Discounting and Doom

1

In *The Precipice*, philosopher Toby Ord argues that safeguarding humanity's long-term future by reducing the probability of catastrophic risks, such as extinction, is among the most pressing moral issues of our time. A significant motivation for Ord's thesis is that future persons–even those who might exist millions of years from now–merit the same moral consideration as people alive today. Just as it is (putatively) immoral to discount the well-being of people geographically distant from us, it is immoral to discount the well-being of those temporally distant from us.

In contrast, when government agencies perform cost benefit analyses of potential policies, they typically discount impacts that will accrue in the far future, if such impacts are considered at all.[1] This essay examines the implications of Ord's argument on temporal discounting in policy analysis. **2** introduces the argument. **3** outlines three reasons why policymakers may

[1]Some examples: The UK government uses an annual discount rate of 3.5%, based on the Ramsey equation (HM Treasury, 2011). The New Zealand Treasury recommends a 5% annual discount rate for most government projects, and a 6% discount rate for others (NZ Treasury 2020). The European Commission uses a 3% rate (Sartori et al. 2014).

(justifiably) be reluctant to consider the far future in quantitative analyses, and suggests ways to overcome this reluctance. **4** concludes.

2

In policy analysis, the temporal discount rate reflects the reduced value of future impacts compared to near-term ones. For example, a 5% annual exponential discount rate implies that some benefit of a policy that will accrue a year from now is worth only 95% of that benefit if it were to accrue today. This is clearly in tension with longtermist philosophy–at an annual discount rate of 1.7%, a policy's impacts within the next 40 years are valued just as highly as its impacts from 40 years onwards in perpetuity.[2] The practical effect of even seemingly low discount rates is that the interests of nearly all future people are entirely excluded from cost benefit analyses.

Ord employs a popular economic model to contrast the standard approach to temporal discounting with his longtermist perspective. The **Ramsey model** explicitly accounts for two common justifications for temporal discounting (Ramsey 1928). First, we would rather increase the consumption of a poorer person by some set amount, than increase the consumption of a richer person by that same amount. This is because there are diminishing marginal returns to consumption. If we expect the median future person to be richer than us today, as a result of economic growth, then we should similarly discount any policies that increase their consumption. Second, people exhibit **pure time**

[2] 1.7% is a rate recommended in the 2023 United States OMB draft Circular A-4. The real solution of $\int_0^a (1-0.017)^x \, dx = \int_a^\infty (1-0.017)^x \, dx$ is $a \approx 40$.

preference, i.e. a preference for benefits that accrue sooner, independently of any economic considerations as discussed above.

These two factors are respectively expressed in the two terms of the Ramsey equation, where ρ is the social discount rate, determined by ηg and δ, representing the two justifications explained in the previous paragraph.

$$\rho = \eta g + \delta$$

As Ord emphasizes, the first term doesn't apply to catastrophic risk reduction measures, because their key benefit isn't consumption increase for future, possibly wealthier, generations; these measures primarily ensure that future generations will exist at all. Ord, like most other longtermists, also rejects the second term, on the basis that pure time preferences display an arbitrary disregard for the well-being of future people.[3] Even if people actually display pure time preferences, these preferences *ought* not be considered in cost benefit analyses. Instead, we should discount future benefits only by the **catastrophe rate**, i.e., the year-on-year risk of extinction. For instance, if there is a 0.1% chance of extinction in the coming year, then effects accruing next year should be discounted by 0.1%, and so on.

If longtermists are correct, how should policy analysis change as a result? Analysts should favor methods placing more weight on the far future, like discounting by a (dynamic) catastrophe rate. Additionally, Ord's stance implies that policy analysts should consider long time periods over which a

[3]For an example of a detailed ethical argument against pure time preferences, see MacAskill and Greaves 2019.

policy might have effects. The recommendation from the United States' 2023 draft Circular A-4 is that analysts should select an ending point such that, "to the extent feasible", all relevant costs and benefits stemming from the policy are considered (Office of Management and Budget 2023, 74). From *The Precipice*, we learn that the appropriate ending point for analysis of policies affecting catastrophic risks is likely *very* far in the future. This is because, if humanity's potential is reached, the resultant well-being gain is many orders of magnitude larger than any effects accruing in the next century. Moreover, if a future civilisation is more-or-less robust to catastrophic risks, the catastrophe rate should decline back to an annual 'background rate' much lower than the current, heightened risk of catastrophe, further increasing the weight placed on catastrophic risk reduction as a policy impact.[4]

3

Despite objections to high temporal discounting rates, policymakers may hesitate to lower them for at least three reasons. The first has been discussed previously by Tyler Cohen and Derek Parfit (1992), who offer a refutation. The second is an obvious pragmatic issue, and the third was coined by Smith and Winkler (2006), although not in the context of longtermism.

1. **Respect for constituent preferences:** Policymakers may feel obliged to respect constituents' pure time preferences, viewed as subjectively irrational or not.

[4]See Appendix A of *The Precipice*.

2. **Accountability and legibility of analysis:** Since estimates of long-term impacts tend to be more uncertain, a lower discount rate may make it easier for analysts to introduce their own biases into their estimates, while simultaneously making it very difficult to refute the subjective probabilities used. This might lead to faulty analyses, and hence suboptimal policy.

3. **Post-decision surprise:** Assume that estimates of the net benefits of multiple policies are *unbiased*. If a government selects the policy with the highest estimated value, we expect the value of the chosen policy to be systematically overestimated. The magnitude of this upwards bias increases in the variance of the cost benefit estimate. In practice, this means that the estimated value of highly uncertain policies, such as catastrophic risk reduction, should be adjusted downwards, even if the government were risk neutral.[5]

Reluctance to consider the far future in policy analysis is not merely a matter of philosophical disagreement over the worth of future people. There are moral and epistemic reasons to be wary of the conclusions drawn by Ord. It seems to me that respect for constituent preferences is not easily dispelled by philosophical rebuttals. Two promising alternatives are a) arguing that the public does not have a pure social time preference (or, at least, a time preference faithfully represented by exponential discount rates), or b) directly

[5]The explanation as to *why* this is the case is beyond the scope of this essay. But suffice it to say that, especially when dealing with highly uncertain impacts, non-obvious issues may adversely affect attempts at quantitative cost benefit analysis.

convincing the public to reject pure time preferences *in the policy analysis context*.

The second and third reasons are concerns about estimating highly uncertain future impacts. Fortunately, the treatment of uncertainty is already a topic of interest among many longtermists. The community can devise suitable estimation methods for highly uncertain impacts and help to train analysts accordingly. Some work that might be useful includes: a) developing easy-to-implement standards for uncertainty analysis, e.g. Excel templates, b) surveying analysts to identify their actual concerns with evaluating impacts to catastrophic risk, and c) lobbying for precise official guidelines and/or recommendations on analysing speculative impacts, including outside of the United States.

4

The Precipice advocates for governments to reevaluate current temporal discounting approaches, emphasizing catastrophic risk reduction. I have argued that this will require longtermists to appreciate and reckon with the factors currently preventing a policy's effects over very long time horizons from being granted due consideration.

References

[1] Canadian Treasury Board. "Canadian Cost-Benefit Analysis Guide: Regulatory Proposal". In: *Ottawa, Canada: Canadian Treasury Board* (2007).

[2] Tyler Cowen, Derek Parfit, et al. "Against the social discount rate". In: *Justice between age groups and generations* 144 (1992), p. 145.

[3] Hilary Greaves and William MacAskill. "The case for strong longtermism". In: *GPI Working* (2019).

[4] Office of Management and United States Government Budget. *Circular A-4*. Draft. 2023.

[5] New Zealand Treasury. *Discount Rates*. https://www.treasury.govt.nz/information-and-services/state-sector-leadership/guidance/financial-reporting-policies-and-guidance/discount-rates. Accessed on 16th June 2023. 2020.

[6] Toby Ord. *The precipice: Existential risk and the future of humanity*. Hachette Books, 2020.

[7] Frank Plumpton Ramsey. "A mathematical theory of saving". In: *The economic journal* 38.152 (1928), pp. 543–559.

[8] Davide Sartori et al. "Guide to cost-benefit analysis of investment projects. Economic appraisal tool for cohesion policy 2014-2020". In: (2014).

[9] James E Smith and Robert L Winkler. "The optimizer's curse: Skepticism and postdecision surprise in decision analysis". In: *Management Science* 52.3 (2006), pp. 311–322.

[10] Her Majesty's Treasury. *The magenta book: guidance for evaluation*. HM Treasury, 2011.