same country at the same time, where one group has double the income of the other, life satisfaction is expected to differ by **0.2** standard deviations. As the standard deviation of life satisfaction *(measured on a scale from 1 to 10*) lies around 2, this means that with every doubling in income, it is expected to see a gain in life satisfaction of **0.4** *(on a scale from 1 to 10).* This seems economically meaningful, as it is a moderate effect. These approximations of the effect size of a doubling on income on average life satisfaction scores are also in line with the data behind figures 1, 2 & 3.

The effect on affect or emotional well-being is weaker. As can be seen in figure 5, a doubling in income is associated with a z-score increase of 0.1. This means that if one compares two groups of people within the same country at the same point in time, where one group has double the income of the other, a measure of affect is expected to differ by **0.1** standard deviations. The direct interpretation of a 0.1 standard deviation increase is harder to bring across, since most affect measures are not computable on a 1-10 scale. What the z-scoring does show is that the correlational effect size of income on affect *(according to some; a closer representation of happiness than evaluative well-being)* is about half of the correlational effect income has on life satisfaction.

---

[2] It seems a bit inconsistent to plot the 30th, 50th, 70th percentile but not the 90th and 10th, instead the 85th and 15th. Diving into the underlying data, it seems that these quantiles make the gist a little less perceivable. For example: the 10th percentile's lines have slopes of 1.72 & 0.52. Choosing these quantiles as top- and bottom is maybe remarkable, but is not misleading or an incorrect representation of the overall data.

[3] The finding that such a plateau only exists for people who belong to the least happy 20% in each income group is clearly a result that has only been spotted after the data collection, it was not present in the formulation of hypotheses or in any of the reviewed literature. That gives the suspicion of HARK'ing. It is not a priori clear that sorting people within each income band from least to most happy and then drawing lines between the different percentiles makes methodological sense.

To quote Killingsworth (2021, data file): *“Any amount of money in the measured range is enough to have a good chance at happiness, and no amount of money can guarantee it. There were plenty of happy people without much money in the study and some people who were miserable despite a big paycheck.”*

2.2: Between countries, macro level

The focus in this section will lie on statement 2: “Between countries at a given point in time, there is a positive association between income and happiness at the macro *(national)* level.” By investigating this statement, one answers the question whether richer countries are happier countries - or: whether happier countries are richer countries.

Statement 2 hasn’t always been part of the Easterlin paradox. If anything, if one would have presented statement 2 to Easterlin in 1974, he would have believed it to be false. This is because with the available data at the time, he came to a different conclusion. To understand this, it helps to see a plot of some of the data that was at hand by then.

Figure 7: *“Mean personal happiness rating and real GNP per capita”* [g]

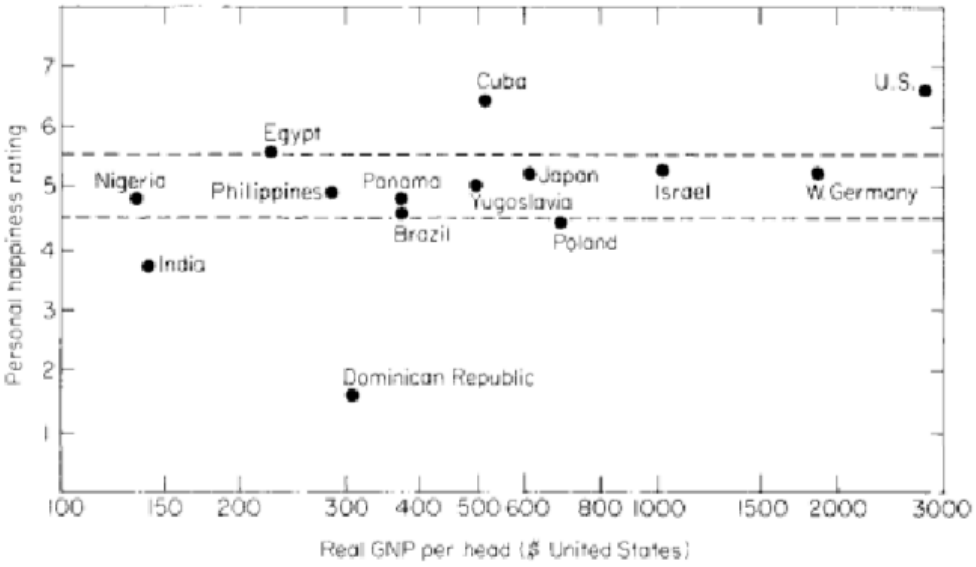

[g] **Source:** *(Easterlin, 1974)* Personal happiness rating *(on a scale from 1 to 10)* and real GNP per head, 14 countries, data from 1959 to 1963, depending on the country. Dashed lines do not represent a regression line of some sort, but the bandwidth in which most countries fall.

The main thing that Easterlin notes from this is that the happiness differences between rich and poor countries that one might expect on the basis of the within-country differences do not seem to be reflected by the actual cross-sectional data. A direct quote: *“if there is a positive association among countries between income and happiness, it is not very clear”*. (p. 108).

If there is one thing in Easterlin’s work from 1974 that has not aged well, it may be this. It goes too far to bluntly say that Easterlin was wrong, since his view seems a reasonable guess from the data.[4] Luckily, the relevant data has risen in quality, quantity and diversity over the years. Instead of 14 countries, we now have data of over 130 countries. We return again to Our World in Data for a visualisation:

[4] Though, with the same data at hand, Cantril (1965, p.194) comes to a different conclusion, namely that there is a positive association between income and happiness between countries at a given point in time.

Figure 8: *"Self-reported life satisfaction and GDP per capita"* [h]

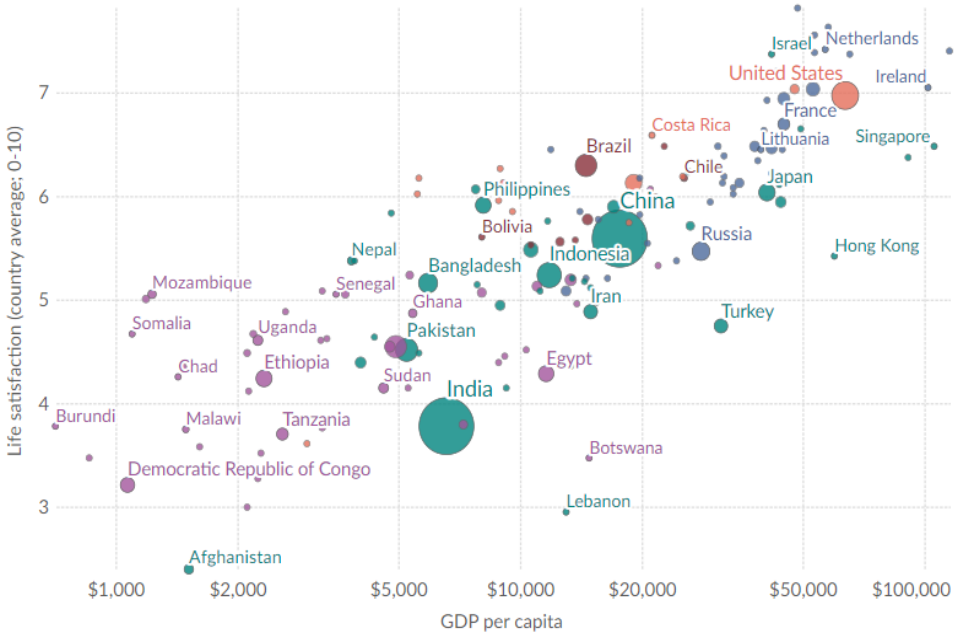

[h] **Source:** *(Our World in Data, 2022)* Self-reported life satisfaction and GDP per capita. Self-reported life satisfaction is measured on a scale ranging from 0-10. GDP per capita is expressed in international-$ at 2017 prices. Data from World Happiness Report (2012-2024) and World Bank (2023).

This yields a correlation coefficient *(unweighted for population size)* of 0.69.[5] This is a fairly high positive correlation. It was therefore inevitable for Easterlin to revise this statement in later work.

Since this, again, is a bivariate plot, it makes no scientific sense to interpret this as to depict a causal relationship, though it is tempting to do so. Easterlin himself has criticised such causal claims to a great extent, and has provided a different hypothesis *(than a causal link)* as to why the association between happiness and (log)income between countries at a given point in time is both noticeable and strong (Easterlin, 2013). More on the issue of causality in section 4.1.

Concluding this section, one can say that the association between happiness and income between countries is positive, and quite noticeably so.

2.3: Between countries over time, macro level

In this section, the focus will lie on statement 3: "Long-term trends in average happiness and average income are not significantly related." By looking into this statement, one answers the question how likely it is that general income growth causes happiness to increase.

This statement is the heart of the Easterlin paradox. It is the statement that has the potential to be the most important - of the three - for economic policy, because it answers what people want to know, namely how far general income growth - beyond income levels already achieved - is likely to increase average happiness. (Layard et al., 2009, p.1). Or, formulated by Inkeles in 1960 (p.18): *"Will raising incomes of all increase the happiness of all?"*

To test this, one needs to turn to time series data. Relevant data was in short supply around the time Easterlin wrote his paper, this is due to the fact that it requires time to gather a

[5] Pearsons´s correlation, own calculation

proper dataset, and since happiness economics was in its infancy during the time of writing, not much time in which there was interest in these kinds of variables had passed. Despite this, Easterlin managed to put together a time series for the United States from 1946 to 1970. Today, many more time series are available, a few of them are visualised in Origins of Happiness by Clark et al. (2017):

Figure 9: *"Self-reported life satisfaction/happiness and GDP per capita"* [j]

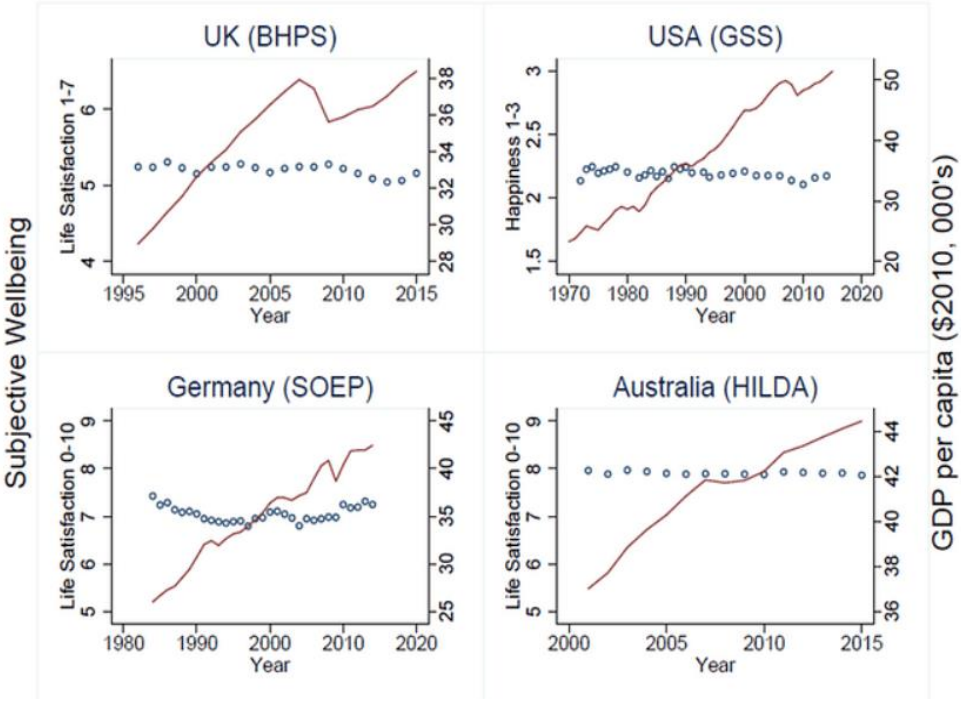

[j] **Source:** *(Clark et al., 2017)* The relationship between GDP per capita and measures of subjective well-being/happiness. Note that the measures on the Y-axes differ and that the time span covered on the X-axes does so too. Data from different sources.

These kinds of figures feel closely connected to the paradox, it hints to the hypothesis that as a nation gets richer, average evaluative happiness does not increase. Though this is not what the paradox is about per se[6], time series from around the globe that are consistent with this hypothesis are included in the majority of Easterlin's work on the matter. (Easterlin 1995, 2010, 2015, 2016; Easterlin et al. 2010).

Individual time series cannot serve as convincing evidence for the third statement. This is because there could be all sorts of different country-specific things happening that can influence both the trajectory of growth in happiness and income. In fact, some time series that have been put to the front as an example of this phenomenon *(that countries have had periods of rapid economic growth without accompanying growth in happiness)* have received sound criticism. To highlight an example:

> The case of Japan has been put to the front *(most recently in Easterlin and Angelescu, 2011)* as an example of a country where mean life satisfaction remained flat in a period where the economy grew remarkably fast. Stevenson and Wolfers (2008) show however, that the survey from which this time series collected their life satisfaction data changed the questioning over time. This makes it hard to confidently

[6] Easterlin sometimes makes it look this way; for example: "The Paradox states that … over time happiness does not trend upward as income continues to grow." (Easterlin 2016, p.4) or "When income goes up, happiness stays the same because of social comparison." (Easterlin & O'Connor 2020 p.14). The adding of 'ceteris paribus' would be a helpful instrument to prevent confusion here.

track changes in well-being over the whole period. In their analysis, the time series was cut up into 4 parts where the life satisfaction survey questioning was constant. Within three of the four periods, the correlation between income and happiness growth is positive.

The time of looking at individual countries has passed, and there is a need for higher quality analysis. The thing that is interesting to know is whether subjective well-being grows faster when GDP growth is higher. So what is needed is a collection of time series data from countries worldwide, with which a comparison between changes in happiness to GDP growth rates can be made. The World Values Survey and the European Values Survey provide this data.

Easterlin and O'Connor (2022a) use this data and their analysis yields the following result:

Figure 10: *"Growth rates of life satisfaction and GDP per capita"* [k]

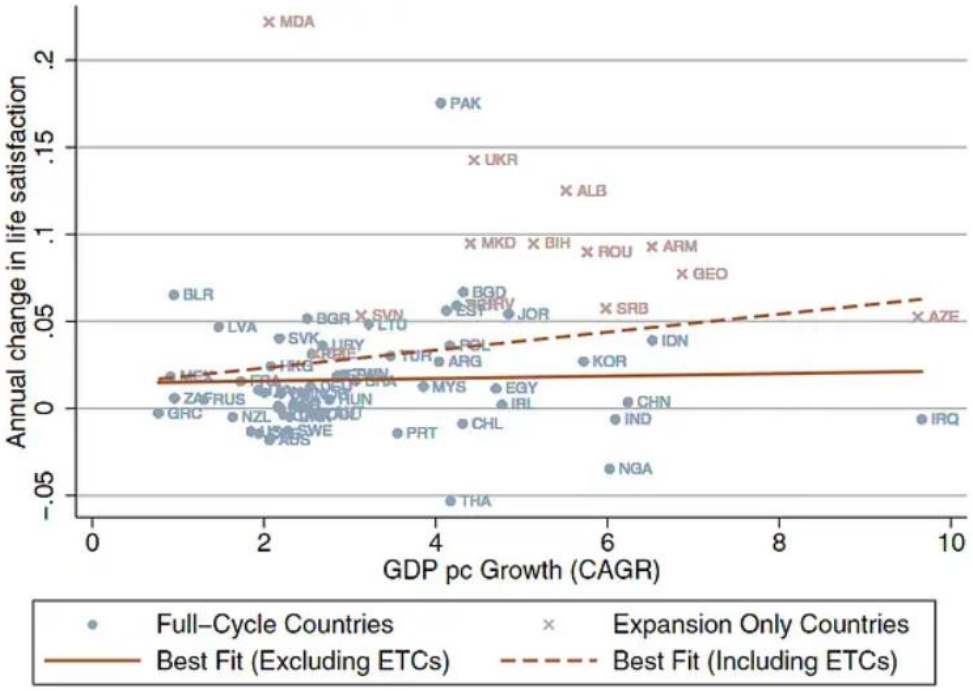

[k] **Source:** *(Easterlin and O'Connor, 2022a)* Growth rates of life satisfaction and GDP per capita with and without expansion only countries. Data from 1981-2019 (exact time span dependent on country) from the World Values Survey and European Values Survey.

This figure presents data of 67 countries, with time series ranging from 12 to 39 years, with an average of 27 years. The data ranges from 1981 to 2019. Easterlin and O'Connor emit 'Expansion Only Countries' in their main conclusions, which means that they are focussing on the solid brown line, *not* the dashed one.

The Expansion Only Countries *(ETCs)* are a set of European countries for which happiness data collection started to take place only after their transition from socialism to capitalism. When the USSR collapsed, these countries' economies saw a big downturn. In the period thereafter *(when data collection on happiness started)*, these countries have only seen economic expansion – that is why they are referred to as 'expansion only' countries - and big rises in happiness levels. Easterlin and O'Connor argue that to make a correct analysis, one should only look at countries who have also been through economic downturns, so through booms and busts, because only then can reliable *trend* lines of both happiness and income be estimated.

This choice might seem controversial at first, especially since when the ETCs are included, a positive relationship between GDP growth and happiness growth does seem to materialise - the dashed line -, but Easterlin and O'Connor are all but the first ones to make this distinction between transition and non-transition countries, see for example Bartolini and Sarracino (2014), Mikucka et al. (2017) and Kaiser and Vendrik (2019). The only part in the literature where this specific choice is criticised is in Kamilcelebi and Veenhoven (2024, p. 12) who state: *"Omitting countries that underwent fast economic growth is questionable, to say the least. We would rather see these cases as a natural experiment that revealed a strong long term impact of economic growth on happiness."* [7]

Easterlin and O'Connor state the following about the dashed line: *"The upward tilt that occurs when including the ETCs is not because they exhibit a positive relation between happiness and GDP. As mentioned, it is because they have both greater income growth and greater happiness growth than the countries with trend estimates. In fact, the ETCs' own income-happiness relation is not statistically different from the relation in the developed countries…"* (p. 10)

Following from this analysis, the authors conclude that there is no economically significant association between economic growth and happiness. Although the regression coefficients are positive *(the coefficient is 0.001 for solid line, and 0.005 for the dashed line)*, the implication of this positive result *(of 0.001)* would be that it would take 500 years to raise happiness *(ceteris paribus)* by one point if the long term GDP growth rate of a country would be increased by two percentage points, rendering this is in fact practically a nil result, according to Easterlin and O'Connor.

This result is not undisputed. There are authors who claim that the positive relationship between economic growth and happiness does in fact exist, therefore denying the existence of the paradox. The interesting thing is that these authors also have time series data at the basis of their analysis. Some quotes from 'paradox deniers':

> *"Both GDP and happiness have gone up in most nations, and average happiness has risen more in nations where the economy has grown the most…."* (Veenhoven and Vergunst 2014, p. 311)
>
> *"We found that rising income is significantly associated with… subjective well-being."* (Diener, Tay, and Oishi 2013, p. 267)
>
> *"… life satisfaction and other measures of subjective well-being rise with income, …. these findings cast doubt on the Easterlin paradox…"* (Sacks, Stevenson, and Wolfers 2012, p. 89)

What astounded me during researching these papers is that some fairly basic mistakes are made, including things that are included in chapter 1. To name three:

- Veenhoven and Vergunst (2014) are using current dollar GDP, which measures the output in terms of value instead of in terms of the quantity. Because of international differences in inflation rates, this renders their attempt to test the happiness-growth relationship less valuable.

- These authors also state that the paradox holds that happiness tends to be stable over time (p. 319). This might be an idea one gets following from looking at

[7] For more information on the choice of emitting the ETCs, one can read pages 6-8 of Easterlin and O'Connor (2020) or the other sources cited.

happiness trends like the ones in Figure 9. What the authors failed to realise is that the paradox does not claim anything about happiness trends per se. The paradox is about the association between trends in economic growth and trends in happiness.

- Sacks, Stevenson and Wolfers (2012, pp. 74-75) also fail to represent the paradox correctly by stating that it would rest on the belief that the relationship between happiness and income was stronger within countries than between countries. This is not the gist of the paradox.

These mistakes strengthen my confidence in that there is need for clarity on what the paradox does and does not entail, as I have tried to provide in chapter 1.

It would be scientific misconduct to not take the analysis of the 'deniers' into account because there are some basic mistakes to be found in some of the papers, we are all fallible humans after all. I will not go into detail in their analysis, but will provide the main reason why, according to Easterlin, these authors come to a different conclusion than the 'paradox believers'.

The most important difference has to do with the length of the time series that are used. This is because for the paradox, **trends** in happiness and income growth are of importance, which only can be measured over a longer period of time. The interesting thing namely is that in the short run, happiness and GDP tend to **do** go up and down together. Easterlin visualises this schematically in the following way:

Figure 11: *"Hypothetical co-movement of GDP and happiness"* [m]

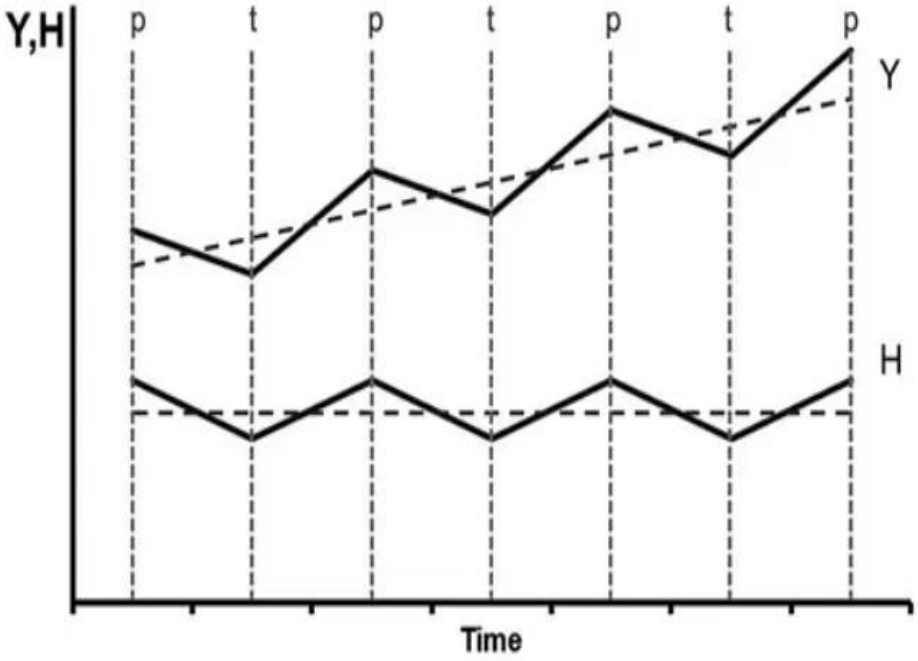

[m] **Source:** *(Easterlin, 2016, p.3)* Y is the hypothetical trend in per capita real GDP and H is the hypothetical trend in average happiness Short runs are depicted as the distance between p and t, and the long run is defined to exist out of a period where a country has gone through both economic busts and booms.

Here, one can see that although in the short run *(the distance between p and t)* income and happiness correlate, this effect disappears when looking at the longer run *(the dashed lines)*. If one would construct an analysis like in figure 10 with time spans that are too short to find trends *(dashed lines),* one is expected to find a positive relationship.

This is precisely what happens in most of the research from 'deniers'[8]. This is also a reason why Easterlin and others emit the abovementioned Expansion Only Countries from their analysis, because the data from these countries spans over too short of a time, schematically between p and t, and is therefore not suitable to estimate *trends* in happiness and income growth.

This difference in time span is crucial to understanding the different views. For example, if one would add time series that cover less than 10 years to the analysis of Easterlin and O'Connor (2020), figure 10 would yield a significant positive relationship. If one looks at 'fluctuations' instead of long term trends, the short-run positive relation is expected to dominate the results. (Easterlin and O'Connor 2020, p. 5) This is also consistent with the research of Bartolini and Sarracino (2014), who find that the nil relationship between trends in happiness and income disappear as the time series length is shortened. In the long run, a period that they define to be at least 15 years, they find no relationship between the trends in income and happiness.

In 2024, Kamilcelebi and Veenhoven have reacted to Easterlin and O'Connor. Together, these two papers[9] are the most recent, highest quality representation of the two different views - paradox believers vs. paradox deniers. In their response, they re-analysed the 2014 paper of Veenhoven and Vergunst, both by adding the most recent data and by changing the analysis somewhat. A major change is that in their most recent analysis, the authors only include countries for which time series of at least 20 years are available. This is thus a reaction to the critique of Easterlin about using time series that are too short, thereby looking at the fluctuation instead of the trend in income and happiness.

Furthermore, K&V use data from the Eurobarometer, the European Social Survey and the Lationo Barometro, next to the data sources E&O use, making their overall dataset more diverse than the one of E&O.

The only lacuna in K&V's analysis compared to E&O is that they do not express economic growth in purchasing power, but use growth of GDP as expressed in constant local currency. As noted above, since there are international differences in inflation rates, this hurts the quality of the analysis, though I cannot oversee to what extent the final results are dependent on this choice *(of not using PPP)*[10]. Their analysis yields a coefficient of **0.0035**, next to the 0.001 E&O found.

[8] There are more reasons why the two 'camps' of authors come to a different conclusion. See Easterlin (2016) for a complete *(though, not impartial)* analysis. For an analysis from 'the other side', read Kamilcelebi and Veenhoven (2024, pp.10-15)

[9] Easterlin and O'Connor (2022a) and Kamilcelebi and Veenhoven (2024)

[10] Kamilcelebi and Veenhoven report that they do not use PPP because this is not available for periods before 1990, while some of the time series they use start more than 20 years earlier. To show that this does not render their analysis useless, they could have included two analyses from 1990 onwards: one with GDP denoted in constant local currency and one using PPP data *(that is available for data from 1990 onwards).* If the two resulting coefficients do not differ substantially, this could provide support for the hypothesis that using constant local currency on the whole dataset does not hurt the confidence in the overall resulting coefficient of 0.0035. For now, as this analysis is not executed, it is difficult to judge to what extent the lack of PPP data influences their results.

Interestingly, K&V also convincingly argue that the coefficient of **0.001** that E&O report is dubious. This is because E&O originally found a coefficient of **0.005**. But this is reduced down to 0.001 in two different ways that the authors dispute, see Kamilcelebi and Veenhoven (2024, p.12-13).[11]

As a non-econometrist, I am not the one to make the final call on whether these choices made by Easterlin and O'Connor are methodologically sound or not, but aim to show that the choices made that bring the found coefficient down to **0.001** from **0.005** are somewhat dubious, and can be criticised. Though, these choices also find support in some parts of the literature.

Concluding this section, one can say that there is mixed evidence regarding statement 3: "Long-term trends in average happiness and average income are not significantly related." It is important to note that the *(counter)*evidence for this statement is of lower quality than that for statement 1 and 2. This has to do with data availability. Accepting or rejecting statement 3 thus need be done with lower confidence than the other two. For now the evidence points to a coefficient of the relationship between GDP growth and happiness growth between **0.001** and **0.005**.[12]

The spread in these findings have to do with the data used, the inclusion criteria and the econometric analysis used during and after regressing. Interpretation wise, this would mean that the average life satisfaction on a scale of 1 to 10 in a country with a long term GDP growth rate of 4 per cent is expected to gain somewhere between 0.2 and 1 point over 50 years ceteris paribus.

Next to that, I want to add two emphases that are important for the remainder of this thesis:

First, that if one accepts the statement, a better formulation is that the two are not related in an economically meaningful way[13], since even the solid brown line in figure 10 technically does show a statistically significant positive estimate *(of 0.001).*

Second, that the words 'long-term' carry a lot of the weight here. I will show in a later part of this thesis that some economic and interpretation problems pop up precisely because of the focus on the long-term.

---

[11] To summarise how they reduce the 0.005 to 0.001:

- The first way is by excluding the Expansion Only countries, which K&V criticise (K&V, 2024, p.12) *(Easterlin & O'Connor 2020, table 2, result (2) -> result (1))*
- The second way is by adding 'interaction terms in each group of countries to the regression in order to allow happiness in each group of countries to have its own constant and relation to growth' (E&O 2020, p.9) *(Easterlin & O'Connor 2020, table 2, result (2) -> result (1)).* The dubious thing here is that this statistical exercise is only reported in their 2020 working paper and not in their 2022 book chapter, where they only mention that a slope of 0.001 was found. Adding controls into the regression shifts the analysis from a bivariate one to one that comes closer to a causal analysis, but this is not the correct test for statement 3. E&O find support for their choice in Kaiser and Vendrik (2019), but Kamilcelebi and Veenhoven refute this support.

[12] This conclusion might feel a bit unsatisfactory. I agree with that sentiment and after reading the literature, I am convinced that it is possible to draw a stronger conclusion by combining the strong aspects of both analyses. This requires an adversarial collaboration between Easterlin and Veenhoven, like we have seen in section 2.1 between Killingsworth, Kahneman and Mellers. I come back to this point in chapter 6.

[13] Even this interpretation can be criticised, if the actual coefficient comes closer to 0.005 than 0.001, the implication of average happiness rising by 1 point on a scale from 1 to 10 *(assuming 4 per cent growth over 50 years)* seems rather substantial, even if actual growth rates over such a period of time are expected to be lower.

3: The proposed mechanisms

In this chapter, the focus lies on the proposed mechanisms behind the paradox. I do not aim to give a complete overview of the history and evidence for each of the mechanisms, but aim to critically reflect on the mechanisms, see if there is any sound empirical evidence to quantify their relative importance and aim to check whether there are any alternative mechanisms underexposed in the current literature.

3.1 Social comparison

This first mechanism is the one Easterlin used in 1974 to explain the found paradox and is also the only one that gets most attention in the current literature. It focusses on relative income instead of absolute income and the status concerns and social comparisons that go with that. Simply put, it states that the happiness return from income growth is a zero-sum game.

Verbally, the simplified argument goes somewhat like the following: every individual derives income related happiness - in more economic terms: utility - from two things: one's absolute level of income *(this gives them happiness via consumption)*, and their relative income in respect to others *(this gives them happiness via status)*. The 'others' here are the reference group, the people with whom she compares herself. An income increase for an individual is beneficial for her happiness, because this both improves her absolute level of income and *(ceteris paribus)* improves her relative income.

In reality however, with growing GDP, there is no such thing as 'ceteris paribus'. As GDP grows, the income of one's reference group grows as well. If you then assume status to be a zero-sum game - where the more status one person has, the less others have – and assume that the importance of the relative aspect of income dominates the absolute aspect, *(i.e. that the benefit on happiness from consumption is mainly derived from the relative, instead of absolute level)* this would then be able to explain the paradox. If one were to accept this mechanism as dominant and real, you could conclude - as Easterlin does - that raising the income of all increases the happiness of no one.

The theoretical background for this mechanism goes back a long way into history, Smith (1776) notes that considerations of appearance, honour, and dignity might influence one's choices regarding consumption, and Marx (1849) mentions the importance of social comparison to explain the social part of utility. One name that need not remain unmentioned in this respect is Veblen (1899), who came up with the term 'conspicuous consumption', which describes the part of consumption that is intentionally used to signal status. More recently, there have been constructed theoretical models to describe that the utility one gets from consuming goods that are conspicuous in nature - like clothing, or vehicles - depends partly on the consumption relative to that of reference group members (Perez-Truglia, 2013).

Empirically, there has built up a large literature showing the importance of relative income for happiness. To highlight some examples: What has been found is that people in the same income cohort are less happy when the incomes of people in their *(assumed)* reference group are higher. (Graham & Felton, 2006; Luttmer, 2005; Kingdon & Knight, 2007). Knight and Song (2006) find that relative income is the most important variable in explaining happiness differences between individuals in rural China. After reviewing some of the literature, Graham et al. (2010, p.259) conclude: *"The evidence suggests that concerns for relative income differences matter and can erode the positive effects of higher absolute income levels on happiness, … . … also at surprisingly low levels of income."* For a review of the literature on the effect of social comparison on happiness, one can read Clark et al. (2008) or more recently, Wu (2020).

Still, acknowledging the importance of relative income, this cannot be the full story. A part of income that one consumes is spent on goods and services that are not deemed conspicuous, for example on healthcare, groceries, housing, insurance and education. Even if income growth *(and accompanying reference group income growth)* would mean that the status benefit is zero for an individual *(as our social comparison theory predicts)*, the extra income also allows spending more on goods and services that are non-conspicuous in nature. There is thus a need for a theory of why this still does not lead to happiness growth in the long term: this is where the second mechanism comes into play.

3.2 Hedonic adaptation

The second mechanism is about hedonic adaptation. This mechanism was not explicitly mentioned by Easterlin in 1974, but has popped up in later literature trying to find explanations for the paradox. The focus here lies not on the relation one's income has with others, but with one's income in the past.

Evolution has caused people to be adaptive creatures. A substantial literature has been built up to show that people are able to internalise changes, and are able to transform a new, changed situation into a new ´baseline´ reference point after some time. For income specifically, this would mean that when an individual experiences income growth, this results in extra happiness in the short run, since this new situation of abundance is compared to the situation before the rise in income. Over time however, people adapt to this change by increasing their internal reference point and aspirations. People experience habituation and get used to the new standard of living, the theory goes.

An implication of that theory is that most people feel like they were less happy in the past and expect to be happier in the future. This is explained by the use of their current aspirations *(or: current reference point)* in judging their living conditions in the past, and the use of that same reference point in predicting their happiness in the future. People thereby forget to take into account their rising aspirations and adjustment of their reference point going into the future (Easterlin, 2001). An example of this phenomenon can be found in Frijters et al. (2002), there, the authors find that people in East Germany did not account for the adaptation to the initial euphoria right after the German reunification. Failing to consider this, they overestimated their future life satisfaction.

Theoretically, this theory finds ample support in evolutionary psychology. The importance of reference points has been stressed by the famous paper of Tversky and Kahneman (1991). Empirically, there is sound evidence of adaptation in many different aspects of life. An interesting example is Barazetta (2014), where a randomised controlled trial is performed to show that the initial well-being increase following a health improvement - increased mobility - disappears after one year, suggesting that people have adapted completely to this change after this time. Clarck and Georgellis (2013) show the existence of partial adaptation after marriage, divorce, widowhood and childbirth. There is some heterogeneity in extent and pace of adaptation, and what has been found is that people need longer to adapt to negative events *(disability, poverty, unemployment)* than positive ones (Lyubomirsky, 2011).

The same has been found in respect to income. Di Tella et al. (2010) found that the majority of the initial effect of an increase in income wears off in the following few years. The evidence from Burchardt (2005) is consistent with this idea, and also shows evidence for greater hedonic adaptation to income increases than to income decreases, as we have seen for other life events.

Take these two mechanisms - social comparison and hedonic adaptation - together, and one can see that they are two sides of the same coin: under social comparison, one takes others as a reference point, under hedonic adaptation, one's past self. The aim of this thesis is not

to study the proposed mechanisms in all their detail, rather, I aim to show that these mechanisms come with some problems that can explain why the debate surrounding the paradox still has reached anything but a satisfactory consensus.

A problem surrounding **social comparison** is that there is a strand of literature that follows Hirschman's (1973) interpretation in that the income of people in the reference group has an *information effect*. This means that the fact that there are richer individuals present means that there are opportunities for oneself in the future to become richer. In this way, the income of others in the reference group would correlate positively instead of negatively with happiness (Senik 2004, Clark et al., 2009a), this could also be the case because rich neighbours help in the provision of local public goods (as suggested in Clark et al., 2009b). Another issue is that some authors, like Kamilcelebi and Veenhoven (2024, p.15) state that people, when evaluating their lives, do not compare themselves to others to the extent Easterlin claims.

The specific issue with **hedonic adaptation** is that the literature finds little evidence of *complete* adaptation to income over time, which would be necessary for adaptation to explain the paradox. The evidence for *partial* adaptation is extremely strong[14], but as pointed out, partial adaptation is relevant for the relationship between happiness and income, and not orthogonally so for the paradox. Even if one decides to go with the sparse evidence of complete adaptation to income *(there are some examples)*, then you are still bound to run into problems. As pointed out by Clark (2016), individuals are happier with higher incomes in the short run, something Easterlin would agree with. Under full adaptation, for rich people to report higher well-being than poorer people *(as is shown in section 2.1)* the rich must have enjoyed more recent positive income gains than the poor. This can only be the case when there is a specific form of rising income inequality; something that has certainly not been the case for all countries during all time periods. In other words: under full adaptation and non-growing income inequality, one would expect the happiness-income cross section to be flat. As seen in figure 5, this is anything but the case. To quote Clark (2016, p.90): *"I believe that this somewhat undermines the case that can be made for adaptation to income as an explanation for the Easterlin Paradox."*

A more **general** problem I see is that the relative importance of the two mechanisms is hard to quantify. For a 'mature' explanation to be accepted, not only should the phenomena that are described be backed up by extensive literature *(which is definitely the case, both social comparison and hedonic adaptation are widely accepted phenomena),* there is also need for compelling evidence that these two are actually specifically relevant for the paradox, and not only for the relationship between happiness and income in general. One way to do so is to provide a reliable estimate to what extent these matter, but this has not happened yet.

What has been put to the front are some back-of-the-envelope calculations. This has been done by combining the finding in Knight and Song (2006), that relative income is at least twice as important than absolute income for happiness, with the finding in Van Praag and Frijters (1999) and more recently in Di Tella et al. (2010), that around 60% of the initial effect of an income increase on happiness fades away after a few years. Putting these findings together, Clark et al. (2008) suggest that 'only around 13 percent of the initial individual effect will survive in the long run at the aggregate level.'. The simplicity of such calculations signals the problem here: it is clear that the mechanisms matter to some extent for the relationship in which a rise in income translates to a change in happiness, but what has not be convincingly quantitatively proven is that the mechanisms are able to explain the paradox, and are not simply factors that dampen the effect of growing income on happiness.

[14] An interesting paper on lottery winnings is Lindqvist et al. (2020)

More **fundamentally**, a problem arises when proposing these mechanisms, namely that they implicitly assume that the Easterlin paradox is about **causation**, while it is one of **association**. The two mechanisms are 'answers' to the question why over the long run, income growth does not cause happiness growth, while in the short run, it does. This question is not the gist of the Easterlin paradox. Instead, it is all about why over the long run, income growth is not associated with happiness growth, while in the short run, it is. By naming the two proposed mechanisms as an explanation for the paradox, one misrepresents the paradox *(and yes, I blame Easterlin for this as well)*. This point is vital and can be seen as one of the most important contributions of this thesis to the debate. I find support for this view in Kelecelebic and Veenhoven (2024, p. 15) who write: *"…the effects of economic growth on happiness are likely to be multi-causal and contingent to situations and can as such not be captured in a single monocausal explanation."* After which they call the proposed mechanisms: *"a deficient account".*

There are possible mechanisms that respect that the paradox is about correlation, rather than causation. I propose two of these mechanisms in section 3.3. Furthermore, I have devoted a separate section to the issue of causality: section 4.1.

There is also one specific form of criticism on the paradox to which Easterlin replies with a theory of how the two proposed mechanisms interplay. I do not agree with the criticism, but neither with the refutation that Easterlin gives. In other words: I think the critique makes little sense, but not for the reasons that are provided in Easterlin and O'Connor (2020, pp.13-14).

The critique goes along the following lines: "If social comparison explains why happiness does not increase in periods of growth, then why does happiness decline in a recession, so with falling GDP? If social comparison were at play, people should not feel less happy because others are also in the same situation."

Easterlin goes on to counter this statement by claiming that *"in a recession income reference levels are no longer determined by social comparison but by comparison with one's pre-recession peak income. …. True, others are in the same boat, but that does not help meet one's monthly mortgage or rent payments. ….. When income goes up, happiness stays the same because of social comparison. But when income goes down, happiness decreases because the former peak income becomes the benchmark. ... In short, when incomes increase, social comparison determines the reference level, but when incomes decrease, comparison with one's … peak income, takes over … ."*
I deem this reply weak, it seems like an ad hoc 'solution' that is not backed up by proper literature[15] and gives the impression that the authors cherry pick the parts of the two mechanisms *(social comparison and comparison to one's past income)* that together are consistent with their theory.

Rather, the correct reply to this critique is that declining happiness during recessions is perfectly in line with the Easterlin paradox. The paradox namely does in fact allow that happiness and income go together in the short run.

More specifically: the fact that happiness growth and income growth do correlate in the short run is essential to the paradox. It is being used by Easterlin to counter the paradox deniers, as his main critique to them is that they fail to analyse the long-term growth rate and rather confuse it with the short-term fluctuation *(and thereby finding a positive correlation between happiness growth and income growth, as discussed in the text right under Figure 11, see*

[15] Easterlin and O'Connor refer to De Neve et al. (2018) in making this claim, but that paper does not show any evidence for the change from social comparison during economic growth to comparing to oneself in the past during economic contraction.

also Bartolini and Sarracino, 2014). Figure 11 shows this schematically, in every period, the short-run correlation is positive, but over the long run, this correlation disappears.

This difference between the short-run and long-run correlates is what I deem essential in explaining the paradox, not so much the contrast between the cross-sections and time series. It is now time to dive deeper into this issue: in the next section I aim to show that a related underexposed mechanism could explain *(part of)* the paradox.

3.3 Underexposed mechanisms

Happiness asymmetry

The starting point for this mechanism *(one could call it a hypothesis, since I will not provide new empirical evidence to provide support)* is the dynamics between income and happiness as depicted schematically in Figure 11. It shows that within each short-run period, income growth and happiness growth are positively correlated. This correlation has empirically been proven robust. The Easterlin paradox states that this association disappears in the long run.[16] This provides a puzzle, since the long run consists by definition out of different short-run periods, in which both are positively correlated. While the mechanisms in 3.1 and 3.2 both are trying to explain this by implicitly stating why economic growth *causes* short run growth in happiness, but not in the long run, the following mechanism does so by explaining only why the correlation disappears, and does not claim that economic growth does not cause happiness growth in the long run. In a way, it looks for a solution 'inside' of the paradox, instead of 'outside' the paradox *(in the world of causation; which the paradox is not about).*[17]

The reason the two dashed lines in Figure 11 do not correlate is because of happiness asymmetry in income changes: the fact that people react more heavily to losses in income than to gains. The theoretical background for this mechanism comes from prospect theory by Kahneman and Tversky (1979) and more specifically from loss aversion*: "the aggravation that one experiences in losing a sum of money appears to be greater than the pleasure associated with gaining the same amount".* This finding has been well established in the literature and Baumeister et al. (2005), after reviewing the evidence, conclude: 'bad is stronger than good'.

De Neve et al. (2018) show this specifically for the 'ingredients' of the paradox and find that:[18] *"measures of life satisfaction as well as positive and negative affect are more than twice as sensitive to economic downturns as compared to equivalent upswings."*. Beja's (2017) findings are consistent with this evidence. The fact that people react asymmetrical to economic changes is what can explain the discrepancy between the short run and long run evidence.

This is consistent with the idea that economic growth *causes* a short run happiness increase, and economic contractions *cause* people to be less happy in the short run, though to a different extent.[19] Focusing on the long term growth rate - as Easterlin does - is then not the

[16] In the conclusion of section 2.3 however it is shown that such a correlation does also exist in the long run, albeit a weak one.

[17] The difference between looking for a solution 'inside' the paradox and looking 'outside' is that when looking 'inside', one looks at the variables on the X and Y axes of Figure 10 and explains why these are the reason for the nil association, or proposes different variables to put on the axes to get to a positive association. Looking 'outside' requires one to find a solution for why income causes happiness growth in the short run, but not in the long run.

[18] They are all but the first to do so, in a similar study by Boyce et al. (2013), an equivalent result was found

[19] But is not a necessary condition for the mechanism to hold

most relevant, effects of greater average trend growth on happiness can be nullified by having more periods of economic decline. The 'solution' to the paradox would therefore be to take into account *how* economies grow in the long run, and not just *how fast*.

This hypothesis would 'predict' that correlation wise, the happiness of countries that have gone through long periods of economic growth without periods of contraction, would have grown faster than countries who have had both periods of economic growth and contraction, no matter the average long term growth. The remarkable thing is that Easterlin does provide us with a set of countries that have done just that: the Expansion Only countries in Figure 10. The fact that exactly these countries show higher happiness growth than 'Full Cycle' countries and at the same time show no *internal* correlation between income and happiness *(meaning that within the expansion only countries, income growth is not positively associated with happiness)* provides evidence for the theory that the paradox hinges on the presence *(and extent of)* economic contractions, not so much on the size of the long term income growth rate.

It is important to stay cautious: I do not provide any new evidence for this hypothesis here, and it is thus far from proven. Furthermore, it is not a 'complete' theory that has been worked out completely. But the fact that the hypothesis is consistent with all the stylized facts[20] without claiming anything about the long run causation of income on happiness proves to me that this mechanism is a promising possible explanation to the paradox and is at least worth examining further. I deem the fact that this theory has not been explicitly proposed or tested by the literature a lacuna that needs to be filled.

Income inequality

Another underexposed mechanism to try to explain the Easterlin paradox is the one of rising income inequality. Oishi and Kesebir (2015) show that the paradox can partly be explained away by the concurrence of economic growth and rising income inequality. Correcting for income inequality makes sense, since the Easterlin paradox specifically claims that 'raising the income of all raises the happiness of no one'. In the paradox, GDP *(sometimes GNP)* growth is used as a proxy for how much 'the income of all' has risen, but is an imperfect one.

It is the presence of rising income inequality that reveals this imperfection: if the economy grows while income inequality does so as well, GDP grows while that need not be the case for the income of the modal person. Showing then that the relationship between happiness and income growth is flat does not prove that raising the income of all does not cause an increase in happiness, for the fact that GDP growth does not signal that the income of all has risen *(to the same extent).* This mechanism has the possibility to again explain the paradox from the 'inside' rather than looking for causal mechanisms. This difference is a subtlety, but a rather important one. That is why section 4.2 comes back to this point, and explains the importance of income inequality in relation to GDP for the paradox more extensively.

Concluding chapter 3, I have shown the proposed mechanisms behind the paradox. The theoretical and empirical background for the two mechanisms are well founded in the literature, albeit with some critical remarks. The issues that I have brought up have little to do with the content of the mechanisms, but rather the relevance of them for the paradox. Since

[20] To name them:
- positive association between income growth and happiness growth in the short run
- little (or no) association in the long run
- prospect theory and the short run asymmetrical reaction to changes in the economy on happiness
- ETCs showing high happiness growth but no within group correlation

the quantification of the relative importance of the mechanisms has not been reliably worked out, I deem it likely that the proposed mechanisms matter for the *(causal)* relationship of income on happiness *(they work as inhibiting factors),* and not per se on the paradox specifically. The paradox is one of association, not of causation, but with the two proposed mechanisms in hand, one implicitly assumes so, thereby misrepresenting the paradox. Therefore, I have proposed two underexposed mechanisms that need not make the implicit assumption that income growth does not cause happiness growth in the long run, but that can explain the finding that income growth is not associated with happiness in the long run while it is in the short run.

4: Economic problems

In this chapter, some economic problems are being discussed that surround the paradox in some way. Understanding these problems can explain why the scientific debate has not resulted in a consensus yet and can also help to get a better picture of what needs to be solved or accounted for moving forward.

4.1 The problem of causality

Though the paradox is one of association and correlation, in interpreting and discussing the matter, the literature sometimes seems to act as if we are talking about causality. In this section I aim to give examples of this, show why it is tempting to do so and why confusing correlation with causation can lead to false claims and interpretations.

To sketch the problem it is needed to return to the start of this thesis, the **first statement:** "Within countries at a given point in time, there is a noticeable positive association between income and happiness at the micro *(individual)* level." The evidence that we have seen for this statement like Figure 4 and 5 can be misinterpreted by stating or implying that this thus shows that more income *causes* higher happiness on the micro level. As stated in the relevant section: from a correlation alone, this cannot be judged. There could be all sorts of endogeneity issues preventing this conclusion from being correct. Take for example *reverse causation*, it could very well be that happier people are more successful in different aspects of life, and are also on average able to generate a higher income.

Empirically, the existence of this reverse causality can for example be found with panel data. This has been tried and the literature finds that indeed, people with happier personality traits are more likely to obtain better jobs (Graham et al., 2004) and that unhappier people are more likely to become unemployed (Clark, 2003). More recently, Elsas (2021) has used an instrumental variable approach on German SOEP data to try to distinguish between the two causal paths and finds evidence of significant effects of happiness on income. Ye et al. (2023) use a unique dataset of Chinese twins to be able to estimate the causal effect of income on happiness and find a large effect. Together, these studies thus point to the fact that the causality between income and happiness goes both ways. Explaining the evidence for the relative strength of both causal paths lies outside the scope of this thesis, but these papers are examples to show that one must be careful with drawing causal conclusions from correlational evidence.

Interestingly, in 1974, Easterlin did take the possibility of reverse causality into account, but did not expect the effect to be relatively large. He therefore concluded: *"On the whole, therefore, I am inclined to interpret the [cross sectional] data as primarily showing a causal connection running from income to happiness."* (p. 104). This conclusion cannot be drawn from the type of evidence that was at hand. Hence, one could regard this conclusion as mere speculation rather than following directly from the available evidence. In any case, Easterlin did forget to consider another form of endogeneity, next to the possibility of reverse causality, namely omitted variable bias. The cross-sectional - correlational - data could also be

compatible with an underlying causal model where a certain variable causes both higher *(lower)* income and higher *(lower)* happiness. Barker (2005) shows that one of these possible omitted variables could be size at birth, linking it both to income and depression. There are many more omitted variables possible, like good health, quality of relationships with coworkers or marital stability. Personality traits could be an interesting candidate, for Mõttus et al. (2024) recently showed that the Big Five personality traits allowed predicting life satisfaction scores with accuracies up to r = 0.8-0.9 in independent subsamples. Again, I will not try to come close to finding the 'actual' underlying causal path here, but want to show that to draw causal conclusions, one needs proper econometric techniques for causal inference.[21] If these are not present, there is need to be extremely cautious in interpretation.

Remarkably enough, Easterlin seems very aware of this when it comes to **statement two:** "Between countries at a given point in time, there is a positive association between income and happiness at the macro *(national)* level.". The correlational evidence in Figure 8 is clear: the association is strong and positive. Were one to wrongly interpret this data as causal then the conclusion were to follow that economic growth *causes* higher happiness. This is the direct opposite of the prior disposition of Easterlin, namely that economic growth *does not* cause higher happiness. That is why Easterlin wrote a paper in 2013 to argue that the causal claim from this correlational evidence is incorrect. Instead, he proposes an omitted variable: historical experience. The argumentation goes as follows: *"Although [the] cross section relationship [of economic growth on happiness] is often taken to indicate causation, it may, … merely reflect historical experience, that is, similar leader-follower country patterns for variables that are causally unrelated. … Suppose, … that, the same group of countries have a head start, and the follower countries in the various parts of the world fall in line in a similar geographic order. The result will be statistically significant international cross section relationships among the various phenomena, despite their being causally independent."* (p.1) Easterlin thus hypothesises that countries that had earlier 'access' to causes that boost GDP, also had earlier 'access' to *(different)* causes that boost happiness. The cross-country correlation therefore would not signal a causal connection between the two, but would be explained by this omitted variable.

The real trouble starts with **statement three. "**Long-term trends in average happiness and average income are not significantly related." The evidence that is at hand here is time series data. Time series data, unlike cross-sectional data, *could* be used for causal inference, if one were to construct a correct underlying causal model. This is not done in Figure 10, or in any of the evidence for or against the paradox. What it depicts is the correlation between two variables, albeit variables that are the average growth rate over a longer period of time. No matter if there is a paradox or not *(section 2.3),* no matter if the solid line in Figure 10 is upward sloping or flat, from this evidence alone, concluding that higher income growth does or does not cause higher happiness is premature and in clearer words: plainly false. The words of Easterlin and O'Connor (2020 p.9) *"The Paradox holds – economic growth does not make people happier"* should be categorised in this category: premature and false. The only implication from their evidence is that *(assuming their analysis to be correct)* economic growth does not go with more happiness. Finding out whether that is because growth does not cause happiness or because of some other *(endogenetic)* reason is extremely relevant, but can only be done with an approach suitable for causal inference. Clark (2014) adds to this: *"There may well be other differences between countries that are correlated with satisfaction. This is a standard omitted variable argument: if something is positively correlated with income but negatively correlated with subjective well-being, then the bivariate*

[21] Wu (2020, p. 4) concludes, after reviewing the micro evidence: *"In general, the effects found in panel studies are smaller than those found in cross-sectional studies. In addition, the effects found in causal analyses, using either exogenous variations in income or instrumental variables, are larger than those found in associational analyses."*

correlation between income and well-being will be biased towards zero. This is the case for hours of work or pollution, to give two examples."

The question then is why an approach suitable for causal inference has not dominated recently. The answer that I deem most probable is that designing such a method comes with extreme difficulties, econometrically speaking. With statement one, exogenous variation in income and IV-approaches are used to approximate the causal effect. This is way harder at the macro level, especially in the long run. Clark et al. 2008 (p. 138) agree with this and state: *"… natural experiments producing exogenous variation in income are only rarely observed, making the issue of establishing the causal effect of income on happiness a major challenge."*

Another way of causal inference could be to build a model as used in hedonic pricing models, where one uses many more macro time-series variables [than economic growth and happiness growth] and controls for the 'right' ones. The impossibility here is that deciding which variables are confounders and should therefore be controlled for to rule out omitted variable bias and which variables are colliders and mediators, for which controlling renders the estimated effect biased, is next to impossible. For more on the matter, one can read the great paper from Wysocki et al. (2022) titled 'Statistical Control Requires Causal Justification'. The crux is that such justification is very hard for the complex relationship the paradox studies. In IV-approaches, a similar problem arises: it is rather difficult to verbally defend that the *exclusion restriction* holds, for the case can be easily made that whatever instrument one chooses, this could potentially affect happiness.

These difficulties in constructing a method better suited for causal inference, makes me draw an uncomfortable conclusion: the technique that Easterlin uses, however limited, might be the best that is presently available. Since detailed econometric challenges are far from my specialism, I cannot judge whether sometime in the future a model *could* be constructed that is able to distil the causal effect of economic growth on happiness on the macro level in the long run. For now, the method used by Easterlin as in Figure 10 might just be the best we can get. The fact that *'paradox deniers'* try to demonstrate their position by stating that the actual long term correlation is positive, instead of prompting a method that claims that economic growth causes happiness in the long run, provides evidence for this supposition.

Because of the scientific and human urge to say something meaningful about the long term causal relationship between economic growth and happiness, it is therefore understandable that it is tempting to make *causal* claims from the *correlational* evidence. Though tempting, it is premature and false. No matter how unsatisfactory it may be, the right and prudent conclusion is that the Easterlin paradox provides little direct insight into the causal effects in this relevant relationship.

4.2 Measuring income

This section asks the question what variable should be on the X-axis of Figure 10. It should be a measure of income on the macro level[22], as that is what the paradox is about. But that the compound annual growth rate of per capita GDP is the right one, is what could be argued about.

Easterlin and O'Connor (2020, p.3) write: *"In what follows "income" always means real income, what money will buy. At the national level, it is typically approximated by per capita Gross Domestic Product … because of the ready worldwide availability of reasonably*

[22] The discussion on what should be the correct measure for income on the *micro* level, either in time series or in cross-sectional analysis like in Figure 3, 4, 5 & 6 falls outside the scope of this thesis.

comparable GDP statistics. Per capita GDP here will always mean real GDP per capita, a country's average quantity of goods and services per person." Real GDP per capita is not a controversial choice, if anything, it may be one of the most used variables in economics.

The possible problem arises in the combination with the variable on the Y-axis, the growth rate of life satisfaction. With life satisfaction, the answer of every person *(in reality, every participant in the study)* on an evaluative well-being question is given an *equal weight* to determine the average. With GDP however, the 'answer' to the income question of every individual by definition does not have an equal weight, for people's incomes are vastly different. This yields the situation that *even if* people would *only* use their income level to determine their happiness level, a rise in GDP would not have to mean that happiness rises as well. To give some intuition: assume a country where year on year, the incomes of the poorest 50 percent of people has gone down by 5 per cent, the incomes of the richest 10 percent of people has doubled, and the sum of these two results in a rise in per capita GDP. Now, even if people would *only* use the change in income to determine their happiness level, it could very well be that the average happiness level does not rise, while GDP does. To then draw the conclusion that *'raising the income of all, raises the happiness of no one'* is wrong, for the incomes of all have not been raised. This is the crux: a rise in GDP need not mean that the incomes of most people have risen as well.

To correct for this, there are two possible options. The first one being correcting for changes in income inequality, like proposed in section 3.3. Oishi and Kesebir (2015) do this and find that this 'solves' a big part of the paradox, Mikucka et al. (2017) find a similar result. Another way that I have not come across in the literature would be to take a different measure altogether, for example the median equivalised disposable household income, deflated using a consumer price index, though in theory, more median income measures are possible. The gist is here that one takes a variable that better presents the income change that represents an 'equally weighted' change in income, which mirrors the equal weighting in the happiness question.[23]

Not correcting for changes in income inequality need not be 'wrong' per se, but it limits what can be said about what the paradox implies. In the current form, one could say that greater economic growth does not seem to be correlated with greater happiness growth.[24] But *not* that raising the income of all does not raise the average happiness, since then you need to take income inequality into account. If you do not correct for this, like Easterlin does not, then a mechanism behind the paradox could be that income inequality needs to be corrected for. Which is a mechanism that is not explicitly proposed by Easterlin, but that I do propose in section 3.3. Instead, the proposed mechanisms like discussed in section 3.1 and 3.2 seem to be a 'solution' to why people over time do not get happier of the increased consumption of goods and services. To have a more consequent story for these mechanisms, it then seems more logical to have a measure of average *(or median)* consumption on the Y-axis (see for example Meyer and Sullivan, 2003). One possibility of why this has not happened yet is that comparable data availability might be the limiting factor at play.

[23] It may be that correcting for changes in income inequality is superior to using a 'median' variable. For example in the following case: if the top 5 percent of incomes gain a lot y/y while the bottom 20 percent of incomes lose a lot y/y while the middle incomes stay more or less equal, this would show up in income inequality measures, while it may not show up *(to that extent)* in the change of the median income.

[24] In all fairness, this prudent result is a conclusion that Easterlin mentions regularly, but also sometimes - wrongly - seems to go beyond.

Concluding, GDP/capita is not a faulty variable per se, but problems arise when combining it with happiness data and with the mechanisms and implications that are currently proposed in Easterlin's work.

4.3 Measuring happiness

Thus far in this thesis, the term 'happiness' has been used a lot, sometimes as depicted by a 'synonym' like subjective well-being or life satisfaction. A complete overview of the definition of these terms will not be given, but some attention on what these different concepts entail, and where they overlap and differ is needed to be able to properly discuss the matter.

Easterlin (2003) states in the first sentence of this paper: *"I take the terms happiness, utility, well-being, life satisfaction and welfare to be interchangeable".* Though it makes life a lot easier to do so,[25] it is short-sighted and too simplistic, especially for scientific literature. Radcliff (2013, p.78) mentions that Easterlin is far but the only one to use these terms interchangeably, and that is rather common to do so.

To give a very concise overview of the literature on this aspect; it seems to be that subjective well-being can be split up into three parts: life satisfaction, affect and eudaimonia. In section 2.1, I have referred to life satisfaction as evaluative well-being *(it is also named 'cognitive well-being' by some authors)*, and sometimes to affect as emotional- or experienced well-being. To not muddy the waters further and for the sake of consistency, these underlined terms will be used in what is to come. These three parts all entail something different and are measured using different techniques.

Evaluative well-being captures a reflective assessment on a person's life or some aspect of it. The most used measures of evaluative well-being are questions on life satisfaction and the Cantril ladder. These two measures are central in Easterlin's work and seem to be the closest representation of his version of 'happiness'. Though, as demonstrated in section 2.1, it is not the only dimension of subjective well-being.

Affect is used to describe a person's feelings or emotional states. The aim here is to measure how people experience life rather than how they remember it (Kahneman and Krueger, 2006). Hedonic well-being is further split up between two dimensions: positive affect and negative affect (Diener et al., 1999), where positive affect captures positive states or emotions and negative affect captures negative states or emotions. These two different dimensions can be seen in Figure 4.

Eudaimonia is a concept that differs from the other two by that it does not focus on a person's reflective evaluation and emotional states but focusses on the good psychological functioning and the realisation of one's potential (OECD, 2013, p.32). It is more related to the concepts of meaning and purpose in life, and is sometimes referred to as flourishing[26]. An example of this would be that of having children. It has been reported that the correlation between having a child and evaluative well-being is negligible (Dolan et al., 2008), and that caring for a child is *(how surprising that may sound)* not per se correlated with high positive affect (Kahneman et al., 2004). Instead, children mainly 'show up' in subjective well-being through *eudaimonia*, since parents do report higher levels of meaning and purpose in life (Thompson and Marks, 2008).

[25] I have implicitly done so as well in the majority of this thesis

[26] Though this term is also used to depict a broader concept that resembles 'the good life'.

Now that the conceptual differences between these three concepts are clear, it is needed to have a look at the empirics. After all, if these concepts entail something different, but their measures correlate to a very high degree, Easterlin's equivocation might be defendable.

This seems not to be the case, as correlations between measures of affect and evaluative well-being are found to be in the range of 0.5 and 0.6 *(for example in Diener, 2009)*, and correlations between evaluative well-being and eudaimonia between 0.2 and 0.3, which is also the range in which affect and eudaimonia seem to be correlated. This indicates that the different dimensions of subjective well-being do indeed entail different phenomena.

Where Easterlin goes wrong is in the following argument*: "All of these questions, in which people are asked to evaluate their lives as a whole, yield quite similar results ... Hence, these measures are typically used interchangeably as indicators of happiness."* (2020, p. 3). The fact that measures of evaluative well-being are highly correlated *(close to 0.9)* need not mean that these represent happiness well or can be used as a synonym for it.

If anything, in the literature, 'happiness' seems to be more closely connected to affect *(or emotional well-being*). For example, in Ng (2022) and Howell & Howell (2008) the word 'happiness' is being used to describe the concept of what I call affect in this thesis. Furthermore, I think that this resonates most with the day-to-day use of the word happiness, as it seems to be used mostly as a description of an emotional state. Some authors see happiness as the combination of both evaluative well-being and affect, but happiness as a synonym for evaluative well-being on its own *(like Easterlin seems to imply)* is a seldom defended position.

Concluding, the Easterlin paradox is thus more about evaluative well-being than the most used connotation of 'happiness'. Whether this matters for the content of the paradox is the next question. What has been found is that affect measures typically have less of a relationship with income than evaluative well-being measures (Graham et al., 2010; Howell & Howell, 2008). This is consistent with Figure 5, where the slope for life satisfaction is steeper than that of experienced *(emotional)* well-being. This can be interpreted as to be favourable for the paradox. The paradox (tries to) show(s) the nil relationship between trends in income and evaluative well-being. Since affect measures typically hold an weaker relationship with income, the prior probability that this association will turn up positive *(the actual relationship between 'happiness' and income, if you will)* is not that high, though there is need of empirical validation of this claim (Deaton 2008; Harter & Arora 2008).

The use of terms and concepts by Easterlin may then not be detrimental for the paradox, I deem it sloppy at the very least. Decades ago, this position might have been defendable given the rudimentary knowledge at the time, but with the subfield of happiness economics reaching a more mature state, the use of terms should be reflective of that. That Easterlin in as recent as 2016 states: *'"Happiness" is used here interchangeably with subjective well-being … as a proxy for all evaluative measures of self-reported feelings of well-being'* is in my view indefensible. Words matter, and to respect the seriousness of the field, choosing them should happen cautiously and consciously.

4.4 Evaluative well-being: scale issues and cardinality

This section zooms in on some of the problems surrounding the measures for 'happiness' in the Easterlin paradox. Following from the last section, these can be more accurately depicted as measures of evaluative well-being.

In the data that Easterlin uses, two different measures of evaluative well-being are used. The World Values Survey asks "All things considered, how satisfied are you with your life as a whole these days?" to which subjects can respond with an integer between 1 and 10. The

Gallup World Poll uses the Cantril Self-Anchoring Striving Scale (Cantril, 1965). In this approach, people are asked to rate their present lives on an imaginary ladder with steps ranging from 0 to 10, where 0 is the worst possible life, and 10, the best possible life. The life satisfaction question is what is used in Figures 1, 2, 8 & 9, and the Cantril ladder is used in Figure 3. As shown in section 4.3, these slightly different measures correlate to a very high degree, making it defendable to regard them as more or less the same.

The distinct quality of these measures is that they are nominally value-neutral in nature, since respondents themselves can decide what the best possible life or maximum satisfaction with life for them entails, and what is needed to reach that. From that, it should follow that when asking participants what their preferred level is that they would like to achieve, they respond with 10: exactly because people can themselves determine what well-being is. A recent analysis by Nilsson et al. (2024) studies these preferred levels and find the following:

Figure 12: *"Distribution of preferred scores on evaluative well-being scales"* [n]

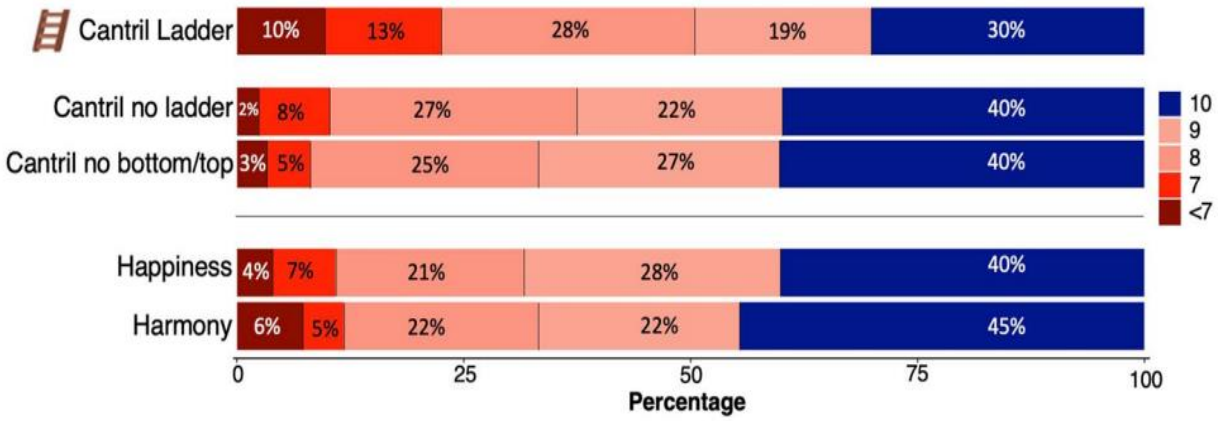

[n] **Source:** *(Nilsson et al., 2024)* The figure shows the distribution of the preferred level: all scores below 7 were grouped into one colour. It shows the preferred levels for different versions of the Cantril Ladder and for a more general happiness and harmony question *("where 10 represents the happiest/most harmonious life for you and 0 represents the unhappiest/least harmonious life for you").*

It is surprising to see that in all scale variations the majority of people do not indicate that they would prefer to score a 10. Averaging the preferred levels results in scores between 8.4 and 8.9. This is contradictory to the belief that people actually see the 10 as 'the best possible/happiest/most harmonious life' for them. The latter is assumed in most analyses that use these measures. This may indicate a misalignment between the scales' theoretical construct and how the scales are being interpreted. The consequences of this misalignment are thus far understudied in the literature.

A second problem surrounding these types of scales is cardinality**.** Cardinality is needed to be able to claim that two people who report 6/10 on the Cantril Ladder *(or on another evaluative measure)* are actually both as satisfied. (Kristoffersen, 2010; Lehberger and Musshoff, 2015). One dimension of cardinality is that the scales are **linear**, thus that there is a linear relationship between the true and reported subjective state. (Plant, 2020, p.8). This means that going from, say 7/10 to 8/10 on the life satisfaction scale is the same change in

satisfaction as going from, say 2/10 to 3/10. If such linearity does not hold, then aggregating scores into an average score, as happens in the Easterlin paradox, makes no sense.[27]

There are authors that claim that subjective well-being scales are not linear. For example, Ng (2008) proposes an arc-tangential relationship, where the distance in happiness at the extremes of the scale increase, i.e. that it is 'easier' to go from a 3/10 to a 4/10 than it is to go from a 1/10 to a 2/10. Plant (2020, pp.16-18) reviews the evidence on linearity and concludes that it is unlikely that individuals use anything other than a linear reporting function, though leaves this option open for future research.

Another issue on the cardinality of evaluative well-being scales appeals to the intertemporal aspect. The problem is that it is difficult to, over time, distinguish between **adaptation** and **scale shifts**. Ng (2008) mentions the importance of differentiating between the two. Adaptation as described in section 3.2 describes that the subjective impact of an event reduces over time. Scale shifts describe that people may change what their scale's endpoints *(the 0 and 10)* represent.

Take for example a woman that indicates to be 7/10 satisfied with her life *(or, scores 7/10 on the Cantril Ladder)*. She then wins the lottery and receives a sum of money worth about five years of income. After a few years, she again reports to be 7/10 satisfied. The first possibility is that she has fully adapted to the monetary gain and is genuinely as satisfied. The second option is that she is actually more satisfied than before winning, but has rescaled. Meaning that she has expanded her satisfaction scale, inflating the level of satisfaction a 10/10 represents.

The two options are not mutually exclusive, it could be that both are at play. Rescaling could potentially be a big problem for the Easterlin paradox, since it leaves the option open that over time, people *are* actually getting happier/more satisfied the faster the economy grows, but at the same time change their reporting style so that it looks like they are not getting happier (Plant 2022). The question then is whether the reported adaptation that we see after all kinds of life events are a result from 'actual' adaptation or from scale shifts.

The issue of rescaling is difficult to settle definitively since the actual underlying states *(how happy/satisfied people really are)* are impossible to directly observe. More surprisingly then, there is little discussion or attempt in the literature to help this debate move forward. Most of the studies on adaptation make the assumption that rescaling over the lifetime simply does not happen, see Luhmann et al. (2012) for a review. Oswald and Powdthavee (2008, p.16) only state that *"There is probably no way to reject such concerns definitively,"*. Plant (2020) is the first to review the limited literature on rescaling, and concludes that the best guess is that people do not rescale over their lifetime.

[27] Imagine that linearity *does not* hold, and that going from a 9/10 to a 10/10 is associated with a *much bigger* change in satisfaction than going from a 5/10 to a 6/10. Now assume that in t=1, there are two individuals. Person 1 reports a 6/10 and person 2 reports a 9/10 on a life satisfaction scale. This averages to 7.5/10. In t=2, person 1 reports a 5/10 and person 2 reports a 10/10, averaging again to 7.5/10. But since the positive difference in satisfaction between 9/10 and 10/10 is much bigger than the negative change between 6/10 and 5/10, average *(underlying)* satisfaction has actually risen. Taking the average of the reported scale values is thus a misleading way of representing the change in average underlying satisfaction from t=1 to t=2. To use an aggregate happiness level over time *(as Easterlin does)*, linearity thus must hold.

Even if one agrees with the analysis of Plant that there is little evidence of rescaling over the lifetime, there is one remaining issue which goes almost unmentioned in the literature[28]: **intergenerational rescaling**. This is the hypothesis that the 'underlying/actual' happiness - Ng (2022) calls such a hypothetical intertemporal and interpersonal comparable direct measure 'happiness units' - for one generation that is associated with, say, a 7/10 on a happiness scale, is different than that the 'underlying/actual' happiness needed to reach a 7/10 for another generation.[29]

As an illustrative thought experiment, one could go back to the Middle Ages. Now imagine a person that reports their happiness as 8/10, and think about the amount of 'happiness units' needed to reach such a level, and the amount of 'happiness units' the average person back then associated with a 10/10. Now travel back to the present, envision a person that rates their happiness as 8/10 today and compare the happiness of these two people. The question then is whether the amount of 'happiness units' that underlies both of these scores is the same or not. I find it very tempting to go for the latter, and assume that the amount of happiness experienced by a '8/10 person' today would result in a higher score *(than 8/10)* when experienced by a person from somewhere in the past. Plant (2020, p. 20) disagrees with my inclination.

Assuming intergenerational rescaling to be an actual phenomenon, which is my disposition, it could be relevant for the paradox if and only if the rescaling between subsequent generations is substantial enough to be picked up within the timeframe that is used in testing the paradox, mostly around 20 years.[30] With the lack of literature on this issue, and specifically the lack of any empirical testing of the existence of the phenomenon and extent to which it happens, it cannot be judged whether this is the case or not. If intergenerational rescaling does happen, it is more relevant to claims about long-term changes in happiness than claims about the short-run, since the extent to which it matters for aggregate happiness scores is dependent on the extent of generational replacement.

Together, it is not exactly clear to what degree these scale issues matter for the Easterlin paradox. This is the case because of the limited literature that has devoted attention to these kinds of problems. For now, none of the presented issues seem detrimental for the gist of the paradox. But as the existence and relevance of the proposed problems becomes clearer with new literature, this will reflect back on the scientific relevance of the paradox.

---

[28] Plant (2022) mentions that there is no literature that addresses this specific concern. I have also not been able to find any that tests this empirically, but Ng (2022, p.35-36) does address the concern into some detail.

[29] Remember, the 'actual/underlying' happiness *(or, any subjective state)* is impossible to observe directly

[30] To give a numerical example: assume a simple two generation model. In t=0, the whole population consists of people of generation Y. This generation needs, on average, 100 'life satisfaction units' to reach an 8/10 satisfaction score. In t=20, the average income has risen substantially and that *caused* generation Y to experience 102 'life satisfaction units' resulting in an average score of 8.2/10. The new generation Z also experiences these 102 units, but needs 110 to reach an 8/10 score *(this is the intergenerational rescaling),* and therefore reports only a 7.2/10 satisfaction. Between t=0 and t=20, generational replacement has taken place *(a part of generation Y has died and is replaced by generation Z).* The trajectory of the aggregate satisfaction score therefore misrepresents that the underlying satisfaction has in fact gone up *(from 100 to 102).* The extent to which this matters for the paradox is thus dependent on the extent of rescaling between subsequent generations and the amount of generational replacement that has taken place *(which is directly related to the time frame one studies)*.

Concluding chapter 4, I have aimed to show economic problems that surround the paradox. Most importantly, there is the issue of causality that needs to be addressed. Next to that, problems regarding the measurement of income and happiness have been presented, and in the last section, problems on the measures used in the paradox have been put to the front. The fact that these problems are all but solved can explain why the debate on the Easterlin paradox rambles on today. The fact that the issues are rather heterogeneous in nature predicts that a multidisciplinary literature is needed to help move the debate one step further.

Chapter 5: Implications for policy

This chapter aims to reflect on what the paradox and the previous chapters can imply for policy. First, a general reflection on policy is given, then, in 5.1, I ask the question whether happiness needs to be promoted in order for societal well-being to be maximised. In 5.2, some attention will be given to the case for degrowth, and a warning will be given on the use of the paradox in propagating such ideas.

It is not directly clear what the implications of the paradox are for policymakers. This is mainly because - as has been shown multiple times in this thesis - the paradox is not about the causation of income on happiness, although it gets represented that way often. Therefore, in the most strict scientific view, the paradox implies nothing for policy, for policymakers are interested in - or *should* be interested in - the causal effects of policy instruments *(like promoting economic growth)* on happiness.

Easterlin and O'Connor (2020) indirectly state that the paradox implies that policymakers should not seek to increase GDP to improve society's overall well-being and that instead, the paradox promotes the movement to go ´beyond GDP´ to measure societal progress, like with measures of happiness. This wish of looking further than merely economic growth to promote societal well-being has been widely supported in the literature, and is not something on which the opinions of the paradox believers and the paradox deniers systematically seem to differ. In other words, believing that the paradox exists seems not to be the decisive factor in whether or not one believes that policymakers should look further than economic growth. Even if the paradox does not exist, the wish of going beyond GDP would still be a defensible position.[31]

Furthermore, the conclusion that because of the lack of a long term correlation between happiness and economic growth, a government that seeks to maximise happiness in the long run should therefore devote their attention to things 'beyond GDP' is drawn too hastily. This is because, as Bartolini and Sarrachino (2014) point out: *"the potential alternatives to GDP may suffer from the same adaptation and social comparison effects that prevent economic growth from having a positive impact on well-being in the long-term.".* To not apply double standards, in order to devote the attention away from GDP to something else, it first needs to be shown that the long term trends in happiness and that 'something else' *(unlike GDP)* **do** correlate. As long as it is not clear that 'something else' is not subject to the same exact shortcomings *(to the same extent)* as economic growth, this policy implication - that policymakers should focus less on economic growth in order to promote happiness - is premature. For example: as partly discussed in section 3.2, and following from Clark (2014), the mechanism of adaptation seems present in different 'non-GDP' aspects of life, including health, social capital and religion. It may therefore be that the Easterlin paradox is more about happiness than about economic growth. I therefore agree with Clark that one should be cautious in

[31] This is shown by Kamilcelebi and Veenhoven (2024) who are the most prominent paradox deniers. In their paper, they still do not disagree with the wish of 'going beyond GDP', as long as the effects of a lower growth rate on overall happiness are taken into account. (p.17)

diverting attention away from economic growth towards ‘something else’ on the basis of the paradox.

Another reason why translating the thesis into concrete policy advice is challenging is the existence of multiple ‘open ends’: hypotheses that have not been accepted or rejected yet. If the theory of the paradox deniers is correct – that there is a positive association between trends in income and happiness -, then the paradox might not even exist. The same holds for when the proposed phenomenon of intergenerational rescaling is proven to be real. In the case of the non-existence of the paradox, it is important to not apply double standards: a positive relationship in trends still does not imply that policymakers should seek to increase economic growth in order to facilitate happiness, for the same reason the nil relationship cannot imply that they should not: the lack of causality.

Other hypotheses that have been covered thus far that might be policy relevant are the two underexposed mechanisms in section 3.3. If happiness asymmetry is proven to be of importance, then the implication for policy would be to strive to minimize the probability of economic *decline*, and shorten the periods in which it happens, rather than to strive to promote the magnitude of economic *growth*. If the proposed mechanism of income inequality solves the paradox, then the message to policymakers could be that economic growth can promote happiness, if and only if the incomes of all are raised in a more or less proportionate matter.

As long as these hypotheses have not been accepted with some confidence, neither should the related policy advice. The distinct value of the paradox therefore lies not so much in its direct implications on policy *(for its lack thereof)*, but in its indirect implications: the insights it can provide into related matters. A discussion that is related to the paradox and that could have consequences for policy is the question whether policymakers should seek to increase happiness in order for societal well-being to be maximised.

5.1 Is happiness worth maximising?

This section asks the question whether a benevolent government, who seeks to maximise societal well-being, needs to maximise happiness. As discussed in section 4.3, Easterlin regards happiness and well-being to be more or less the same thing, and equivocates between the two concepts (See Easterlin 2003; 2013).

First, if one **accepts** the equivocation, it is still not exactly clear what it means to ‘maximise happiness’ on the macro scale. The measure that has been used mostly in all that has been covered thus far is the mean value of a subjective *(evaluative)* well-being score. Is this the thing on which the focus should lie? If so, what to think of the differences between correlates on the short versus the long run? If one solely focuses on things that seem to matter in the long run, one thereby foregoes correlates in the short run, of which economic growth is an example. It could be argued that policy interventions that temporarily boost happiness can still be of value, since every increase in happiness matters, no matter when it materialises. Furthermore, it can be argued that the *median* happiness score matters more than the *mean*. If so, one should focus on the distribution of - and inequality in happiness scores. Interestingly enough, Clark et al. (2014) show that within-country happiness inequality has fallen in the majority of countries that have experienced positive income growth. If the negative relationship between economic growth and happiness inequality is proven to be robust, aiming for the maximum median happiness therefore might allow a bigger role for economic growth. A different view is that the focus should lie on some form of total happiness, which could be operationalized as the amount of ‘happy life years’: the average happiness one has in life times the amount of years one lives. Veenhoven and Hagerty (2006) seem to propagate this measure.

This can even be taken one step further, by multiplying the mean of happy life years by total population. For this version of total happiness, life expectancy and population growth play as an important role as does mean happiness. Economic growth would then be relevant not only if it is associated with mean happiness growth, but also if it would be of causal importance for either or both longevity and population growth.

Second, if one **rejects** the equivocation, the goal of maximising societal well-being need not mean maximising happiness *(mean, median or total),* since the two are different concepts. The view that well-being is not only constituted by happiness is a common view in the well-being philosophical literature (Hersch, 2018). One could argue that some *objective* aspects should be part of the operationalization of well-being, no matter their relevance or influence on happiness, which is *subjective* in nature. There are many examples possible that could be worth pursuing in the context of maximising well-being because they have intrinsic or moral value to well-being, irrespective of their influence on happiness. Examples of such objective aspects are the minimalization of child mortality, the reduction of crime, or the improvement of the overall health of the population: even if it would be the case that these aspects hold no relationship with happiness *(for whatever reason, for example because of adaptation)* they would still matter for societal well-being if one's view of well-being is broader than happiness alone. The value of economic growth then hinges not only on the direct relationship it has with happiness, and also not only on its indirect relationship *(through mediating factors)* on happiness, but also on the relationship it has with objective factors that are deemed intrinsically or morally valuable for well-being. It is for this reason that Hersch (2018) claims that no matter the empirical side of things, no matter the direction of correlation between happiness and income, Easterlin may best be ignored when it comes to policy. I find it hard to argue against this view.

5.2 The Easterlin Paradox and the case for degrowth

According to Kamilcelebi and Veenhoven (2024), the message that economic growth will not make us any happier has been welcomed by critics of capitalism and by proponents of degrowth. This thus shows that the Easterlin Paradox may serve as an argument in propagating such ideas. Irrespective of my view on these movements, I argue against the use of the paradox in defending them.

In what has been covered, many supportive reasons can be found for my standpoint, for example the possible non-existence of the paradox, the existence of many 'open ends', the issue of causality and the previous section which argued the possibility of a total irrelevance of the paradox for whatever policy measure (of which degrowth is an example). Instead of repeating these arguments lengthy here, I want to add an orthogonal argument for the non-use of the paradox in discussing the case for degrowth: namely that in doing so, one commits a logical fallacy.

If one assumes that the paradox exists, and views the correlational evidence as of causal importance - thus that economic growth does not *cause* growth in happiness – this still cannot serve as evidence for the claim that, therefore, economic decline will not cause declining happiness. If anything, the literature shows that in the short run, economic decline is associated with happiness decline, and that this positive relationship is twice as strong as the positive short run relationship between happiness growth and economic growth. To then extrapolate the *(possible)* nil relationship between economic growth and happiness in the long run into the negative domain is premature – there is no data for countries that have experienced a long run negative growth rate – and rests thus on a logical fallacy. The Easterlin Paradox has no role in advocating for degrowth.

Concluding this chapter, all possible policy implications that follow from the paradox are surrounded by problems and/or uncertainty. It could be argued that the paradox implies nothing for policy. Therefore, I have presented a discussion that is related to the paradox that is of importance for policy, namely the question whether happiness is worth maximising in order to maximise societal well-being. Furthermore, I have explicitly argued against the use of the paradox in defending the case for degrowth.

Chapter 6: Research Agenda

In this chapter, a concise overview is presented that contains outstanding issues that are of scientific relevance for future research. I do not aim to give a grand overview of issues that are of importance for happiness economics or the causes of happiness or well-being in general, but limit the agenda to things that follow directly from what has been covered.

- On the empirical side of things, there is still no consensus on statement 3: whether the line that follows from analyses like the one behind Figure 10 has a positive slope or not, i.e. whether the paradox even exists or not. To settle this debate, I recommend an adversarial collaboration between Easterlin and Veenhoven. Instead of criticising each other's papers and telling why their 'scientific opponent' is wrong *(this process has been going on for more than 20 years now),* they should learn from Kahneman and Killingsworth in section 2.1. The latter duo decided to, after discovering that their findings were incompatible, combine their intellectual capabilities in order to seek truth, rather than to 'win' the debate. There is potentially substantial added value to be found in constructing a methodology that combines the strong aspects of both their analyses,[32] both concerning data aspects as concerning statistical aspects, that is agreed upon in advance. Furthermore, in this collaboration, I deem it necessary for both authors to reflect on this issue of causality, as this is the most pressing issue that is underrepresented in their work on the matter. Lastly, they should also give attention to the aspects where they agree upon, as it may turn out their disagreement is less severe than one would initially think.

- Assuming statement 3 to be correct, it is unclear what the mechanisms behind the paradox exactly are, and what their relative importance is. There is little literature that tries to quantify the relevance of the mechanisms on the paradox specifically. The established underexposed mechanisms of happiness asymmetry and income inequality deserve proper testing. Furthermore, it should be addressed that if one finds a mechanism that 'solves' the paradox, it need not mean that the others cannot be relevant. It could be the case that the others still work as inhibiting factors in the way income does or does not translate into happiness. In other words: even if the slope in Figure 10 is found to be positive, attention to the mechanisms could still explain why the line is not steeper, and is therefore of value.

- The mechanism of adaptation deserves special attention, since finding out whether and to what extent it is also relevant to the things that go 'beyond GDP' could be policy relevant, see chapter 5. Clark (2014) tried to review the literature, but found that it is limited in size. Now that ten years have passed, there again seems to lie value in a synthesis of the literature that has built up during this time. In light of measuring well-being and happiness, it seems relevant to find out what components of well-being adapt faster than others. Lastly, the connection with

[32] That of Easterlin and O'Connor (2020) and Kamilcelebi and Veenhoven (2024), as discussed in section 2.3.

psychological literature seems relevant in this respect if it can help answer questions about whether people can slow adaptation to good events and speed recovery from bad events.

- The limited research on rescaling and intergenerational rescaling deserves expansion. This is because the extent to which these are present are of fundamental importance in the way one should analyse and interpret the long run data that consists of results from subjective well-being scales *(for which people can decide themselves what the begin- and endpoint represent).* As discussed, Plant (2020) does a step in the right direction, but this limited attention is still not proportional to the possible consequences it could have on the use of such scales.

Chapter 7: Concluding thoughts

Concluding this thesis, I want to give some attention to the research question that has been posed at the start: *"What is the current state of research into the Easterlin Paradox and what are the remaining economic issues to tackle to help the debate move one step further?"*

In chapter 2, I have tried to show the current scientific standing on the matter. There is consensus on statement 1 and 2, but the problem starts with statement 3. Despite more than fifty years of empirical work, the answer to the question whether trends in income and happiness correlate cannot be given with enough confidence. It is important to resolve this, as discussing all that has been said in response to the existence of the paradox might collapse once its foundation *(statement 3)* turns out to be rather wobbly. In chapter 6, I have laid out a possibility to do so: an adversarial collaboration.

Furthermore, there are underlying issues regarding the proposed mechanisms. It has not been proven that they are of importance specifically in explaining the paradox. A different view is that they are merely inhibiting factors in the relationship between income and happiness. There are underexposed mechanisms that also can be of relevance in explaining the paradox *(partly).*

No matter the (non-)existence of the correlation and the relevance of the mechanisms therein, the elephant in the room seems to be the issue of causality. Surprisingly, in the discourse around the paradox, this problem – that restricts what can be implied from the paradox – is mentioned very seldom.

The water gets even more troubled with the use of words like happiness, subjective well-being and evaluative well-being, as these concepts get thrown together to mean more or less the same thing. This blurs the discussion and makes it difficult to participate in a focused scientific debate. The implications for policy are as vague, and I have argued against the use of the paradox in concrete policy advice and in defending the case for degrowth.

It is for these kinds of reasons that the discussion around the Easterlin Paradox has not resulted yet in a satisfactory consensus. There are many steps to take in order for a consensus to be reached, some of which I have tried to describe in chapter 6.

All in all, these open ends might leave one with a dissatisfying aftertaste. Though I empathise with this feeling, one has to accept that given the current state of affairs, making clear what cannot be implied or claimed is more suitable than what can. In doing so, one foregoes

exciting but hasty conclusions and shows rigorous scientific prudence, which, one can dream, might just be the source of true happiness.[33]

Acknowledgements

First and foremost, I want to thank my supervisor, prof. dr. Johan Graafland. I am grateful for the fact that he has devoted his scarce time to read my texts and ideas and deeply value the substantial effort that has been put in trying to improve the thesis.

The same goes for the second reader, em. prof. dr. Lans Bovenberg, who was the supervisor for my bachelor thesis. I appreciate the time it takes to read and reflect on what I have written.

The fact that I feel taken seriously by men with such major academic track records when discussing economic and philosophical ideas fills me with great honour and humility.

Lastly, I want to thank the one person that has been the most important causal driver for my own happiness and well-being the past months: Marte. Your support throughout the process has impacted me greatly.

References

Abramovitz, M. (1959). The welfare interpretation of secular trends in national income and product. In The Allocation of economic resources: Essays in honor of Bernard Francis Haley (M. Abramovitz et al.). Stanford, California: Stanford Univ. Press.

Barazetta, G. (2014). Health improvements and adaptation: Evidence from a randomized experiment. *Journal of Health Economics, 35*, 149-158.

Barker, D. J. P. (2005). The developmental origins of adult disease. *Journal of the American College of Nutrition, 23*(6), 588S-595S.

Bartolini, S., & Sarracino, F. (2014). Happy for how long? How social capital and economic growth relate to happiness over time. *Ecological Economics, 108*, 242-256.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2005). Bad is stronger than good. *Review of General Psychology, 9*(4), 323-370.

Beja, E. L. (2017). Subjective well-being approach to the Easterlin Paradox. *Applied Research in Quality of Life, 12*(4), 907-925.

Boyce, C. J., Wood, A. M., Banks, J., Clark, A. E., & Brown, G. D. A. (2013). Money, well-being, and loss aversion: Does an income loss have a greater effect on well-being than an equivalent income gain? *Psychological Science, 24*(12), 2557-2562.

Bradburn, N. M. (1969). *The structure of psychological well-being*. Aldine Publishing Company.

Burchardt, T. (2005). Are one man's rags another man's riches? Identifying adaptive expectations using panel data. *Journal of Economic Inequality, 3*(3), 197-218.

Cantril, H. (1965). The pattern of human concerns. New Brunswick, New Jersey: Rutgers Univ. Press.

[33] What about those million dollars you got offered at the start, you may ask? Well, the insights you have gained by reading this thesis are worth at least that much. May it cause you a great deal of happiness.

Clark, A. E. (2003). Unemployment as a social norm: Psychological evidence from panel data. *Journal of Labor Economics, 21*(2), 323-351.

Clark, A. E. (2014). Happiness, habits and high rank: Comparisons in economic and social life. In Helliwell, J., Layard, R., & Sachs, J. (Eds.), *World Happiness Report 2012* (pp. 65-83). New York: Earth Institute.

Clark, A. E. (2016). Adaptation and the Easterlin paradox. In K. F. Zimmermann (Ed.), *Handbook of labor, human resources and population economics*. Springer.

Clark, A. E., & Georgellis, Y. (2013). Back to baseline in Britain: Adaptation in the British Household Panel Survey. *Economica, 80*(319), 496-512.

Clark, A. E., Flèche, S., & Senik, C. (2014). Economic growth evens out happiness: Evidence from six surveys. *The Review of Income and Wealth, 60*(3), 596-619.

Clark, A. E., Frijters, P., & Shields, M. A. (2008). Relative income, happiness, and utility: An explanation for the Easterlin paradox and other puzzles. *Journal of Economic Literature, 46*(1), 95-144.

Clark, A. E., Frijters, P., & Shields, M. A. (2017). *Origins of happiness: Science of well-being over the life course*. Princeton University Press.

Clark, A. E., Kristensen, N., & Westergård-Nielsen, N. (2009a). Job satisfaction and coworker wages: Status or signal? *Economic Journal, 119*(536), 430-447.

Clark, A. E., Kristensen, N., & Westergård-Nielsen, N. (2009b). Economic satisfaction and income rank in small neighbourhoods. *Journal of the European Economic Association, 7*(2-3), 519-527.

De Neve, J.-E., Clark, A. E., De Keulenaer, F., Krekel, C., Layard, R., & Ward, G. (2018). Life satisfaction and sensitivity to growth and recession across Europe. *Review of Economics and Statistics, 100*(3), 362-375.

Deaton, A. (2008). Income, health, and well-being around the world: Evidence from the Gallup World Poll. *Journal of Economic Perspectives, 22*(2), 53-72.

Di Tella, R., Haisken-DeNew, J. P., & MacCulloch, R. (2010). Happiness adaptation to income and to status in an individual panel. *Journal of Economic Behavior & Organization, 76*(3), 834-852.

Diener, E. (2009). Subjective well-being. In E. Diener (Ed.), *The science of well-being: The collected works of Ed Diener* (pp. 11-58). Springer.

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.

Diener, E., Tay, L., & Oishi, S. (2013). Income and happiness. *Psychological Science, 24*(1), 1-9.

Dolan, P., Peasgood, T., & White, M. (2008). Do we really know what makes us happy? A review of the economic literature on the factors associated with subjective well-being. *Journal of Economic Psychology, 29*(1), 94-122.

Easterlin, R. A. (1974). Does economic growth improve the human lot? Some empirical evidence. In *Nations and households in economic growth* (pp. 89-125). Academic press.

Easterlin, R. A. (1995). *Will raising the incomes of all increase the happiness of all?*. Journal of Economic Behavior & Organization, 27(1), 35-47.

Easterlin, R. A. (2001). Income and happiness: Towards a unified theory. *The Economic Journal, 111*(473), 465-484.

Easterlin, R. A. (2003). Explaining happiness. *Proceedings of the National Academy of Sciences, 100*(19), 11176-11183.

Easterlin, R. A. (2010). *Happiness, growth, and public policy*. Economic Inquiry, 48(2), 311-332.

Easterlin, R. A. (2013) "Cross-sections are history." *Population and Development Review* 38: 302-308.

Easterlin, R. A. (2015). *Happiness and economic growth: The evidence*. In L. Bruni & P. L. Porta (Eds.), *Handbook of research methods and applications in happiness and quality of life* (pp. 369-386). Edward Elgar Publishing.

Easterlin, R. A. (2016). *Paradox lost?*. Review of Behavioral Economics, 4(4), 311-339.

Easterlin, R. A., & Angelescu, L. (2011). Modern economic growth and quality of life: Cross-sections of large surveys. In M. Maggino (Ed.), *Quality of life and the millennium challenge: Advances in quality-of-life studies, theory and research* (pp. 43-58). Springer.

Easterlin, R. A., & O'Connor, K. J. (2020). The Easterlin Paradox (IZA Discussion Paper No. 13923). Institute of Labor Economics (IZA).

Easterlin, R. A., & O'Connor, K. J. (2022a). The Easterlin paradox. In *Handbook of labor, human resources and population economics* (pp. 1-25). Cham: Springer International Publishing.

Easterlin, R. A., & O'Connor, K. J. (2022b). Explaining happiness trends in Europe. *Proceedings of the National Academy of Sciences* 119.37

Easterlin, R. A., McVey, L. A., Switek, M., Sawangfa, O., & Zweig, J. S. (2010). *The happiness-income paradox revisited*. Proceedings of the National Academy of Sciences of the United States of America, 107(52), 22463-22468.

Elsas, M. (2021). Happiness and income: The causal direction matters. *Journal of Economic Behavior & Organization, 183*, 1-16.

Graham, C., & Felton, A. (2006). Inequality and happiness: Insights from Latin America. *Journal of Economic Inequality, 4*(1), 107-122.

Graham, C., Eggers, A., & Sukhtankar, S. (2005). Does happiness pay? An exploration based on panel data from Russia. *Journal of Economic Behavior & Organization, 55*(3), 319-342.

Graham, C., Chattopadhyay, S., & Picon, M. (2010). The Easterlin and other paradoxes: Why both sides of the debate may be correct. In E. Diener, J. F. Helliwell, & D. Kahneman (Eds.), International differences in well-being (pp. 247-288). Oxford University Press.

Gurin, G., Veroff, J., & Feld, S. (1960). *Americans view their mental health*. Basic Books.

Harter, J. K., & Arora, R. (2008). The impact of time spent working and job fit on well-being around the world. *Journal of Occupational and Organizational Psychology, 81*(1), 25-46.

Hersch, P. (2018). Ignoring Easterlin: Why Easterlin's correlation findings need not matter to public policy. In M. Rojas (Ed.), *The Economics of Happiness: How the Easterlin Paradox Transformed Our Understanding of Well-Being and Progress* (pp. 123-145). Springer Nature.

Hirschauer, N., Lehberger, M., & Musshoff, O. (2015). Happiness and utility in economic thought: Or, what can we learn from happiness research for public policy analysis and public policy making? *Social Indicators Research, 121*(3), 647-674.

Hirschman, A. O. (1973). *The changing tolerance for income inequality in the course of economic development*. World Bank.

Howell, R. T., & Howell, C. J. (2008). The relation of economic status to subjective well-being in developing countries: A meta-analysis. *Psychological Bulletin, 134*(4), 536-560.

Inkeles, A. (1960). *Industrial man: The relation of status to experience, perception, and values*. Harper & Row.

Kahneman, D., & Deaton, A. (2010). High income improves evaluation of life but not emotional well-being. Proceedings of the National Academy of Sciences, 107, 16489-16493.

Kahneman, D., & Krueger, A. B. (2006). Developments in the measurement of subjective well-being. *Journal of Economic Perspectives, 20*(1), 3-24.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263-291.

Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., & Stone, A. A. (2004). A survey method for characterizing daily life experience: The day reconstruction method. *Science, 306*(5702), 1776-1780.

Kaiser, C., & Vendrik, M. C. M. (2019). Different versions of the Easterlin Paradox: New evidence for European countries. In M. Rojas (Ed.), *The economics of happiness: How the Easterlin Paradox transformed our understanding of well-being and progress* (pp. 27–56). Springer Nature.

Kamilçelebi, H., & Veenhoven, R. (2024). Economic growth and happiness in nations between 1946-2019: The demise of the Easterlin Paradox

Killingsworth M. A. (2021). Experienced well-being rises with income, even above $75,000 per year. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(4)

Killingsworth, M. A., Kahneman, D., & Mellers, B. (2023). Income and emotional well-being: A conflict resolved. *Proceedings of the National Academy of Sciences of the United States of America*, *120*(10)

Kingdon, G. G., & Knight, J. (2007). Community, comparisons and subjective well-being in a divided society. *Journal of Economic Behavior & Organization, 64*(1), 69-90.

Knight, J., & Song, L. (2006). Subjective well-being and its determinants in rural China. *China Economic Review, 17*(2), 395-411.

Layard, R., Clark, A. E., & Senik, C. (2009). The determinants of happiness. *Journal of Economic Literature, 47*(4), 1153-1172.

Lindqvist, E., Östling, R., & Cesarini, D. (2020). Long-run effects of lottery wealth on psychological well-being. *Review of Economic Studies, 87*(6), 2703-2726.

Luhmann, M., Hofmann, W., Eid, M., & Lucas, R. E. (2012). Subjective well-being and adaptation to life events: A meta-analysis. *Journal of Personality and Social Psychology, 102*(3), 592-615.

Luttmer, E. F. P. (2005). Neighbors as negatives: Relative earnings and well-being. *Quarterly Journal of Economics, 120*(3), 963-1002.

Lyubomirsky, S. (2011). Hedonic adaptation to positive and negative experiences. In S. Folkman (Ed.), *The Oxford handbook of stress, health, and coping* (pp. 200-224). Oxford University Press.

Marx, K. (1849). *Contribution to the critique of political economy*. Progress Publishers.

Meyer, B. D., & Sullivan, J. X. (2003). Measuring the well-being of the poor using income and consumption. *Journal of Human Resources, 38*(Supplement), 1180-1220.

Mikucka, M., Sarracino, F., & Dubrow, J. K. (2017). When does economic growth improve life satisfaction? Multilevel analysis of the roles of social trust and income inequality in 46 countries, 1981-2012. *World Development, 93*, 447-459.

Mõttus, R., Realo, A., Allik, J., Ausmees, L., Henry, S., McCrae, R. R., & Vainik, U. (2024). Most people's life satisfaction matches their personality traits: True correlations in multitrait, multirater, multisample data. *Journal of Personality and Social Psychology, 126*(4), 676-693.

Ng, Y. K. (1996). Happiness surveys: Some comparability issues and an exploratory survey based on just perceivable increments, *Social Indicators Research*. 38, 1–27.

Ng, Y. K. (1997). A case for happiness, cardinalism, and interpersonal comparability, *The Economic Journal*, 107(445), 1848–1858.

Ng, Y. K. (2008). Happiness studies: Ways to improve comparability and some public policy implications, *Economic Record*. 84(265) 253–266.

Ng, Y. K. (2015). Some conceptual and methodological issues of happiness: Lessons from Evolutionary Biology, The Singapore Economic Review, 60(4).

Ng, Y.-K. (2022). *Happiness—Concept, Measurement and Promotion*. Springer Link.

Nilsson, A. H., Eichstaedt, J. C., Lomas, T., Schwartz, A., & Kjell, O. (2024). The Cantril Ladder elicits thoughts about power and wealth. *Scientific Reports, 14*(1), Article 2642.

Nordhaus, W., & Tobin, J. (1972). Is growth obsolete? In *Economic growth, 5th Anniversary Series* (pp. 1-80). National Bureau of Economic Research. New York, NY: Columbia University Press.

Oishi, S., & Kesebir, S. (2015). Income inequality explains why economic growth does not always translate to an increase in happiness. *Psychological Science, 26*(10), 1630-1638.

Organisation for Economic Co-operation and Development (OECD). (2013). OECD guidelines on measuring subjective well-being. OECD Publishing.

Oswald, A. J., & Powdthavee, N. (2008). Does happiness adapt? A longitudinal study of disability with implications for economists and judges. *Journal of Public Economics, 92*(5-6), 1061-1077.

Perez-Truglia, R. (2013). A test of the conspicuous consumption model using subjective well-being data. *Review of Economics and Statistics, 95*(2), 459-474.

Plant, M. (2020). *A happy possibility about happiness (and other) scales*. Happier Lives Institute.

Plant, M. (2022). Will faster economic growth make us happier? The relevance of the Easterlin Paradox to Progress Studies. Happier Lives Institute.

Robinson, J. P., & Shaver, P. R. (1969). *Measures of social psychological attitudes*. Institute for Social Research.

Sacks, D. W., Stevenson, B., & Wolfers, J. (2012). Subjective well-being, income, economic development and growth. *Proceedings of the National Academy of Sciences of the United States of America, 109*(49), 19953-19958.

Senik, C. (2004). When information dominates comparison: Learning from Russian subjective panel data. *Journal of Public Economics, 88*(9-10), 2099-2123.

Smith, A. (1776). *An inquiry into the nature and causes of the wealth of nations*. Strahan & Cadell.

Stevenson, B., & Wolfers, J. (2008). Economic growth and subjective well-being: Reassessing the Easterlin Paradox. *Brookings Papers on Economic Activity, 2008*(1), 1-87.

Stevenson, B., & Wolfers, J. (2013). Subjective well-being and income: Is there any evidence of satiation? *American Economic Review, 103*(3), 598-604.

Thompson, E. H., & Marks, A. D. G. (2008). Parental status and well-being: A research update. *Journal of Marriage and Family, 70*(3), 612-629.

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *Quarterly Journal of Economics, 106*(4), 1039-1061.

Van Praag, B. M. S., & Frijters, P. (1999). The measurement of welfare and well-being: The Leyden approach. In D. Kahneman, E. Diener, & N. Schwarz (Eds.), *Well-being: The foundations of hedonic psychology* (pp. 413-433). Russell Sage Foundation.

Veblen, T. (1899). *The theory of the leisure class: An economic study of institutions*. Macmillan.

Veenhoven, R., & Hagerty, M. R. (2006). Rising happiness in nations 1946–2004: A reply to Easterlin. *Social Indicators Research, 79*(3), 421-436.

Veenhoven, R., & Vergunst, F. (2014). The impact of economic growth on happiness. In J. F. Helliwell, R. Layard, & J. Sachs (Eds.), *World happiness report 2014* (pp. 311). Sustainable Development Solutions Network.

Wilson, W. R. (1967). Correlates of avowed happiness. *Psychological Bulletin, 67*(4), 294-306.

Wu, F. (2020). An examination of the effects of consumption expenditures on life satisfaction in Australia. *Journal of Happiness Studies*. Advance online publication.

Wysocki, T. A., Lawson, D. F., & Koski, M. H. (2022). Statistical control requires causal justification. *American Psychologist, 77*(1), 79-92.

Ye, S., Li, J., & Liu, H. (2023). The causal effect of income on happiness: Evidence from Chinese twins. *Economic Modelling, 115*, 105988.