How to address sustainability risk in a dangerous universe

While our galaxy is full of planets that could harbour life, there is no evidence that any of them besides our own actually do. This observation suggests that the universe is a dangerous place and that we may be facing civilisational risks that could, in a worst case outcome, result in extinction. Sustainability policy (broadly conceived) deals with these risks. It follows that the goal of sustainability policy is to ensure that we don't drop the baton of civilisation during our leg of the relay race of human history (Llavador, Roemer, and Silvestre 2015 and Cowen 2018 advocate a similar approach).

Global warming is the paradigmatic example of a sustainability risk, and much of our thinking on how to deal with such risks has emerged from the global warming debate. Yet, reflecting upon how this debate has gone suggests that: i) the cost-benefit analysis (CBA) framework that we now use to discuss sustainability policies has serious flaws; and ii) these flaws distort policy analysis. The flaw in the CBA framework (in this context) is that it combines the question of whether to address a given sustainability risk and the question of how to address that risk into a single analysis. Consequently, factors that affect whether to act have massive—but not necessarily sensible—implications for how to act. A framework that explicitly separates those two questions will therefore lead to better policy analysis and better policies.

A better framework for sustainability policy

Dealing with sustainability risks such as those arising from climate change require enormous investments now to create benefits in the far and far far future. CBA would therefore seem to be the natural framework for thinking about whether and how to respond to such risks. In this framework, one evaluates policy options (including doing nothing) to address a risk by: i) using a discount rate to convert that policy's future costs and benefits into a present discounted value now (for example, with a discount rate of 6%, $1 one year from now is worth $0.94 today); and ii) choosing the option with the highest net benefit.

In the context of sustainability policy, however, CBA suffers from a fatal flaw: there is no discount rate that yields both a sensible policy outcome and a sensible balance of intergenerational costs and benefits. That is, if one uses empirically grounded market discount rates that provide a sensible trade-off between the present and the distant future from our perspective, then the discounted value of avoiding even horrendous consequences in the far future is too low to justify much in the way of action now. So, we will underinvest in preventing human extinction (this is a bad policy outcome). To avoid that outcome, one can use the hammer of necessity to smash down market discount rates to a societal discount rate that is low enough to justify action now. But, in this case, the present discounted value of future benefits implied by this low discount rate is so high that maximising net benefits means that we end up in the ridiculous position of running policy now in the interest of the eventual people of the far far future.

A simple example suffices to illuminate this conundrum. For individual decision-making, market data implies that a sensible discount rate is about 6% (maybe not—as I write this—in the middle of a global pandemic, but in equilibrium once we recover). Using this rate, a benefit stream of $1 a year (for ever) starting at the human scale future of 25 years from now (say for retirement) is worth $3.55 today. But, that benefit stream starting 100 years from now (that is, on time-scales needed to deal with sustainability risks) is worth only $0.035 today. By shifting out the start of the future benefit stream from the personal scale to the intergenerational scale, the value of that benefit stream falls by 99%. Obviously, it will be difficult to motivate people now to deal with intergenerational sustainability risks using market discount rates.

Confronting this dilemma, the highly influential Stern Review's analysis of climate change policy applied a combination of economically rigorous argument and vigorous hand-waving (see here and here) to push the individually sensible discount rate of 6% down to a societal discount rate of 1.4%. Using this discount rate, a $1 a year benefit stream starting in 100 years is worth $17.45 today (500 times higher than it would with the market discount rate). So, a low societal discount rate will definitely push us towards a more future-focused policy outcome.

Indeed, it will inevitably push us too far. With a 1.4% discount rate, a $1 a year income stream starting 200 years from now is worth $4.26 today. This is 60,000 times higher than the value an individual would put upon it and is in fact a higher value than an individual would put upon that same income stream for their own retirement. Societal discount rates therefore force us to place a ludicrous weight upon the utility of eventual people 200 or more years into the future (people of whom we know nothing) when making policy decisions now.

The solution to this conundrum is to separate out the decision of whether to act on a sustainability risk from the decision of how to deal with that risk if action is taken.

In my view, much of the motivation to deal with sustainability risks arises from a moral intuition that we have an ethical obligation to pass on a thriving civilisation to future generations. In this case, this obligation binds now and is independent of any discount rate calculation. And if we don't have to conjure up a low societal discount rate to justify taking action, then we can avoid the pitfall of placing too much weight upon the far far future when deciding what to do.

While departing from an explicitly ethical premise seems different from the "let's just add things up" premise of the CBA framework (where the ethics are in reality just hidden in the math), it is much more in line with how people usually take decisions that have an important ethical component. For example, suppose that a friend asks your advice on whether they should murder someone who is annoying them. You could use the CBA framework to think this through: on the one hand, you do get an immediate benefit; on the other hand, you also run the risk of incurring a stream of costs in the future; so, my advice to you is to be careful to use the correct discount rate to ensure that you don't murder that annoying person in error. But that is not how most people would approach this matter (based upon my admittedly informal survey).

The more standard response would be to say that murder (and other such actions) are off the table of ethically acceptable options independent of the discount rate, and that you should choose the best option of those that remain (taking discount rates into account).

Applied to sustainability policy, this approach would begin with an explicit discussion of what sort of world we wish to leave to future generations and the cost of achieving various possible outcomes (using market discount rates to choose the minimum cost method to achieve a given outcome). Making this discussion explicit rather than having it implicitly in the form of a debate over discount rates, risk aversion parameters, utility function specifications, etc., etc., will be more honest and open, which is a good thing all by itself. It is also more likely to get the support needed to implement sustainability policies given their inevitably substantial costs (and it's not like the current approach has worked in that respect).

Implications for policy

This approach will have significant implications for sustainability policy. Let me briefly consider some of the more important implications here.

To begin with climate change, this approach will tilt the policy response towards adaption (that is, coping with a warmer planet) and away from prevention (keeping global temperature at its current level). Intuitively, the benefits of prevention go to future generations but the costs are borne by current generations. It follows that a policy built upon the assumption of a low discount rate will place an enormous weight upon those future benefits and so tilt towards prevention. But, a policy built upon limiting sustainability risks to an acceptable level directly (without getting to that objective by imposing a low discount rate) will put less weight on far future benefits and more weight on current costs, and so will tilt towards adaption.

A shift towards adaption will naturally lead to a much greater emphasis on promoting economic growth and technological development. Wealthier and more technologically advanced societies are obviously in a better position to adapt to the ramifications of climate change. Furthermore, technological advances and economic growth will also create capabilities and resources that can be used to deal with other extinction risks.

Human civilisation faces sustainability risks aside from those stemming from climate change (nuclear war, giant asteroids, deadly epidemics, etc.). And these risks are far more severe because the earth is (at the moment) a single point of failure. Any event that poses an extinction risk to civilisation on earth is also an extinction risk to civilisation full stop. The ultimate goal of sustainability policy must then be to eliminate this single point of failure risk by creating a self-sustaining human civilisation in space. Over the time-scales we are already dealing with in regards to climate change, this is definitely an achievable goal.

So, the optimal sustainability policy will (in very general terms) entail; i) reducing carbon emissions to maintain an acceptable level of environmental quality and to reduce the tail risk of a climate disaster that imperils civilization (while promoting research to better understand those tail risks); ii) a tilt towards dealing with climate change via adaption rather than prevention; iii) an emphasis on promoting economic growth and technological development; and iv) supporting a vibrant space program.

While no country is designing its sustainability policy from the perspective I advocate here now, some countries are closer to this optimum than others. In particular, the US under the Trump administration comes closest to an optimal sustainability policy. The US is reducing emissions almost as rapidly as the EU while doing more to promote economic growth and technological advances. Crucially, the Trump administration is also supporting the development of a vibrant private sector space industry and building the legal foundations for the exploitation of space resources.

Conclusion

The universe is a dangerous place. Any pre-technological civilisation will inevitably be destroyed by natural exogenous single point of failure extinction risks (the dinosaurs, for example, famously lacked an effective space program that could protect them from giant asteroid risk). It follows that in order to survive over the long run a civilisation—including our civilisation—needs to get off planet. But to get to the space-faring civilisation stage we need to successfully pass through our current technological stage, a stage in which we face not only exogenous single point of failure risks but also endogenous single point of failure risks such as those arising from climate change. A sound approach to sustainability policy will help us to successfully navigate through this perilous time in human history.

♣♣♣

Notes:

- *This blog post expresses the views of its author(s), not the position of LSE Business Review or the London School of Economics.*
- *Featured image by William Bossen on Unsplash*
- *When you leave a comment, you're agreeing to our Comment Policy*

Kevin R. James is an economist in LSE's Systemic Risk Centre. His research focuses upon financial market efficiency and effectiveness, systemic risk, corporate finance, and asset management. (@kevinrogerjames)