How cost-effective are efforts to detect near-Earth-objects?

Toby Newberry

Global Priorities Institute | May 2021

GPI Technical Report No . T1-2021

GLOBAL PRIORITIES INSTITUTE

UNIVERSITY OF OXFORD

How cost-effective are efforts to detect near-Earth-objects?[1*]

1. Introduction

Near-Earth-objects (NEOs) include asteroids and comets with orbits that bring them into close proximity with Earth. NEOs are well-known to have impacted Earth in the past, sometimes to catastrophic effect.[2] Over the past few decades, humanity has taken steps to detect any NEOs on impact trajectories, and, in doing so, we have significantly improved our estimate of the risk that an impact will occur over the next century. This report estimates the cost-effectiveness of such detection efforts. The remainder of this section sets out the context of the report, as well as its motivation. Section 2 explains the methodology used for making estimates. Section 3 estimates the cost-effectiveness of past NEO detection efforts, and makes some comments on the prospects for future efforts.

NEOs fall into two basic categories: asteroids and comets. Asteroids are essentially lumps of rock, and typically orbit somewhere between Mars and Jupiter (in the region known as 'the asteroid belt'). While they range in diameter from one or two metres up to hundreds of kilometres, smaller asteroids are far more common than larger ones, and there are only a handful with diameters greater than 10km. The total population is in the millions, with some 20,000 or so having orbits that qualify as 'near Earth' (JPL, 2019). Almost all NEO detection efforts so far have focused on asteroids.

Comets are mixed lumps of rock and ice. They typically follow highly elliptical orbits, which extend into the outer planets, and only rarely come close to the Sun. Comets tend to be larger than asteroids, and most have diameters in the hundreds of metres - although, once again, there are very few larger than 10km. The total population of comets is more difficult to assess than asteroids, but scientists have identified close to 800 near-Earth comets to date.

The destructive potential of a given NEO depends on its mass, composition, speed relative to Earth, and material details of the site where it impacts (or whether it explodes in the air). Unsurprisingly, an impactor that is larger and faster will tend to cause greater destruction. The composition of both the NEO itself, and the impact site, have implications in terms of an impactor's climatic effects. From the perspective of this analysis, however, it makes sense to categorise NEOs by mass (as approximated by diameter) in the first instance. Of the various features mentioned, mass is the best proxy for destructive potential: the other features explain less of the variation in destructive potential and/or are harder to discern in advance than an NEO's diameter.[3]

This report builds directly on work by Jason Matheny (2007). As part of a broader inquiry into strategies for reducing extinction risk, Matheny (2007) estimates the cost-effectiveness of a hypothetical plan to

[1*] I am grateful to Joao Fabiano for his assistance in fact-checking this report. I am also grateful to Will Macaskill, Toby Ord, Carl Shulman, Matthew van der Merwe, and Anders Sandberg for comments on an earlier version. Nevertheless, this is still a working draft, and likely contains errors (for which I take full responsibility). Please send helpful comments to tobias.newberry@philosophy.ox.ac.uk.

[2] Most famously, the Chicxulub impactor caused the extinction of around 75% of species, including all non-avian dinosaurs.

[3] This broad characterisation of different NEOs, as well as the factors that inform their destructive potential, comes primarily from Ord (2020), informed by Stokes et al. (2017).

detect and (if necessary) deflect all near-Earth asteroids over the next century. The aim here is to update and amend Matheny's estimate in three ways. First, this report incorporates more recent estimates for some of the relevant data, including the risk of acatastrophic or existential impact over the next century. Second, this report focuses primarily on estimating the cost-effectiveness of real-world detection efforts over the past few decades, as opposed to hypothetical future proposals. Third, this report addresses considerations of NEO-deflection (rather than detection) differently to Matheny. Where Matheny incorporates the costs and benefits of deflection into his central estimate, this report focuses almost exclusively on detection. We treat the prospect of deflecting or otherwise mitigating the negative effects of an Earth-bound NEO as exogenous to the central estimate.

The human future has the potential to be unimaginably vast, and, as a result, is plausibly of enormous moral importance. At the same time, it may seem unlikely that anything we might do in the present day could reasonably be expected to have an effect over the very long term. The case of NEO impact provides a counterexample to scepticism of this sort. By demonstrating that our efforts in this area have had, and may continue to have, a robust and positive effect over the future of humanity, this report makes the case that affecting the long-term future in predictable ways, and towards favourable ends, is well within our power.

2. Methodology

Cost-effectiveness estimates depend on two pieces of information: the cost of what is being estimated, and the benefit it provides. For NEO-detection efforts, the relevant cost is just the time and money spent on scientific equipment, personnel, and so forth. The relevant benefit is more difficult to make precise, but can be glossed as 'the expected value of learning the results of a given detection programme'. Here, we use the tools of expected utility theory, together with relevant empirical evidence, to inform an estimate of the same.

We make two versions of this estimate, using two different assumptions - where the assumptions differ in whether or not NEOs are taken to present a non-negligible risk of human extinction. According to the 'catastrophic impact assumption', NEOs present a risk of catastrophic loss of life, as well as many other negative effects, but do not genuinely threaten human extinction. In this case, the amount we would be willing to pay for NEO detection efforts depends primarily on the number of lives we expect to be lost in the immediate aftermath of impact, and in the years to decades that follow. Following NASA (1992), this report defines a 'catastrophic impact' as the impact of an NEO with diameter greater than 1km, but less than 10km.[4]

According to the 'existential impact assumption', NEOs that are sufficiently large present a non-negligible risk of causing human extinction outright. In this case, the amount we would be willing to pay for NEO detection efforts depends primarily on the expected size of the entire human future. Even though an impact that could cause human extinction would also have catastrophic shorter-term effects, the expected size of the future is so large that it ends up dominating these in the estimate. This report defines an

[4] Of course, under this assumption an impactor of diameter10km or greater would also have catastrophic (but not existential) effects. The reason we restrict this category to the 1-10km range is largely pragmatic, since the best recent estimates of impact risk use this categorisation (e.g. Ord, 2020). At the same time, including 10km or greater NEOs here would not affect the estimate very much, since NEOs of this size are extremely rare, and, on the catastrophic impact assumption, would not have effects that are all that more significant than those in the 1-10km range.

'existential impact' as the impact of an NEO with diameter 10km or greater (for similar definitions see, e.g. Ord, 2020; Matheny, 2007).

The estimates below that correspond to these two assumptions should be interpreted as follows. For the catastrophic impact assumption, there is a clear scientific consensus that a 1-10km impactor would have truly devastating effects. Because of this consensus, the relevant estimate should be interpreted as a robust lower bound on cost-effectiveness: everyone agrees that impactors of 1km diameter or greater would at least cause widespread destruction and significant loss of life. For the existential impact assumption, there is no such consensus. Despite the salience of asteroids as a vector for human extinction in popular culture, there are reasons to think that even a very large impactor would not mean the loss of the entire human future: humanity thrives in an extraordinarily diverse set of ecosystems, and has made numerous technological advances that might aid in surviving a major impact event. As a result, the relevant estimate should be interpreted only as a tentative upper bound on cost-effectiveness: NEO detection efforts *could* be this cost-effective, but only if we think NEOs really do pose a non-negligible extinction risk. Here, we assume that an impactor of diameter 10km or greater has at least a 1% chance of causing human extinction.

3. Estimates

3.1 What we have achieved

Since the mid 1990s, international spaceguard programmes have been working to track NEOs, with the aim of identifying any on impact trajectories. To date, scientists working as part of this collaboration have tracked over 95% of asteroids of diameter 1km or morein near-Earth orbit, including, with high likelihood, all asteroids of diameter 10km or more. Here, we estimate the cost-effectiveness of the programme under each of the two assumptions stated earlier. The total cost of the programme was estimated at around $70 million (USD) in 2013 (Mainzer et al., 2011; U.S. House of Representatives (2013)).[5] A reasonable estimate of the total cost to date, incorporating more recent figures, is around $600 million (USD) (Dreier, 2019).

To estimate the benefit provided by the programme under each of our two assumptions, we follow an identical set of steps. First, we ask the high-level question: how much would we be willing to pay to learn the results of the programme? This question can be answered using the tools of expected utility theory, where the programme effectively has two possible outcomes: either it finds an asteroid on an impact trajectory, or it doesn't.[6] Second, we estimate the probabilities of these two outcomes, using available empirical evidence, but excluding evidence provided by the programme itself. Third, we estimate the utility values of the two outcomes, as determined by a similar set of empirical claims. Together, these latter two estimates allow us to answer the high-level question, which, when combined with the cost estimate given in the previous paragraph, gives an estimate of the overall cost-effectiveness.

Under the catastrophic impact assumption, the probability of the programme finding an asteroid on an impact trajectory is just the risk of catastrophic impact over the next century, excluding evidence from the programme itself. This figure is relatively well-characterised as being around 1 in 5000, which is both the

[5] Note that this figure may not capture all relevant costs. The subsequent revision of this report will include a closer investigation of the cost-estimate.
[6] We ignore the possibility that it finds more than 1 such NEO, since this is highly unlikely, and, in any case, would have minimal effect on the final estimate.

number that NASA used in its own calculations prior to significant NEO detection efforts (see, e.g. NASA, 1992), as well as our current best-guess at the average risk per century (Ord, 2020). It derives from the historical track-record of impacts of differentsizes, and does not incorporate information about the next century in particular. Similarly, the equivalent probability under the existential impact assumption is just the risk of a 10km (or larger) NEO colliding with Earth over the next century, which is also well characterised at around 1 in 1.5 million (Ord, 2020). The probability of the programme finding no such Earth-bound NEO, for each assumption, is just the relevant complement.

The utility values of the different outcomes, undereach assumption, can be estimated using two pieces of information: the expected damage caused by the relevant category of impactor, and the extent to which detecting an Earth-bound impactor in advance would militate against that damage. For example, one might think that an impactor in the 1-10km range would cause an expected 2 billion deaths, but that advance detection would give us a 50% chance of deflectingthe impactor, or otherwise reducing the expected death toll to near-zero.[7] In this example, the value of the outcome where we detect the asteroid in advance is very high, since it saves 1 billion lives in expectation. By contrast, the value of theother possible outcome, where no Earth-bound asteroid is found, is close to zero under both assumptions (at least as far as the effects of impact are concerned).

Based on the above, we can express the benefit provided by past NEO detection efforts as follows:

$$Benefit = P \times U$$

Where *P* is the probability of an impact of the relevantkind taking place in the next century, and *U* is the utility of detecting this impact in advance. Thisis just a standard expectation value, where the terms corresponding to the second possible outcome drop out (since the associated utility is zero). We can further express *U* as the product of the expected damagedue to impact *D*, and the chance of averting this damage due to advance detection *A*. The overall cost-effectiveness (CE), in dollars per life saved, is then given by:

$$CE = \frac{Cost}{P \times D \times A}$$

As noted, the total cost of the programme has beenestimated at $600 million (USD), and the relevant values for *P* under each assumption are stated above. It remains to fill in values for the expected damage due to impact, and the chance of averting this damagedue to advance detection.

In their 1992 report on the Spaceguard Survey, NASA estimated that an impactor in the 1-10km range could cause the deaths of one quarter of Earth's population.Depending on when in the next century the impact occurs (dated from 1992), this is one quarter of somewhere between 5 and 10 billion people. On the assumption that an impact is equally likely to occur in any given year, we can approximate this as about 1.9 billion people (¼ of 7.5 billion). Thiscan then serve as our estimate of *D* under the catastrophic impact assumption - though it is worth noting that the figure is highly uncertain, that 'expected deaths' does not capture all the effects of impact, and thatusing a point estimate, rather than a probability distribution, is also a simplification. More generally, the correlation between an NEO's diameter and its destructive potential is somewhat noisy, meaning that any claim of the form 'an impactor of diameter X would cause the deaths of Y people' should be understoodas illustrative, rather than authoritative.

[7] This might include constructing extensive safety bunkers, evacuating the areas likely to suffer the worst effects, and so on.

Under the existential impact assumption, the expected damage due to impact is 1% of the entire human future, since the existential impact assumption involves the claim that an impactor of diameter 10km or greater has at least a 1% chance of causing human extinction. If we further assume, conservatively, that humanity remains Earth-bound over the long-run, this is 1% of somewhere in the vicinity of 10^{14} lives (Newberry, 2021). If we instead assume that humanity's future is bounded by the size of the Solar System, the relevant figure is 1% of around 10^{27} lives.[8]

The chance of averting the damage due to impact A is determined by factors like the viability of deflecting an NEO, surviving an impact winter, and other strategies for enduring an impact event. Rather than entering into a detailed discussion of these factors, this report simply assumes that this chance is non-negligible, and uses a place-holder value of at least 5%, or 1 in 20, in the estimates below. The reasons for taking a light touch to this question are (1) that the evidential base for precise claims about A is relatively poor[9], and (2) that questions of deflection, in particular, come with attendant concerns about information hazards (see, for example, Ord, 2020, p68).

The table below summarises the key information from the preceding paragraphs:

	Catastrophic Impact Assumption	**Existential Impact Assumption**
Cost of detection efforts	\$600 million (USD)	\$600 million (USD)
Probability of impact (*P*)	1 in 5000	1 in 1.5 million
Damage due to impact (*D*)	1.9 billion lives lost in expectation	10^{12} lives lost in expectation (Earthbound scenario) 10^{25} lives lost in expectation (Solar System-bound scenario)
Chance of averting impact due to advance detection (*A*)	1 in 20	1 in 20

We can now calculate the cost-effectiveness of the programme under each of our two assumptions, using the formula given earlier. Under the catastrophic impact assumption, the programme involves paying \$600 million for a 1 in 5000 chance of a 1 in 20 chance of saving 1.9 billion lives in expectation. This is a cost-effectiveness of approximately \$31,600 (USD) per life saved in expectation. Under the existential impact assumption, the programme involves paying \$600 million for a 1 in 1.5 million chance of a 1 in 20 chance of saving at least 10^{12} lives in expectation. This is a cost-effectiveness of approximately \$18,000 (USD) per life saved in expectation. If we instead use the assumption that humanity's future will be bounded by the size of the Solar System, rather than just the Earth, then an investment of just \$1 (USD) could be expected to save around 500 million future lives. This number may seem astonishing, but reflects the genuinely astronomical scale of what is at stake: it should not be especially surprising that protecting the entire future of our species turns out to be a valuable investment.

For comparison, it is common for government programmes to value individual life-*years* at more than \$100,000, and put the value of a statistical life at around \$9 million. Moreover, it costs only \$3,500 or so to save a life by donating to the most effective traditional charities. This suggests that NEO detection efforts

[8] This is also conservative, in that it ignores the prospects of interstellar or intergalactic settlement, as well as digital persons. Each of these would dramatically increase the expected size of the future (see Newberry, 2021).
[9] E.g. no impact event has ever successfully been averted.

have been considerably more cost-effective than standard government evaluations, and may have been many orders of magnitude more cost-effective - if we think there is a reasonable chance that a large enough impact could directly cause human extinction.

3.2 What might yet be achieved

The work of spaceguard programmes is a clear example of how NEO detection efforts have been highly cost-effective in the recent past, but there is also some reason to believe such interventions remain relatively cost-effective today. This section considers two ways in which the existing work could be extended: detecting and tracking any remaining asteroids, or extending detection efforts to include comets as well.

As noted earlier, over 95% of near-Earth asteroids with diameter >1km have already been tracked. In fact, if the best recent models of total asteroid population are accurate, existing spaceguard programmes have actually tracked *all* such asteroids: the remaining uncertainty, and corresponding risk, stem from the model itself. We might therefore improve our risk estimate by continuing to test and improve the model, and tracking any asteroids that may have fallen through the cracks.

The best recent estimate of catastrophic impact over the next century, incorporating evidence from existing spaceguard programmes, is around 1 in 120,000 (Ord, 2020). Using the same estimate of the damage due to a catastrophic impact as in the preceding section, this risk represents close to 16,000 lives lost in expectation over the next century. Theoretically, work to improve the model, and track remaining asteroids, could reduce this number to near zero. If we assume that this work could be achieved for around $1.2 billion[10] (USD), then its estimated cost-effectiveness works out to around $75,000 per life saved in expectation.

The best recent estimate of an asteroid impact leading to human extinction over the next century is around 1 in 150 million (Ord, 2020). Making a similar set of steps to those in the preceding paragraph, this works out to around 670,000 lives lost in expectation. If we again assume this work could be completed for $1.2 billion (USD) or less, the associated cost-effectiveness works out to just under $1,800 per life saved in expectation.

A second way that existing NEO-detection work could be extended concerns comets, rather than asteroids. The total risk of catastrophic impact posed by comets is plausibly similar to that posed by the remaining asteroids, at around 1 in 120,000 over the next century (Ord, 2020). As a result, we can bench-mark the cost-effectiveness of this work in exactly the same way as for the risk of catastrophic impact from asteroids. However, there is some reason to think that comets would be significantly more expensive to track than asteroids: comets spend only a small fraction of their orbital periods near the Sun, and are extremely difficult to detect at other times. As a result, comet detection is likely to be significantly less cost-effective than asteroid detection, in the absence of major advances in detection techniques.

[10] This figure comes from an extrapolation of costs to date, based on Dreier (2019).

4. References

Dreier, C. (2019). 'How NASA's Planetary Defense Budget Grew By More Than 4000% in 10 years.' *The Planetary Society.*

JPL (2019). Discovery Statistics - Cumulative totals. http://cneos.jpl.nasa.gov/stats/totals.html

Mainzer, A. et al (2011). 'NEOWISE Observations of Near-Earth Objects: Preliminary Results.' *The Astrophysical Journal*, (743)2.

Matheny, J. G. (2007). 'Reducing the Risk of Human Extinction,' *Risk Analysis*, 27(5).

Morrison, D. (1992). 'The Spaceguard Survey: Report of the NASA International Near-Earth-Object Detection Workshop,' *NASA*.

Newberry, T (2021). 'How many lives does the future hold?' GPI Technical Report.

Ord, T. (2020). *The Precipice*. London: Bloomsbury.

Schulte, P., et al. (2010). 'The Chicxulub Asteroid Impact and Mass Extinction at the Cretaceous-Paleogene Boundary,' *Science* 327(5970).

Stokes, G. H. et al (2017). 'Update to Determine the Feasibility of Enhancing the Search and Characterisation of NEOs.' Near-Earth Object Science Definition Team.

U.S. House of Representatives (2013). *Threats from Space: A Review of U.S. Government Efforts to Track and Mitigate Asteroids and Meteors (Part I and Part II)* [Hearing]. U.S. Government Printing Office.