Will faster economic growth make us happier?

The relevance of the Easterlin Paradox to Progress Studies

Michael Plant

June 2022

Happier Lives Institute

Contents

This is a transcript of a talk by Michael Plant at the Moral Foundations of Progress Studies workshop at the University of Texas in March 2022. Or rather, it's a re-recorded and edited version of the talk that was subsequently produced for a Global Priorities Institute reading group on 'progress' and then updated in light of many helpful comments from that seminar. The original slide deck can be viewed here.

---

1. Introduction

As I understand it, Progress Studies is a nascent intellectual field which starts by asking the question, "Since we seem to have gotten a lot of progress over the last couple of hundred years, where did this come from, and what can we do to get more of it?" (Vox, 2021).

Progress Studies has been popularised by academics such as Tyler Cowen and Steven Pinker. However, the Easterlin Paradox presents a real challenge to the claim that if we want more progress, we just need to improve the long-run growth rate - a view that Cowen argues for in his book *Stubborn Attachments*. This is a possible version of Progress Studies and the one I'm responding to.

So what is the Easterlin Paradox? Quoting Easterlin and O'Connor (2022), the Easterlin Paradox states:

> *At a point in time, happiness varies directly with income both among and within nations, but over time the long-term growth rates of happiness and income are not significantly related.*

There is a common view that economic growth is going to make our lives better, but the Easterlin Paradox challenges this.

What's paradoxical is that at a given point in time, richer people are more satisfied than poorer people and richer countries are more satisfied than poorer countries, but over the course of time, countries which grow faster don't seem to get happier faster. In other words, if I get richer, that will be good for me, but if we all get richer, that won't do anything for us collectively.

While subjective wellbeing (self-reported happiness and life satisfaction) has gone up in previous decades, the challenge of the Easterlin Paradox is that countries which grow faster do not seem to be getting happier faster; growth *per se* seems unrelated to average subjective wellbeing. If the paradox holds, the result would be striking and significant. It would suggest that, if we want to increase average wellbeing, we must not rely on growth, but go back to the drawing board and see what really works.

There's been quite a bit of debate over the nature and existence of the Paradox. The topic first emerged in 1974 when Richard Easterlin published a paper called, *Does Economic Growth Improve the Human Lot?* It's been particularly challenged by Stevenson and Wolfers (2008), who claim the paradox is an illusion and growth is making us happier. However, after looking into this myself, I actually think that Easterlin has the better half of the debate and the paradox does propose a real challenge to the idea that economic growth alone will make us happier.

My main purpose here is to explain what the Easterlin Paradox is and why - despite doubts - we need to take it seriously. My second purpose is to show that we can work out how to improve subjective wellbeing in society and make some tentative suggestions about this. However, this project is only starting to be taken seriously and there is lots more work to be done.

2. Evidence for the Paradox

So where does the Easterlin Paradox data come from? It's based on survey questions such as:

- Taking all things together, how would you say things are these days?
- Would you say you're very happy, pretty happy, or not too happy?[1]
- All things considered, how satisfied are you with your life as a whole nowadays, from one, dissatisfied to ten, satisfied?[2]
- The Cantril Ladder[3], which asks individuals to imagine a ladder with different steps where the top step is the best possible life, the bottom step is the worst possible life, and individuals are asked which step they're on.

What makes the paradox a paradox is that happiness is related to income *at a time* (in the cross-section) but not *over time* (in the time series).

To set up the paradox, let's quickly look at the cross-sectional data first...

[1] US General Social Survey

[2] World Values Survey

[3] Gallup World Poll

There are quite a few bits of evidence which suggest that richer people are happier and/or more satisfied than poorer people. The 'classic' result is Kahneman and Deaton (2010) indicated in the figure below.

If you look at the thick dotted line, which is the Cantril Ladder of Life, you see that the higher your income the more satisfied you are. Notice too that the x-axis is logarithmic, so it doubles at each point. But if you look at positive affect, which was measured using a basket of indicators of positive emotion, then this starts to plateau around $65,000 dollars of household income. There's a similar story for measures of stress.

What people took from this finding was that if you get richer you'll get more satisfied, but happiness - which is distinct from satisfaction - flattens out at some point.

Figure 1: The relationship between income and happiness (Kahneman and Deaton, 2010)

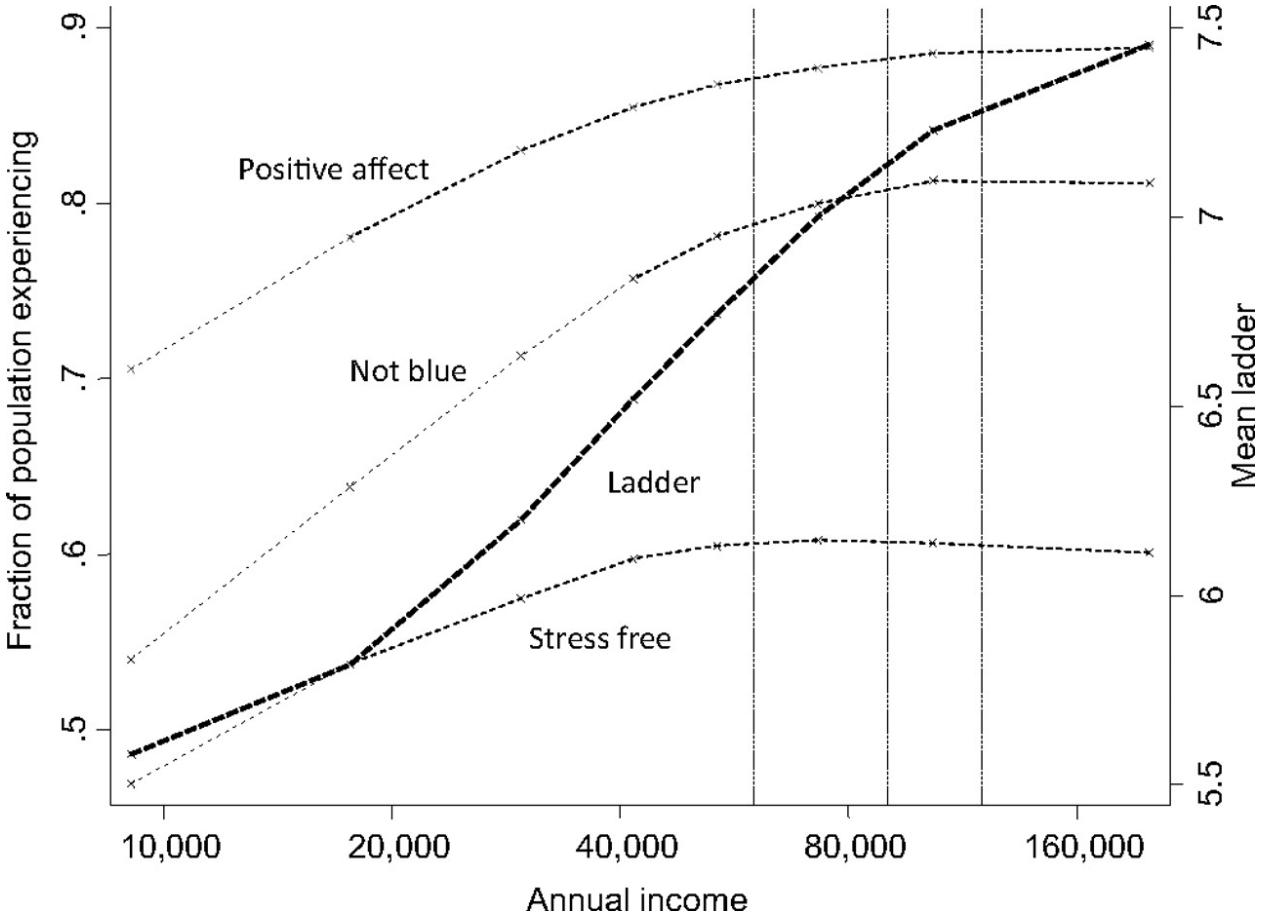

Since then there's been new evidence from Killingsworth (2021) which has caused quite a stir. Using smartphone data where people report how they feel in the moment, Killingsworth finds that both experienced wellbeing and life satisfaction keep going up with larger amounts of household income.

Again, note the logarithmic x-axis.

Figure 2: The relationship between household income and wellbeing (Killingsworth, 2021)

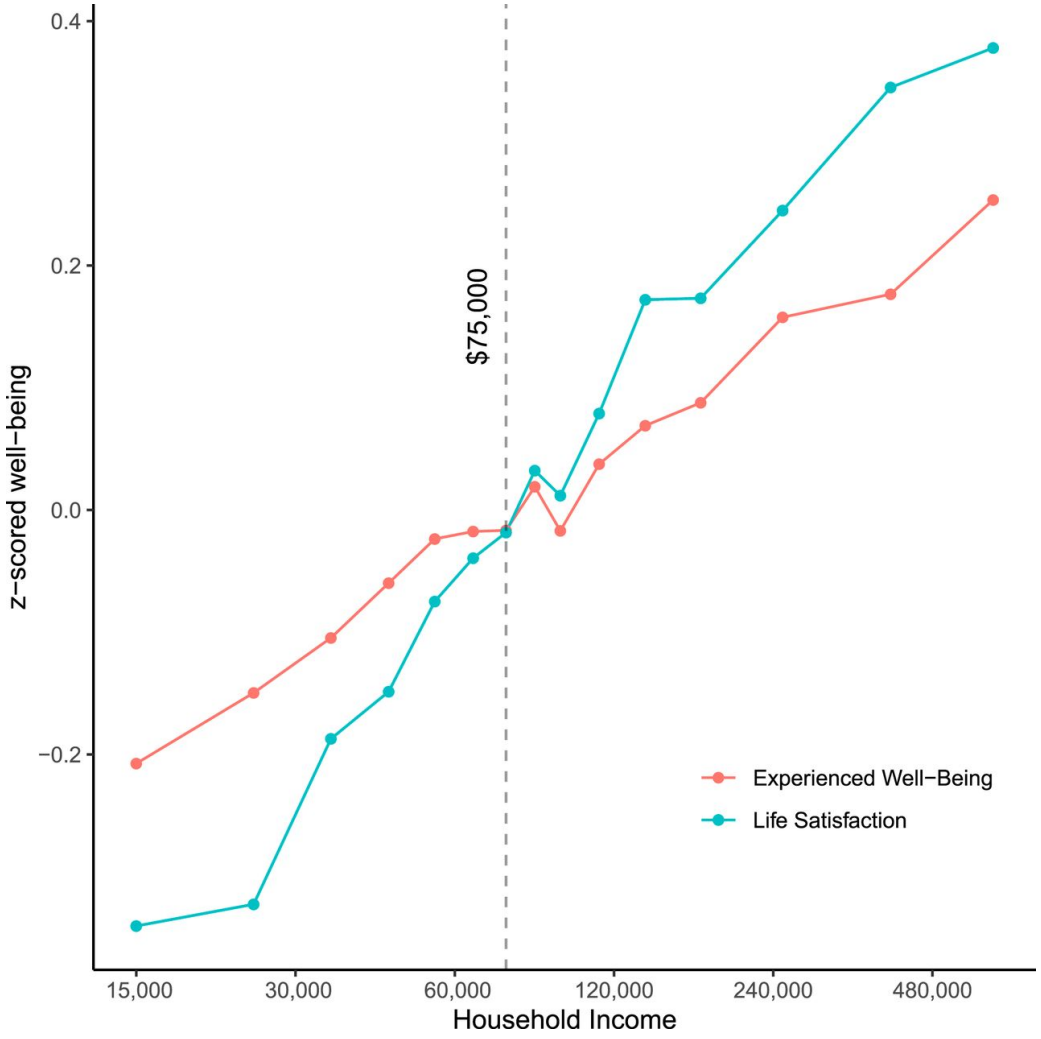

That's an indication of the relationship between money and wellbeing at a time *within* a country. What about at a time *between* countries?

For this, we can look at self-reported life satisfaction across the world as shown in this figure from Our World in Data using data from the World Happiness Report 2020. The colours represent different levels of life satisfaction. At a glance, there is a more or less believable pattern, with the more developed countries being more satisfied than the less developed countries.

Figure 3: Self-reported life satisfaction, 2020 (Our World in Data)

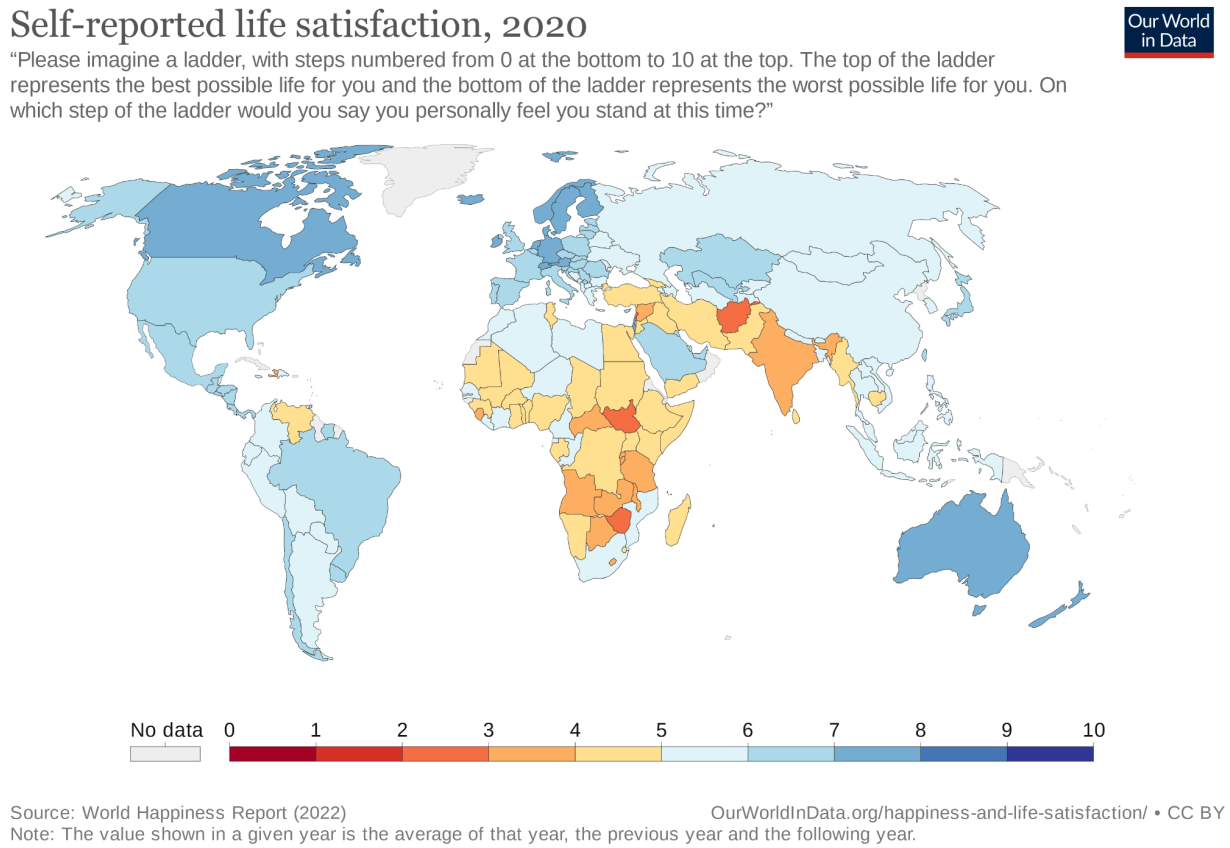

So, more money leads to more happiness, right? Well, if you look at what happens over time, we get a different, indeed paradoxical picture.

Below are a few different bits of evidence from a few rich countries — different data sets where the white dots represent the average subjective wellbeing (using slightly different measures) and the red line is GDP per capita. The figure is from the excellent book *Origins of Happiness* (Clark et al., 2019). We can see that GDP per capita goes up sharply over time whereas the subjective wellbeing scores seem to go nowhere in particular.

This is the sort of thing which motivates the paradox. It's better to be richer at any given moment, but as we all get richer, it seems like average wellbeing doesn't increase.

Figure 4: The relationship between GDP per capita and subjective wellbeing (Clark et al., 2019)

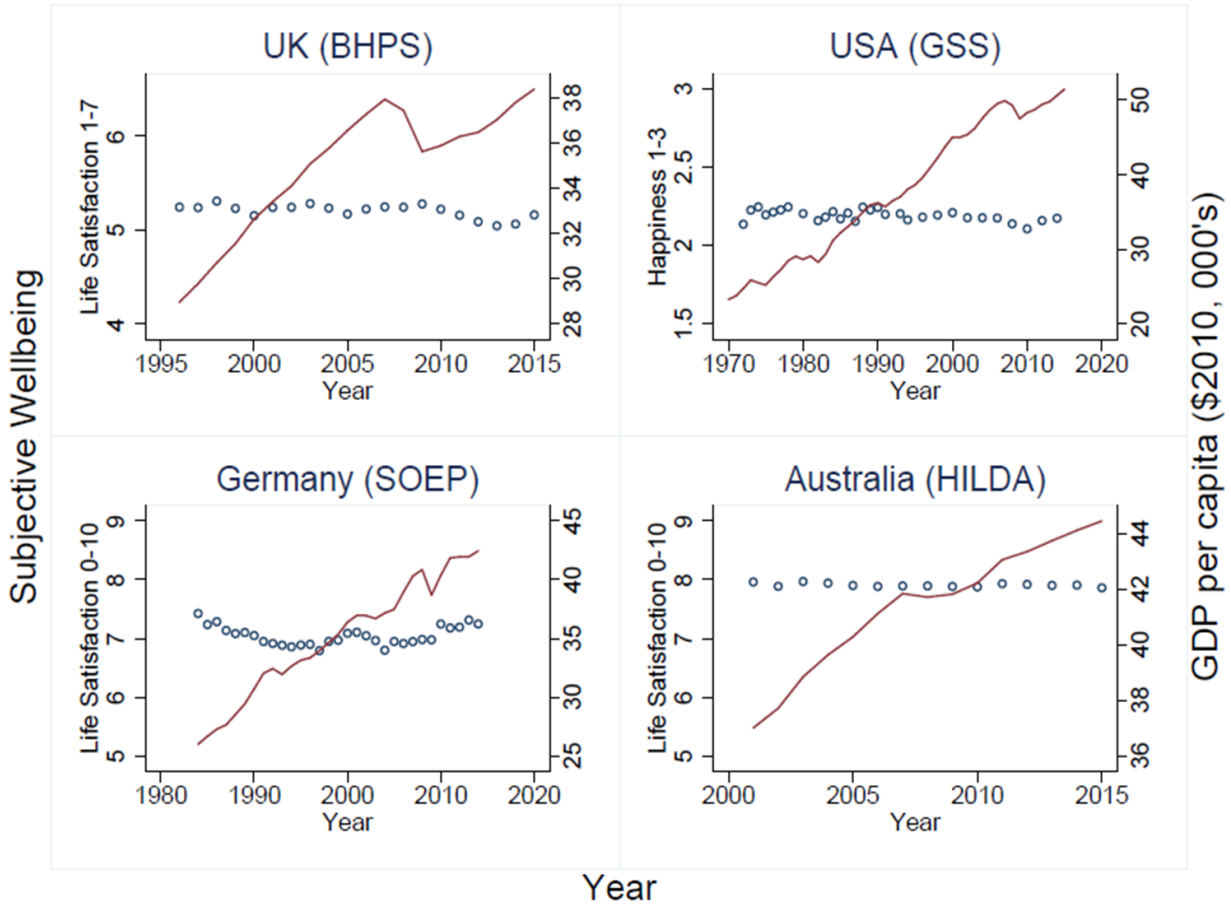

Now it's worth really stressing here that the Easterlin Paradox is about what happens over the long run. It's really not at all controversial that if you're looking at the short run over the business cycle, you would expect wellbeing to go up in a boom and down in a bust. But the interesting thing - to really test the effect of our increasing resources - is to see what happens over the longer term.

As this figure from Easterlin and O'Connor (2022) indicates, you need to look over longer time periods to know if you're in situation Y or situation H. You've got to look beyond the initial fluctuations to really test the paradox, so we can't just look at only a handful of years.

Figure 5: Short-term fluctuations and long-term trends in happiness (H) and income (Y) (Easterlin and O'Connor, 2022)

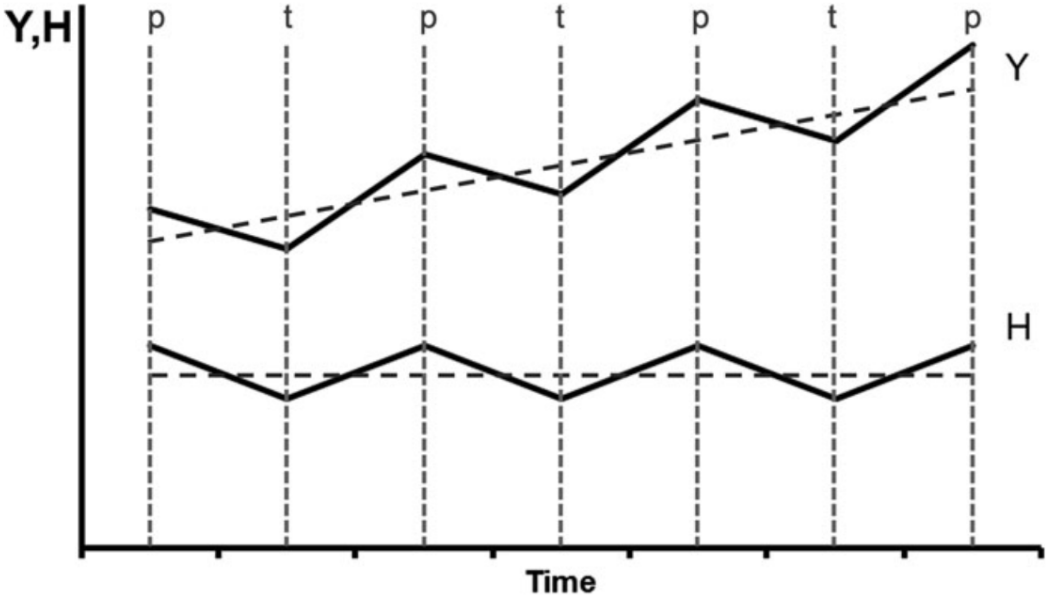

Easterlin and O'Connor (2022) argue, I think convincingly, that to test the paradox we need to look over the long run. We also need to look at many different countries, because if we look at an individual country, there could be all sorts of different things which are only happening there during a given period of time. So, as a better test of this question, we need to look over a wider basket of countries and then compare changes in subjective wellbeing to GDP growth rates. And we really want to know whether subjective wellbeing grows faster when GDP growth is higher.

Now, a challenge in doing this is the data we have. We've only started to measure subjective wellbeing fairly recently in the course of human history, since the 1950s, and there are not many long-running datasets for different countries. In the available data sets, a problem that Easterlin and O'Connor (2022) make a big song and dance about is that there are lots of European countries that are only

observed post-transition from socialism to capitalism. When the USSR collapsed there were big convulsions in those countries - society is sort of torn asunder and then put back together. There's also a big drop in the economy as a result of this which then improves over time. Clearly, there's something a bit odd about looking at countries which are recovering after a trough ('expansion-only countries') compared to those which are going completely through a trough ('full-cycle' countries).

In response, Easterlin and O'Connor (2022) distinguish between expansion-only countries versus full-cycle countries - those which have boom and bust, not just boom. This might seem like a controversial assumption they make just to support their conclusion, but it strikes me as appropriate given what they are trying to test. Other researchers make the same point (e.g. Bartolini and Sarracino, 2014; Kaiser and Vendrik, 2019).

Below, we can see some more figures from Easterlin and O'Connor (2022)...

Figure 6 presents data from the World Values Survey of over 54 countries and with an average 28-year duration comparing the compound annual growth rate to changes in life satisfaction. First, you can see that there are countries all over the place, and if you look at the thick brown line, which excludes the boom-only countries, then you can see that the relationship is basically nil. In other words, a higher rate of GDP growth doesn't correlate with countries getting happier. There is a positive relationship if all the expansion-only countries are included but, again, that doesn't seem the right test because we want to see what the relationship is over the longer run when economies expand *and* contract.

Figure 6: Growth rates of life satisfaction and GDP per capita with and without expansion only countries. WVS/EVS Data. 1981-2019 (Easterlin and O'Connor, 2022)

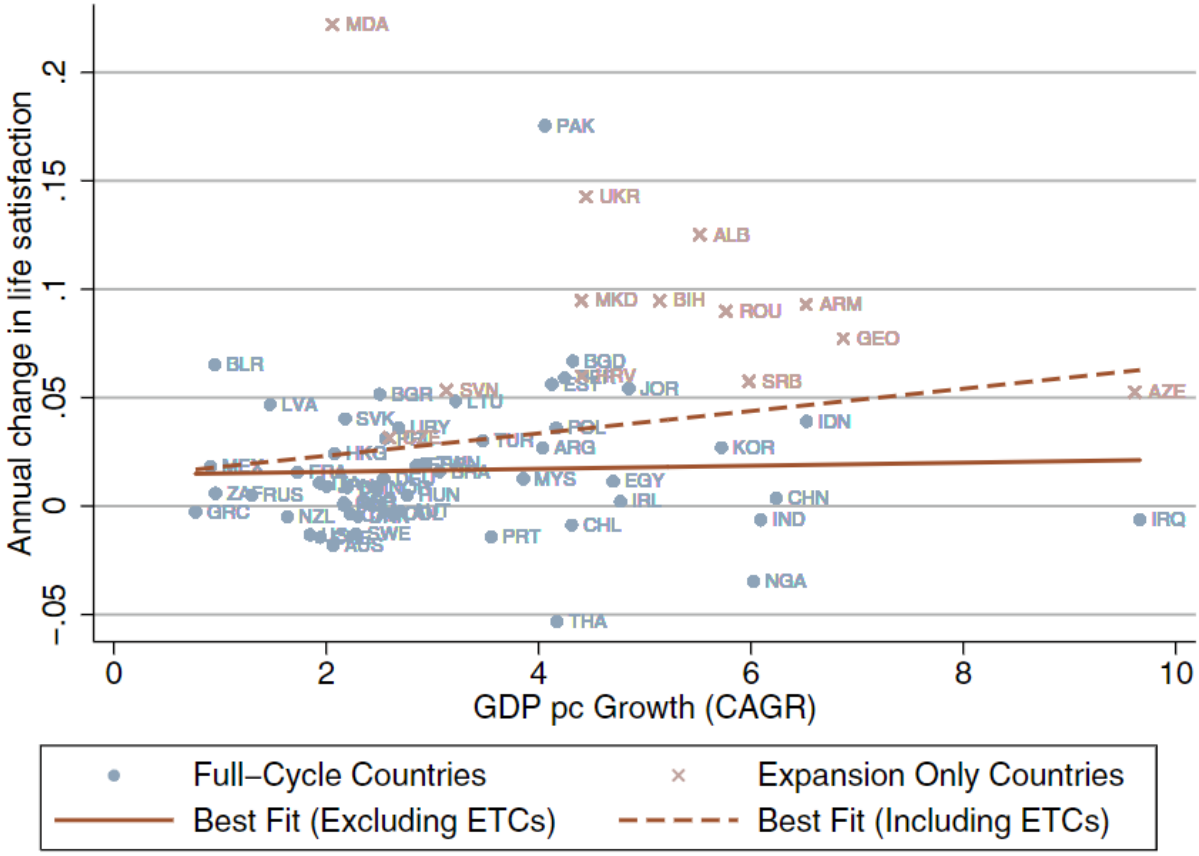

Of course, lots of other things can be happening in countries besides economic growth, a point Easterlin and O'Connor make. Those other things are probably what will explain the changes in happiness.

Below we can see a similar data set from the Gallup World Poll which includes a wider range of countries over a shorter time period than the World Values Survey. It's a pretty similar story. The thick brown line for the full cycle countries is, again, somewhat underwhelming. Easterlin and O'Connor say that the regression lines for full cycle countries are positive, but statistically and "economically insignificant" (I think it would be better if they said "practically insignificant" but never mind).

Figure 7: Growth rates of best possible life and GDP per capita with and without transition countries. Gallup Data. 2005-2019. (Easterlin and O'Connor, 2022)

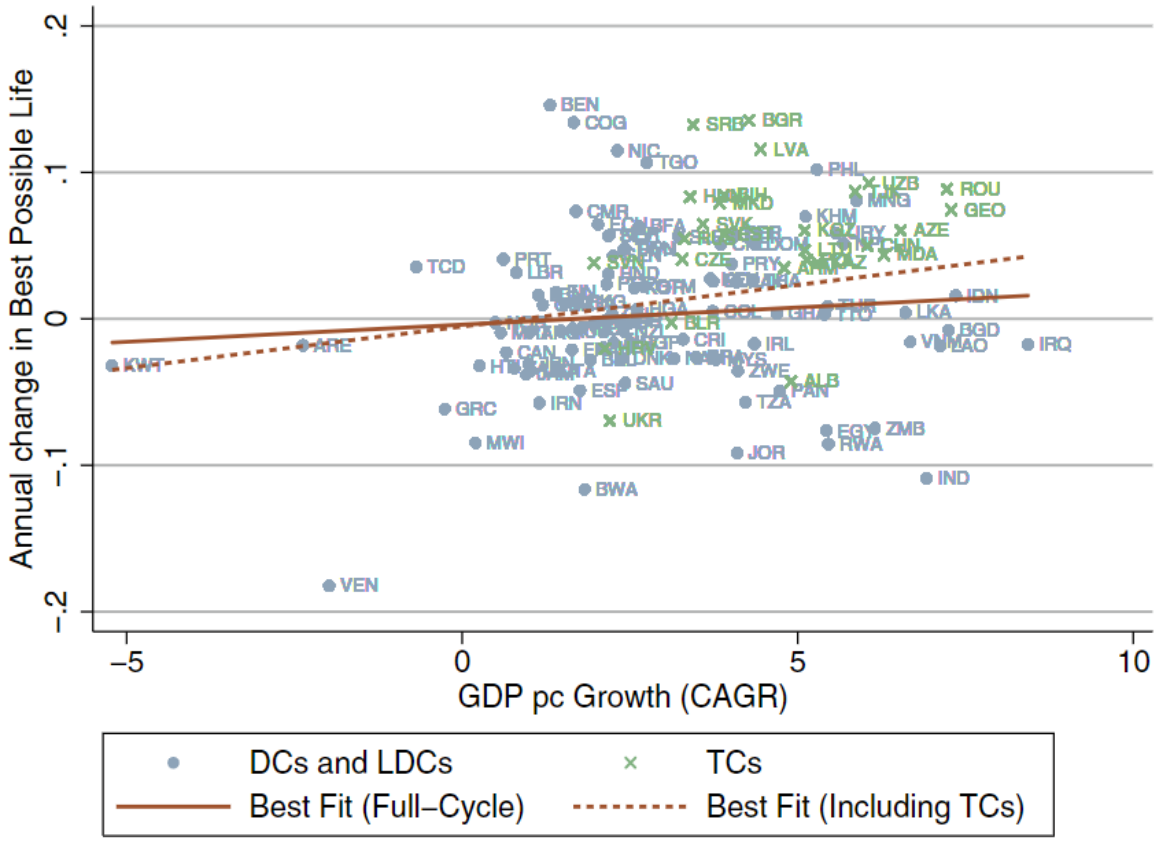

What does "economically insignificant" mean? From their regressions, Easterlin and O'Connor (2022) estimate that an extra 1% percentage point of GDP growth rate would take a thousand years to raise life satisfaction by one point on the World Values Survey data, and 500 years to raise life satisfaction by one point on the Gallup World Poll data.

This might be a bit abstract, so imagine we have a country with a starting average life satisfaction of 6/10. Then, given an increase from a 2% to 3% growth rate we can ask, how long is it going to take for the country to go from 6/10 to 7/10? Well, on the World Value Survey, it would take a thousand years and on the Gallup World Poll, it would take 500 years.

If you take Tyler Cowen's view that we only really need to worry about economic growth - because then we'll have so many resources we can make our lives go better and address all of society's other problems - then this is a really serious challenge posed by Easterlin and others who have argued for the paradox. Their evidence is that growth rates *per se* have not, over recent decades, shown any interesting and detectable relationship to average wellbeing.

It's worth pointing out that these figures are the best test we can do for the Easterlin Paradox right now. To test it, we need to look at the longest data sets with the widest number of countries, and the Gallup World Poll and the World Values Survey are those two data sets. This is why I think Easterlin and O'Connor (2022) have the final word on this, at least for now.

3. Explaining the Paradox

At this point, you may be wondering, "Well, okay this is peculiar. Can we explain this?" There are indeed a couple of explanations people have offered, which make the paradox seem less mysterious.

The first explanation draws on the effect of social comparison. The effect of additional money on how we feel about our lives is not just about how wealthy we are in absolute terms, but how wealthy we are compared to other people (Clark et al., 2008; Latif, 2015). It's not that surprising to think that we will never get used to being richer than other people, because it indicates status. We are social beings and sensitive to status.

The other part of the explanation appeals to hedonic adaptation and the fact that we do get used to lots of things (Diener, Lucas, and Scolon, 2006; Graham and Oswald, 2010). To put it bluntly, we are basically naked apes and have evolved to survive without needing many material belongings. As we've gone from having iPhone 5s to iPhone 6s, to iPhone 7s, to iPhone 8s and so on, these have not made a lasting improvement to our wellbeing because we just quickly get used to them. I know economists will find this conclusion almost heretical, but so much the worse for economics.

We do keep noticing the social comparisons - it's annoying to me that you have the fancier iPhone - and so, between these, we can make sense of the fact that we're all getting richer but not really noticing it, because we're paying attention to the relative social comparisons instead.

What's perhaps in tension with this is the finding by De Neve et al. (2018) that life satisfaction does seem to go down during recessions. If we're all getting richer and that doesn't make us happier on average, shouldn't it also make no difference if we're all getting poorer?

I'm not sure there is a tension here. We've already said, over the short run, economic cycles can affect wellbeing: part of the comparison we make is not just to others now, but to our earlier selves. Recessions seem to be worse than growth which Easterlin attributes to the fact that people plan their economic lives around things getting better, so it's suddenly a real crunch point when you aren't able to pay for the thing you've committed to - your mortgage, rent, your car, your children, etc.

Maybe by this point, you are starting to believe that the Easterlin Paradox might be a real thing. Should we be surprised that we were surprised? Well, after a bit of reflection, perhaps not so much.

Research from psychology suggests that we are not so good at predicting how we will feel (Wilson and Gilbert, 2005; Kahneman et al., 2006). These are known as failures of *affective forecasting*.

One of the biases described by Daniel Kahneman is a *focusing illusion*: when we think about other people's lives, there are certain details which are more salient and drive our mental simulations, more so than they really should. Kahneman argues that part of the reason we overestimate the effects of money is because it's so easy to visualise.

Another bias is *immune neglect*, where we don't account for the fact we will adapt to some things, but not others. As an example, imagine how you would feel driving around in your new Ferrari. If you did get a Ferrari, you'd be driving around thinking, "Wow I've got a Ferrari" for the first few weeks, and after that, you'll just be thinking "Oh my god, when is this light going to turn green" like everyone else on the road. That's the difference between how we imagine things are and then how they are once we've adapted to them.

Further evidence that our predictions about happiness are not so hot comes from the figure below. You can ask people in a country "how happy are you, taking all things together?" and then ask "what percentage of people in your country would say the same?" The results are rather divergent. In the US, 90% of people say they are "very happy" or "rather happy," but in fact, US citizens guess that only 50% of their fellow countrymen say they are "very happy" or "rather happy" — quite a large difference between what people say about themselves and what other people say.

Figure 8: The relationship between predicted levels and actual levels of happiness in the US population

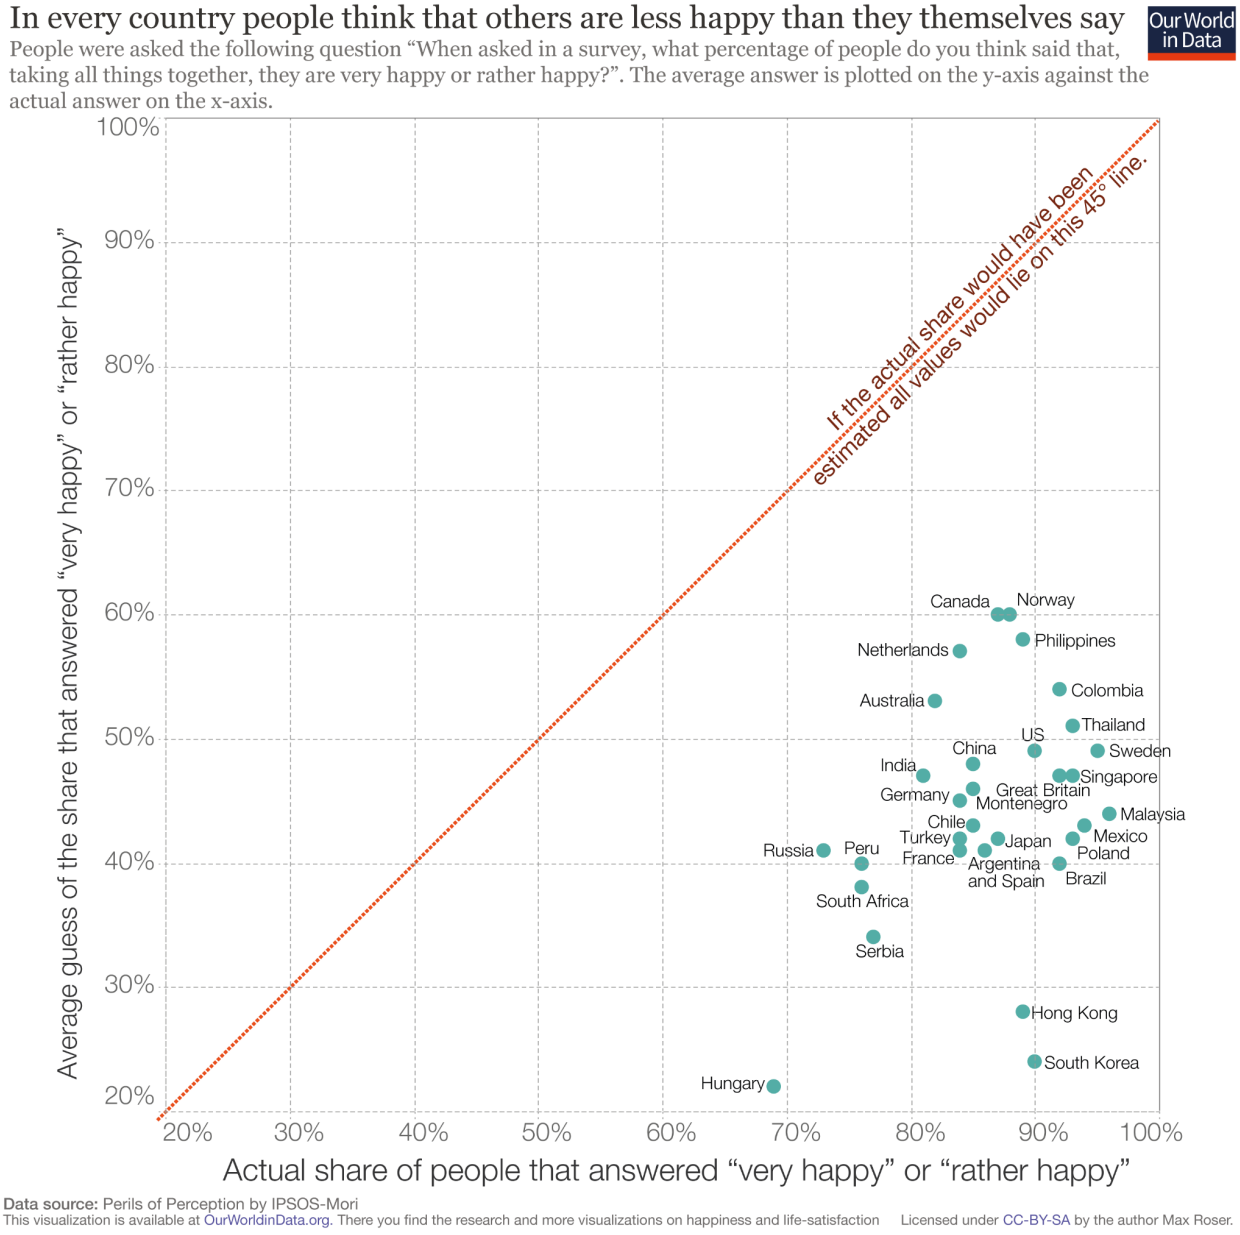

4. Objections

Let's turn now to some criticisms of the Easterlin Paradox. I'll consider four.

4.1 Can we trust subjective wellbeing data?

The paradox is based on subjective wellbeing data: self-reports where we ask people "How satisfied are you with your life?" and so on. But can we trust them? Someone might say something like "We knew these were stupid all along and the Easterlin Paradox proves it."

I understand this as a reaction from those who are not familiar with the ins and outs of subjective data. But it's much too fast. As you might expect, lots of people have worried about the question of whether self-reports are *valid*, that is, whether they are accurate, and it's been investigated in great detail by social scientists over the years (see OECD 2013 for an excellent review). The basic idea is that you don't just automatically believe that your measures are valid, you have to go out into the world and see if they give you believable results - the sort of results they would give you if they really were working.

It turns out that higher levels of happiness and life satisfaction are correlated with: higher income (at a time!), reports from others (if you're happy, your friends will also say you're happy), better health outcomes, being in areas of low pollution, being near green spaces, being in relationships, having a job, and being in low areas of crime. Furthermore, people are more likely to commit suicide if they have lower levels of life satisfaction, and are also more likely to leave jobs or relationships and so on (Diener, Ingelhart and Tay, 2013; Kahneman & Krueger, 2006; OECD, 2013). In other words, it turns out that asking people how happy they are really does seem to tell us how happy they are.

As an aside, the underlying theory in the philosophy of science is *construct validation*, where we test if a measure is valid by seeing if it behaves in the way you expect it to given your best understanding of what you're trying to measure (Alexandrova and Haybron, 2016). It's the same approach which is used in science to test measures of anything, not just self-reports of feelings.

At this point, if we've already accepted that our measures of happiness and life satisfaction work in general, it's very puzzling to say they don't work when it comes to measuring a particular thing. It's a bit like saying your thermometer works in your living room but not your kitchen, or your bathroom scales don't work once you've eaten chocolate cake. At the very least, the sceptic needs to provide a story here for us to believe them.

Perhaps the sceptic thinks, for some reason, we should give up on happiness data altogether - they don't buy the construct validation story. But, why stop here? We ask people for subjective ratings all the time: Uber drivers, restaurants, job satisfaction, mental health diagnoses, pain scores, etc. Are these all nonsense too? Surely not.

4.2 Are subjective wellbeing (SWB) data comparable over time?

A second criticism is that, although we can generally trust SWB data, we can't take the numbers at face value. More technically, the worry is about *cardinal comparability*: if two people said they were 3/10, then report 5/10 later, can we assume they started at the same level of SWB and had the same size increase? The specific version of this criticism that the 'growth-advocate' needs to make is that people are genuinely getting happier and/or more satisfied, but they are changing how they report those feelings over time. For example, perhaps in response to rising economic standards, a score of 8/10 in 1990 might represent a lower level of SWB than the same score of 8/10 in 2020. Same number, different level of happiness.

To address this issue, we need to get quite a long way into the weeds. Those who aren't worried about this can skip ahead to the next objection.

There are quite a few things to say here. Unlike validity, this is not a well-studied topic. When I came to look at it for my PhD, I struggled to find much of a literature on it. There were bits and pieces, but nothing that seemed to convincingly offer an overall assessment of the issue (see Plant 2020 where I try to offer one). Often the only thing researchers will point to is that we should expect *classical measurement error*: if errors are random, they will 'wash out' as random noise. However, we might wonder if errors are non-random: what if different people, or people in different countries, use different scales or they use different scales over time?

We need to distinguish *hedonic adaptation* from *scale shifts.* Hedonic adaptation occurs when you require more and more of something to feel as happy, and can be linked to increasing aspirations. For example, you might need to go to fancier and fancier restaurants to have the same feeling of satisfaction. This is different from actually enjoying the fancier restaurants more and more but consistently rating your enjoyment as 7/10 because your conception of how enjoyable a 10/10 restaurant would be has changed. That's a *scale shift.* For the purpose of testing the Easterlin Paradox, it's not a problem if we're getting used to stuff - indeed, the claim is that the iPhone 10, etc. didn't lastingly make us happier - but it would be a problem if we were actually getting happier but changing our reporting style so that it *looked like* we weren't getting happier.

We can't decisively rule out scale shifts. Obviously, it's in the nature of something subjective that we can't objectively check it. As with assessing validity, we have to reason about what's likely - the fact something is possible doesn't mean it's likely.

In broad terms, I think we should be somewhat reassured about the cardinal comparability of subjective data. As one piece of evidence, see Figure 9 below, taken from a YouGov poll. Here, American individuals were asked to give ratings from 0 (very negative) to 10 (very positive) for different words like: "very bad", "terrible", "outstanding", "excellent", and "perfect."

The 'bumps' represent the proportion of people that give each answer. The overlap isn't perfect, but it's pretty good. If you ask people to score 'perfect', basically everyone says it means 10/10. If people thought this task was meaningless, they'd answer at random, and the lines would be flat. So, it seems that people are sensibly able to compare verbal labels and numerical labels and do so in the same sort of way.

Building on this, the map of life satisfaction scores around the globe we saw earlier in Figure 3 suggests that people are using the same scale. My pet theory on this is that the way we intuitively use 0 to 10 scales is by taking 10 to be the highest realistic level (i.e. the happiest a person can realistically be) and 0 as the lowest (i.e. the least happy a person could realistically be) (Plant 2020). We do this, I claim, so that our answers always fit within the limits and so we *can* use the same scales as other people and over time. If we didn't do this, it would make it very difficult for our answers to be understood. If everyone used bespoke scales which changed from moment to moment, so 7/10 one day meant 2/10 the next, we wouldn't do a very good job of communicating. But the purpose of language is of course to communicate.

Figure 9: Assigning numerical values to words (YouGov)

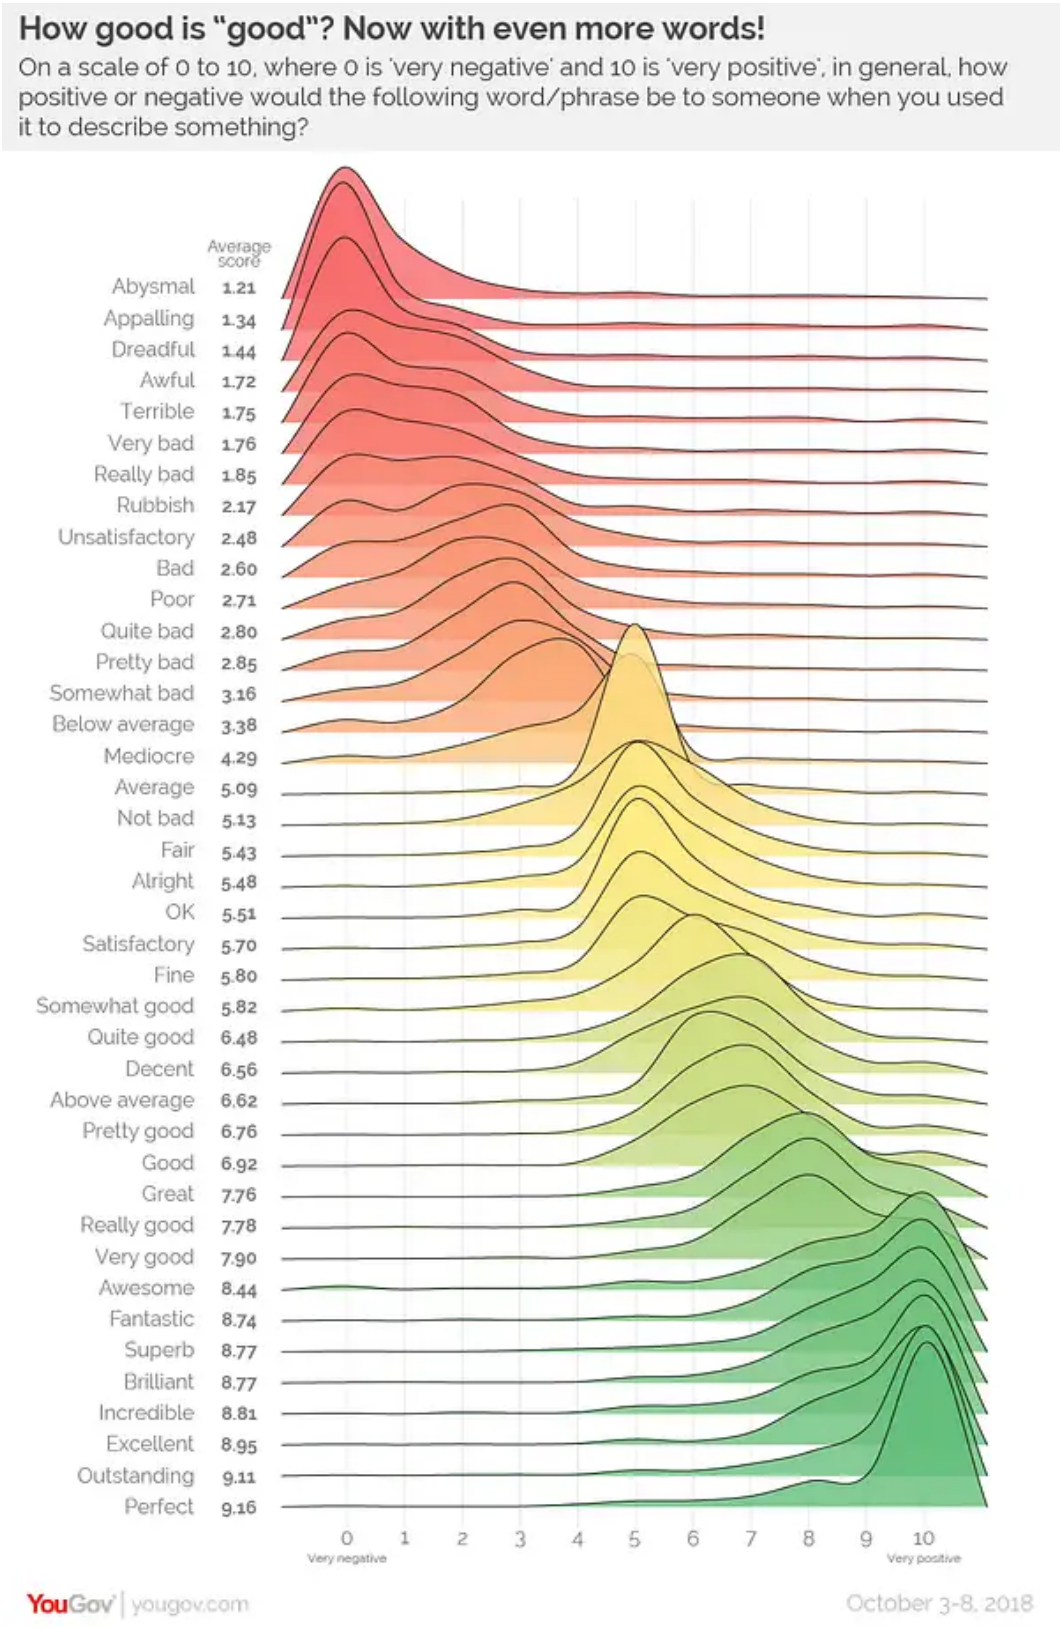

None of this decisively rules out the possibility that *not only* are their rising aspirations and hedonic adaptation - we need more stuff to feel as good - but also that we are changing our scale use over time.[4] There's some evidence from memory data that individuals keep the same scales over their own lives (this is from Prati and Senik 2020, which I discuss in Plant 2020). Scale shift could still happen generationally: an 8/10 to someone born in 1950 represents a lower level of whatever subjective thing is being measured than an 8/10 to someone born in 2000. I can't think of any research that addresses this specific concern. It doesn't strike me as particularly likely, though. It implies that if me, my parents, and my grandparents were each to say we were 10/10 happy, we would assume I would be happier than my parents, who would be happier than their grandparents. Of course, the meaning of particular words changes over time, but we notice this from context. I assume we would have noticed if there were substantial differences in how different ages communicated the intensity of their feelings. If anything, I'm inclined to believe that it would go the other way from what the growth-advocate needs: older people are more reserved, so their 10/10 would be more intense than mine.

An alternative to arguing about possible scale changes would be to take a general theory of how happiness works, how economic growth changes our lives and society, and whether we should expect it to increase or reduce happiness as a result. For my money, the most promising option is to conceive of happiness and unhappiness as "Mother Nature's" reward and punishment mechanisms for evolutionary fitness. In this light, we want to consider humanity's environment of evolutionary adaptation, i.e. from about 100,000 years ago, to the present day but it's not obvious this analysis favours the growth-advocate. Notably, Hidaka (2012) argues that depression is rising as a result of modernity, and points to the fact that "modern populations are increasingly overfed, malnourished, sedentary, sunlight-deficient, sleep-deprived, and socially-isolated".[5]

4.3 There is no paradox

The third criticism is that there is no paradox. This challenge comes from Stevenson and Wolfers (2008) and Sacks, Stevenson and Wolfers (2012) who argue that Easterlin has misinterpreted the data and, on closer examination, growth *is* associated with higher wellbeing.

The reply that Easterlin and O'Connor (2022) make is that Stevenson, Wolfers, and co. are looking over too short a time horizon. They point out that the critique looks at segments of ten years and to

[4] The response Easterlin and O'Connor (2022) offer is that if people were getting happier over time, we'd expect to see 'ceiling effects', where lots of people end up saying they are 10/10 because they've 'run out of room on their own scale'. This doesn't address the concern: there's no need for people to reach the end of the scale to scale shift. I thank the audience of the Global Priorities Insitute's Progress Studies seminar for pressing this point.

[5] A further paradox to resolve would be how to combine the apparent Easterlin result that happiness is rising - but not as a result of growth rates - with the apparent result that depression is rising - as a result of modernity. One way to render these consistent would be if the distribution of happiness had widened over time; then there might be more very unhappy people but an unchanged average. I have not checked this and I am unsure what else to say on this issue.

really test the paradox requires looking over a longer time period, which is what Easterlin and O'Connor (2022) do themselves. Easterlin and O'Connor (2022) write that they don't really understand why Stevenson and Wolfers are using these short time segments rather than the longer ones.

As noted before, I think Easterlin and O'Connor (2022) have the better side of this debate and we really do want to be looking over the longer time period. After all, the longer time period is exactly what we're trying to test and what we find on the longest available datasets indicates a surprisingly non-existent relationship between subjective wellbeing and GDP.

4.4 The Paradox doesn't apply to poor countries

A fourth criticism is that the paradox may apply to rich countries, who have already reached a certain level of wealth, but does not and will not apply to poor countries. This is half an objection, really, as it concedes the paradox does occur in richer countries. The motivation for it is that the richer countries are more satisfied but presumably, they became more satisfied by becoming richer.

This issue is open to debate. The figure below, from Easterlin and O'Connor (2022), shows different data sets looking at subjective wellbeing over time from China, Japan, and India which have all experienced rapid economic growth over different time periods; for China that's 1990 to 2018, or around 30 years. Over that period we can see life satisfaction goes down and then up again. Japan is around a 30 year period as well, and we can see that the relationship is basically flat too. And finally, India in recent time periods has seen life satisfaction go down.

This leaves the question as to how they should be interpreted, which I think is very much open for debate. The answer put forward by Easterlin is that times of enormous economic expansion are accompanied by societal upheaval. People leave villages, where they have good family networks, and go to work in cities, which are often dirty and unfriendly, and where they don't know anyone.

Easterlin's analysis, in the case of China, suggests that the social support systems were taken away and only later do they start to get put back in. That's the human story to all of this — you may experience lots of growth, but that is often less important to people than tears in the social fabric of society.

A key challenge is that there really isn't a lot of data to go on. There are some data on some larger countries which have developed a lot, but there are not lots of data on other less developed and less satisfied countries, and there might not be for another 20 years until we know how and whether they change over time.

Figure 10: Happiness in three formerly poor countries during subsequent periods of rapid economic growth (Easterlin and O'Connor, 2022)

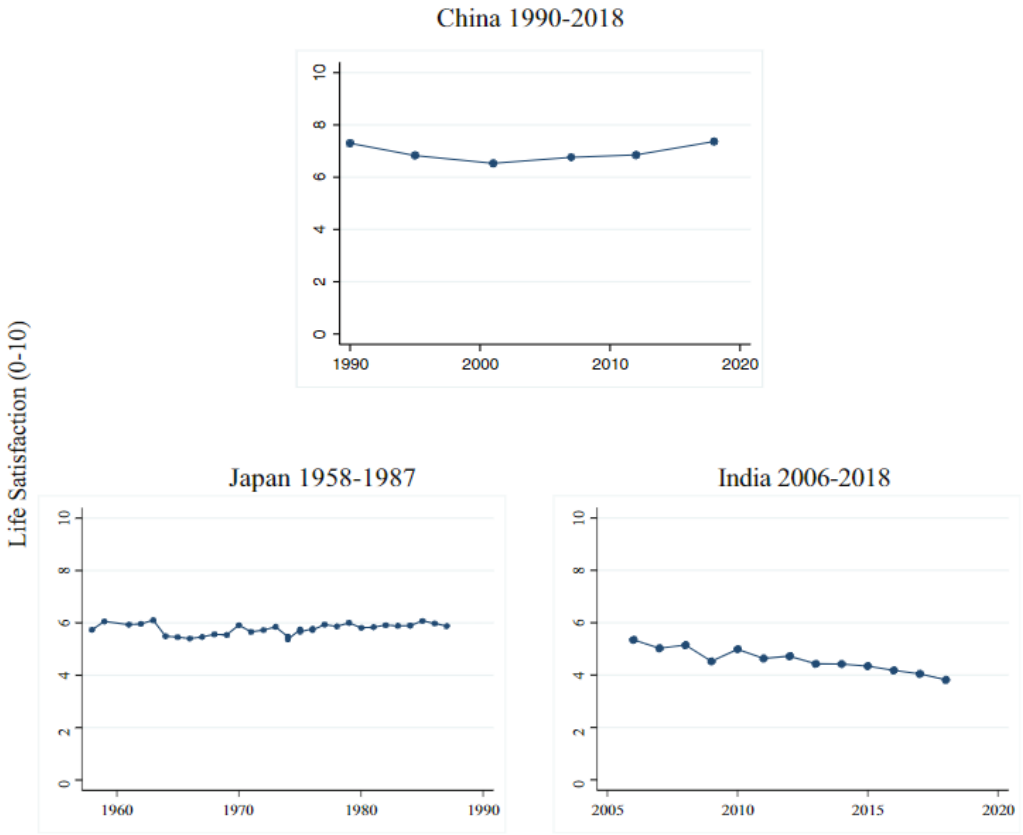

Also in this context, the research by my own organisation, the Happier Lives Institute, finds that cash transfers to the very poor — those on the global poverty line — actually do have a small but significant effect on subjective wellbeing, one that continues over several years (McGuire, Kaiser, Bach-Mortensen, 2022).

So it could be the case that, for the very poor, economic growth does do quite a bit to raise wellbeing, but at a certain point money stops having a significant effect in absolute terms. In some sense, this is the story we all seem to accept: that we do need resources, but only up to a point, and after that point we're just showing off. Hence, we should focus on how society is organised, as opposed to how wealthy it is. The issue is that the empirical evidence on subjective wellbeing hasn't identified exactly where that point is. So we're back to speculating.

5. What can we do to improve the human lot?

My thesis so far has been negative. I've been arguing that the common belief that economic growth will improve the human lot seems to be false.

Is there something more positive to say about how to increase aggregate wellbeing? This, I have to confess, is something of a new question and the growing field of wellbeing science is only starting to bear fruit and identify how our priorities should differ from the status quo. Depending upon your perspective, you can either see that as frustrating or as an exciting opportunity to do valuable academic research.

The Global Happiness Policy Report (2018), says:

> *It must be recognized that in this first Global Happiness Policy Report, written when there are still relatively few examples of rigorously evaluated happiness policies to review, our catalogue will have many tentative entries, as well as many places where the relevant policies remain to be developed and supported with evidence.*

The literature here is small, with notable contributions including Diener et al. (2019), Easterlin (2021), and Frijters and Krekel (2021).

In other words, this is all quite new. If you're trying to directly measure outcomes in terms of people's subjective wellbeing, while many factors are still going to matter, you might get a substantially different story for some issues, such as the value of economic growth. There is clearly lots more work to be done to figure out where our intuitive stories about wellbeing and where our evidence-based stories about wellbeing come together and come apart.

Here are some tentative suggestions. Unsurprisingly, the overarching one is that we should measure outcomes and set priorities using subjective wellbeing. Whatever your view is on what matters morally, presumably people's subjective wellbeing (their happiness and/or life satisfaction) is going to form some part of that story, and you're going to want to account for the best evidence regarding it.

More concretely, in his 2021 book, *An Economist's Lessons on Happiness*, Easterlin suggests that job security, a comprehensive welfare state, getting citizens to be healthy, and encouraging long-term relationships would increase average wellbeing. All of those seem fairly plausible to me.

We should also take mental health and palliative care more seriously — as suggested by the Global Happiness Policy Reports (2018, chapter 3; 2019, chapter 3). We could also consider improved air quality, reduced noise, more green and blue space (blue spaces being water), and getting people to

commute smaller distances (Diener et al. 2019). Social interactions could be enhanced via urban design, reducing corruption, increasing transparency, supporting healthy family relationships, and maybe even things like progressive taxation (Global Happiness Policy Report 2018, chapter 6). These seem sensible but are tentative suggestions. I'm not really committed to any of them as concrete methods of improving wellbeing. These are just to get the juices flowing and indicate there could be a path ahead.

6. Wrapping up

Let's take stock. I take it that the Easterlin Paradox poses a challenge for the view that a laissez-faire approach to policy will lead to progress. If you were inclined to the view, "let's just focus on economic growth and forget about everything else", I think you should be pretty sceptical that this will improve the average person's wellbeing.

That's not to say that we can't improve people's wellbeing, but economic growth alone doesn't seem to do that much and we ought to be thinking about which particular interventions are most cost-effective.

Of course, more or less economic growth might be relevant for other sorts of global priorities. If you're taking a longtermist perspective you might think that a far larger economic pie is better for engaging in activities like space colonisation. On the other hand, you might wonder if a faster rate of economic growth is more likely to lead to existential catastrophes, for example by speeding up the takeoff of unsafe artificial intelligence. The Easterlin Paradox is just referring to average wellbeing, rather than these other sorts of concerns.

It's also worth noting that the Easterlin Paradox seems to apply more clearly to richer countries. It still may apply to poor countries, but the evidence is less decisive for them (see Objection 4).

Finally, a wellbeing policy agenda based on subjective wellbeing data rather than our intuitions about what *we think* would improve people's lives, is only really starting. In years to come, perhaps we will have more things to say about how we can achieve progress in terms of our collective quality of life. This is a topic which I invite people who are interested in Progress Studies to engage with.