

Population as a global priority?

Hilary Greaves, University of Oxford

August 2017

The United Nations Population Division's latest report predicts a global population of over 11 billion by 2100. That is the 'medium' projection, based on standard demographic transition theory. There is also a 'low' projection, in which the total fertility rate is lower by half a child per woman; here, population peaks at 8.7 billion mid-century, returning to 7.3 billion by 2100.

Would the 'low' scenario be better than the 'medium' one? If so, why exactly?

Such questions, of course, are only of practical importance if there are things we can reasonably do to influence the path of population size. Skeptics think that the only available tools to this effect are morally impermissible (for example, forced sterilizations). But this is too pessimistic. Tax and benefit policies also affect reproductive behaviour, as do moral arguments. Not all incentivisation is coercion.

The conventional wisdom is that significantly lower population sizes would be clearly better, on the grounds that reducing population would both increase average well-being in our own lifetimes, and help preserve the state of the planet for future generations.

But we must not accept this uncritically. It is far from obvious whether reducing population size would increase average well-being levels; natural resource constraints would have to be traded off against economies of scale. More fundamentally, it is also far from obvious that *average* well-being is the thing to maximise in any case. The question is not a purely technical matter: it involves value judgements. A more plausible candidate for optimisation is *total* well-being, added up across all the persons who are blessed with existence. Additional lives have intrinsic value, over and above their effects on pre-existing lives.

This 'total view' will in general recommend larger population sizes than the 'average view'. Its advocates, however, do take very seriously the possibility that overpopulation in the short term could be short-sighted. What this view cares about is total well-being added up *across all time*, not merely the total well-being of people whose lifespans overlap with ours. If overpopulation in the short term in any way decreases the prospects for a flourishing human civilisation extending far into the future, this would generate a strong argument for lower population by any reasonable lights. Concerns about climate change, other modes of environmental degradation, and extinction risks are all potential fits for this bucket.

Extinction risks provide arguably the dominant consideration. If all goes well, humanity can expect a tremendously bright future, continuing for millions or even billions of years. But a number of things could wipe us out prematurely – potentially within the next century or two. Such premature extinction, and the loss of future it implies, would be *astronomically* bad.

If an event is astronomically bad, then, under conditions of uncertainty, even a tiny decrease to its probability can itself be an extraordinarily large improvement. Such effects can indeed dominate expected value calculations. This has been increasingly recognised, for example, in discussions of "fat tails" in climate change.

We should ask, therefore, whether the medium or instead the low projection leads to lower chances of premature extinction. Here too, conventional wisdom might seem to favour lower population: for example, overpopulation is thought to incubate and exacerbate conflicts over scarce resources. But this piece of conventional wisdom is not really about extinction-level risks. A conflict over water resources in Gaza is an extremely bad thing, but there are no plausible routes from such relatively local conflicts to human species extinction. Nor are there plausible routes even to radically reduced long-term prospects for human flourishing in the far future, though short-term negative impacts may be very high.

The major anticipated sources of truly extinction-level risks are instead technological. They include scenarios where value-unaligned AI usurps the resources that humans would need to survive, or advances in synthetic biology allow malicious individuals to spread a species-terminal pandemic. Experts working in extinction-risk mitigation generally do not view population size as a major driver. The question is nonetheless worth asking: would higher population increase extinction risk, decrease it, or neither?

Here is a reason to think the answer might be “increase”. Suppose some potentially species-lethal technology comes online. Will humanity survive? One key parameter is the state of what we might call ‘social capital’ at the time the threat is faced. If we have appropriate incentive systems and governance structures to facilitate global group rationality, our survival prospects are significantly better.

Technological development is brought about by researchers. In a world with more people, there will naturally be more researchers, and so the pace of technological development will speed up. The same is probably true of social progress: more people means more social innovators, and so (probably) faster social progress. The crucial question for our purposes is comparative: would population increase speed up social progress *relative to* technological progress (thus improving the state of social capital at the critical time), or vice versa?

There are some reasons to believe the latter. Social progress in particular is also limited by the physical passage of time. Processes of international cooperation, for example, advance at a rate determined in part by the schedule of political elections and international summits, and a 5-year political cycle does not suddenly become a 4-year one when population size increases by 25%. And some aspects of social progress depend on the simple replacement of one generation with the next. Although some processes of technological development too have to be carried out in series rather than parallel, this effect seems to be smaller.

Here, then, is a sketch of a novel line of argument for favouring the ‘low’ over the ‘medium’ scenario, having nothing to do with more traditional concerns about resource scarcity or environmental degradation. It is far from watertight, but seems to merit some thought. But of course, if our conclusion is that the fate of humanity might depend on speeding up social progress relative to technological progress, there are likely to be more direct and effective levers for this than population size. That seems to merit more thought.