

Achieving Impact on a 1,000-Year Scale

Published on Apr 26, 2017 by Niko Canner



Over the next thousand years, a succession of human societies will determine the future of life in this part of the universe.

We will shape and be shaped by our evolving impact on the environment. We—and the other living things around us—will contain or be overcome by our increasing destructive powers. We will have learned to coexist with intelligences beyond human intelligence. Our ability to alter life, including the genetic properties of human life, will have bestowed new powers and responsibilities, which we will have matured to steward, or failed to. We may well have created and found ecosystems beyond our planet that will support life of our kind.

It is entirely reasonable to argue that these several tests will be passed or failed within a shorter time span; perhaps a few hundred years. It's fairly certain that these tests will be decisively passed or failed—not just by avoiding specific bad events, but by creating a framework in which good outcomes can consistently be achieved—within a thousand years. If, for instance, there isn't a sustainable framework for containing destructive power within a few hundred years, robust enough to withstand the continued exponential improvement in weapons, the odds of avoiding catastrophe in the subsequent centuries is vanishingly small.

This reality represents an immense concentration of agency in the span of only a few dozen generations, a tiny sliver of time on the evolutionary scale. Represent the history of human life—from Mitochondrial Eve to today—as a foot-long ruler, laid up against a history of life on earth nearly as tall as Kilimanjaro. On this scale, the time span of societies based on agriculture is the length of a housefly, and the span of a thousand years is the width of the serrated edge of a quarter.

To take this challenge seriously, we must begin by conceptualizing what human society needs to achieve in order to create a positive outcome. A minimum set of achievements should include the following:

- **Capable of selecting a deliberate ecological path.** The climate will undoubtedly change

within the next 1000 years, and though our technologies will develop, we cannot expect to ever be able exert full control over climate. The consequences of unforeseen and unintended large-scale climate shifts could be catastrophic for human society and for the biosphere. The “minimum achievement” in this space should be to reach a comparable level of ability to select an ecological path to the level we’ve reached in macroeconomics over the past century. For all the imperfections in macroeconomics—known to us all in the wake of the recent financial crisis—we’ve reached a fundamentally superior level of capability to where we were in the early 20th century. We can understand the likely consequences of different choices and have the institutions, measurement tools and interventions that make paths of macroeconomic evolution tamer and more benign than they were in an earlier era. The tradeoffs among different alternatives remain painful, as will undoubtedly always be the case for ecological paths as well, but those trade-offs can be broadly dimensioned, understood and made, albeit in broader and rougher strokes than we might wish for.

- **Contain the impact of advancing destructive technologies.** One thousand years ago, the earliest military uses of gunpowder were just emerging in China. Human destructive capacity will continue to increase exponentially over the next thousand years—and even if rate of that exponential increase is smaller than what we’ve seen in the last two hundred years, the cumulative impact will be immense. Politics will shift as smaller-scale actors become capable of achieving increasingly large-scale effects. The “minimum achievement” in this domain is a correspondingly exponential improvement in containing the risks associated with this trend line - through a combination of limitation in the spread of weapons, surveillance, preventative and defensive technologies, as well as reduction in the prevalence of the motivation for violence.
- **Become resilient to catastrophic events.** Over the time frame of a millennium, it is very likely that the magnitude of the most catastrophic event—be it war, other forms of violence, “big bang” natural events (e.g., an asteroid collision) or “slower build” natural events (e.g., a series of pandemics)—will dwarf any events in our recent experience. While the probability distribution of the frequency and magnitude of catastrophes can be altered in a range of ways, even with optimistic assumptions there is still a high bar for the “minimum achievement” of resilience sufficient to withstand the worst likely events. Suppose for instance that the ex ante probability that the world avoided a nuclear catastrophe during nearly five decades of the cold war was 90% (i.e., in nine worlds out of ten, things turned out as well in this way as in our actual history), and suppose that over the next millennium, we cut the risk of such catastrophic outcomes, from all sources, in half - a 5% chance of a catastrophe of a magnitude that it could wipe out a significant fraction of civilization in each fifty-year period. Under these assumptions, the odds are only 1 in 3 of avoiding this kind of truly major catastrophe through the year 3000. The “minimum achievement” here is building the resilience - biological, ecological, and institutional—to be able to reestablish a livable world and renew civilization in the aftermath of a catastrophe orders of magnitude beyond our worst historical events.
- **Establish ethical reciprocity with multiple forms of awareness and intelligence.** Whatever stance one takes towards predictions about development of artificial intelligence in the next handful of decades - from Ray Kurzweil’s view that “the singularity is near” to the views of AI skeptics - it is difficult to imagine that we could be more than a century or so away from having meaningful non-human intelligences, virtual or embodied. There’s a great deal of

dialogue about ensuring that artificial intelligences behave ethically toward us, especially if these intelligences come to exceed human capacities. Certainly this should stand as one of the minimum achievements.

Within a century, it is likely artificial intelligence to achieve a level in which entities have ethical claims about how they should be treated and how they should be enabled to advance their ends. Over the time frame of a thousand years, these questions become still more complex, in a layered world of intelligences making other intelligences of a potentially wide range of natures and of forms. The ethical questions of how to create a society that can honor such a wide range of being and beings are of parallel import to the ethical questions at the foundation of the invention of human rights as a construct and our progress in eliminating slavery as a politically and socially accepted practice. The “minimum achievement” in this domain is the avoidance of a moral catastrophe in which we disregard the palpable ethical claims of other intelligences, or these intelligences disregard the ethical imperative to care for our interests.

- **Navigate the psychic and social dynamics of vastly expanded fields of human choice.** As biotechnologies, material technologies, information technologies and other relevant domains advance, more and more givens about human life—our personalities, our cognitive and emotional abilities, our lifespans—will become variables that we can choose. A society in which people can make choices of these kinds—or in which such choices can be made for them by others—could function in very different ways than the societies we’ve known. The “minimum achievement” here is again the avoidance of catastrophe: of any form of dystopia in which the expanding domain of what can be chosen leads to forms of society that are ethically unacceptable.
- **Develop sustainable and robust systems of governance.** With the expansion of our agency comes expansion in the necessary scope of governance. Nations were large enough constructs to solve for the needs of earlier ages. Even today, many of the questions that matter most to human well-being—war and peace, security from terror, public health, climate, the functioning of markets—reach well beyond the scale of nations. The continued increase of our capacity to influence the world will continue to enlarge the sphere of needs for collective action of different kinds. At the same time, the means for forming powerful non-state actors of different kinds (from the Gates Foundation to Al Qaeda)—identifying like-minded people, building relationships at a distance, organizing, gathering resources, acting collectively—have deepened at a rapid rate, and will continue to deepen. Centuries out, human society will need to act collectively on the vast questions considered above and maintain sufficient stability among a vast array of diverse powerful actors of with different value systems. We will need to create the foundation of law and its enforcement that guarantees the most fundamental rights and needs amidst this kaleidoscopic flux of institutions at different levels and scales. The “minimum achievement” here is simply to avoid a level of volatility that unleashes broadly destructive consequences.

Each of these represents an agenda that, while much larger than, say, placing a man on the moon, is still defined and to a meaningful degree bounded in its scope.

Highest-level agendas of this kind convey a direction that the *very big we* of human societies across

multiple generations need to pursue. What could it look like to take effective agency at this level of the very big we? What does it mean to act purposefully toward an end that extends so far out spatially, temporally and in relation to the size of any given individual's personal capabilities? These questions serve as beacons, telling us the work that we must begin at once, work that will demand more than the best of what we know, now and for generations to come.