

Chapter 3. Overpopulation or Underpopulation?

Toby Ord

Overpopulation has been one of the major global concerns of the last fifty years. The rapid rise of the world's population means many more mouths to feed, raising the possibility of mass famines, and making overpopulation a major humanitarian concern. The increased labour supply to non-Western nations and the projected increases in immigration have made it a major geopolitical concern. The projected increase in resource usage and associated pollution has made it a major environmental concern.

These are the main lenses through which rising population has been viewed. But they have missed a major part of the story.¹ In addition to its well known costs, increased population brings many *benefits*. It means more scientists to discover how our world works, more inventors and thinkers to help solve the world's problems, and more workers to put these ideas into practice. It means more great writers, musicians, and artists to explore the human condition, and to share their masterpieces. It means more people, more sons, daughters, fathers, mothers, who get a chance to experience the world in all its richness—more beneficiaries for all the creations and improvements we hope to bring about.

It is impossible to have a nuanced and mature approach to population without considering these benefits alongside the costs. Do the costs strongly outweigh the benefits? If so, overpopulation is big problem and we will need to work out how to prevent or manage it. Are the costs and benefits roughly balanced? If so, we should put less attention and effort into reducing population growth, for these resources could be better spent on other pressing issues. Or do the benefits outweigh the costs? If so, we may be facing *underpopulation*—having too few people—and strange as it may sound, we might face a duty to increase our population.

In this chapter, I aim to explore the question of the ideal number of people for the earth to bear. As the benefit of a larger population is a neglected area, it is there that I spend most of my time. I focus on the benefits of increased population for the creation of information goods, and on the intrinsic value of the new people themselves. I then consider several important questions about the various

¹ Of course there are some well-known exceptions including Ester Boserup's (1965) historical analysis of the links between population density and technological change and Julian Simon's (1981) argument about wealth creation countering dwindling resources.

environmental limits, and on the distribution of the population increases. Finally, I consider how things might change in the future—how we might be able to relax some of the limits or increase some of the benefits.

No one person has the expertise to arrive at a final answer to the question of whether the benefits outweigh the costs, and I do not attempt to do so. My aim is instead to try to frame population policy in a way that takes accounts of benefits as well as costs, and in doing so, to help us confront the right questions.

Instrumental Benefits

The benefits of increased population can be split into two categories. The first is the instrumental benefits brought about by the presence of additional people. For example, if additional people produce new inventions which improve the lives of everyone, then this is a clear benefit of a larger population and must be weighed against any drawbacks. The second category is the intrinsic value of the lives of the additional people. It considers that it is good for those people to exist and to experience the joys, loves, and excellences that life can offer. It is much less clear how to assess this category. It involves many philosophical questions and there is some debate about whether bringing new people into the world can have intrinsic value at all. Let us therefore begin with the more concrete case of instrumental benefits and confront the philosophically charged intrinsic benefits later.

There are many ways in which creating additional lives can create instrumental benefits. Imagine a childless couple who, after some thought, decide to have a child. This child will make its parents lives better in various ways, chief among them is probably the special relationships they will form which the couple wouldn't have otherwise been able to experience. The child will of course have a financial cost to them, and will take up time that would have been spent on various leisure activities, but presumably the joy and satisfaction for the parents are often enough to outweigh the costs.

The child will also affect many other people in positive and negative ways through its personal relationships, its work, its purchasing decisions, its charitable contributions, and its environmental impact. I don't know how to determine whether these benefits outweigh the costs, but it would be irresponsible to deny that there are benefits being produced which need to be weighed against costs, and to not seek out the advice of specialists who could try to determine this.

One type of instrumental benefit deserves some special attention. *Information goods* may be the most important type of instrumental benefit. These include physical

items like books, CDs, and DVDs where the value lies in the information encoded in the item. They also include downloadable software or music, where the pattern of information is once again key and where no underlying physical medium is actually traded. I shall also take the term to include other forms of valuable information, such as scientific discoveries, technological inventions, and novel systems for government or policy.

We can contrast information goods with *material goods*, such as hammers, cars, oil, or food. Consider the differences between hammers and (recorded) songs. Hammers are made of matter, while songs are fundamentally patterns which could be encoded in many different substrates. Each hammer must be laboriously made, perhaps in a smithy, or in a factory. In contrast, a lot of work goes into writing and recording a song, but once recorded, it can be very cheaply copied. Each hammer benefits a single user, so its value is roughly independent of the world's population. In contrast, each song can benefit everyone who desires it—potentially millions or billions of people. Thus unlike a hammer, a song becomes much more valuable the larger the population.

This is the key feature of information goods. Their value increases with the size of the population which both has access to them and could derive some benefit from them. In many cases there is a global appetite for the information and global distribution, so the value of an information good often increases with the size of the global population.

Over the centuries, we have shifted an increasing part of our economy to the production of information goods. They are now ubiquitous. They can be artistic, such as novels, poems, songs, films or recipes. They can be technical, such as inventions, designs, and new techniques. They can be political, such as new political systems, ideology, or policy approaches. They include all the software that we use on more than a billion personal computers. They include all scientific discoveries, and more than that, they include all academic research in all disciplines.

Adding more people to the world will increase the number of people working to create information goods. This will increase the number of such goods, the variety, and also the peak of quality in many different areas. Additional medical researchers will develop cures for additional diseases; additional programmers will write new and better software; additional activists will develop new and better reforms to the political process. Additional musicians will create new pop hits, and new masterpieces, launching new genres and styles. Of course they will also produce much more dross, but we are entitled to focus on the peaks instead of the troughs because we can by and large choose our preferred package of aesthetic information

goods. With more people there will even be radio stations and critical reviews in more niches, increasing the quality and relevancy of the music, novels, and films that reach you.

Of course it is not clear that twice as many people would produce twice as much value from information goods. Some of these goods will have only local reach or relevance, and many will have diminishing marginal value. We tend to pick the low hanging fruit of science and invention first, snapping up the easy to acquire improvements and leaving the more challenging or less beneficial ones for later. Our cultural improvement will have diminishing marginal returns due to the crowding of our schedules. There are only so many plays or films we can watch, songs we can listen to, or books we can read. Better ones will crowd out worse ones, but if there are limits to the upper end of how much you can like a piece of art, then we will have diminishing returns as we move towards a situation in which you only watch, listen to, and read works that you consider masterpieces. However, we are still a long way from that point, and it sounds like a valuable point to reach!

If the diminishing returns are a major concern, one could also reap the benefits of information goods in another way. If we doubled the population, then instead of doubling the output of information goods, we could halve the input of labour to reach our current output. We could have information workers work 20 hours a week instead of 40, and still have access to the same rate of scientific, technological and cultural output that we have today. This wouldn't change the proportion of people working to produce information goods, so it wouldn't lower the productivity of the material goods economy, it would simply allow us to have a lot more leisure time to spend as we wish.

A final way in which additional people could provide large benefits through information goods is that in addition to being information producers, they would be information consumers. This means that the markets for information goods would expand, making it more profitable for anyone to begin producing information goods. This would allow for more lavish budgets on mainstream films and albums, but perhaps more interestingly, it would mean more people in each niche of aesthetic taste. This would mean that some niches which are currently too small to have many works of art produced for them would have enough market power to entice artists and production companies to deliver more works. The same is true of non-aesthetic niches: there will be more good non-fiction works on currently obscure areas.

Intrinsic Benefits

Our lives contain significant value. While all our lives are marked in some way by hardship, grief, and pain, they are also filled with hope, joy, love, excitement, and contentment. In almost all cases, we judge that the good aspects outweigh the bad, and find intrinsic value in our lives.² The total value of someone's life consists of its instrumental value (the aggregation of all the intrinsic value they add to other people's lives) and its intrinsic value (the value of the life for the person himself or herself). We have discussed the first part, but we would be neglecting a major part of the question if we didn't consider the second.

It is common to hear overpopulation campaigners speak of the need to lower the population by a billion or more. I am yet to hear any discussion in the mainstream media of the lost intrinsic value in not having these people on this earth. What is the intrinsic value of a billion people? Consider, if you will, the population of the islands that make up the United Kingdom and the Republic of Ireland. Think of the currently existing people who live there, in the small villages, the bustling towns and the thriving cities. Now add all the people who ever lived there before them, through the many centuries, with their different ways of life, all the way back in time until humans first arrived there. There have been a great many inhabitants—hundreds of millions—but the total is less than a billion.

Now consider whether there would have been something of value lost if they had never existed—if the islands had remained uninhabited and the rest of the world had gone on as usual. Of course there were various instrumental benefits for the rest of the world caused by these people (and harms too), but to my mind at least, there is also a great internal value: the intrinsic value of all of those hopes and joys and passions. Even if these people had been completely isolationist and had added nothing of instrumental value to the world, it would have been a great loss had they never existed.

This is the scale of the loss of intrinsic value there would be if we were to reduce the population by a billion. It is thus quite astonishing to see how casually reductions such as these are invoked in debates on population.

Consider your own life, and whether there is some value in it. Setting aside the effects you have on others, wouldn't the world be poorer without you—without

² Many philosophers now use the technical terms 'final value' or 'noninstrumental value' where I use the less precise but more readily understood 'intrinsic value'.

your dreams and passions and experiences? If we had decided to radically slow population growth in the past, you might not be here. If you think your life is good, then something would have been lost.

I think that considerations like those above are often forgotten when talking abstractly about population. You are population too. As are your mother, father, brothers, sisters, friends and lover. However, the term 'population' can be dehumanising. It makes us forget that we are talking about you and me: about individuals with their own richly textured lives. It makes us forget about the intrinsic value in us all, and to just focus on the aggregate costs such as pollution or crowding.

These questions about the intrinsic value of additional lives have recently received much attention within moral philosophy.³ It first achieved prominence in the late 19th century in Henry Sidgwick's (1907) writings on utilitarianism. He pointed out that there is an important distinction between trying to increase the total happiness and trying to increase the average happiness (ibid: 415). If we are considering an ethical question where the population is fixed, then there is no significant difference between the total and the average: any action that increases one raises the other in the same proportion. However, if the population can vary, then the two come apart.

To use a heavily simplified example, suppose we could bring about one of two different outcomes:

- (A) 5 billion people living very happy lives
- (B) 10 billion people living lives that are almost as happy as those in (A)

The average happiness is greater in (A), while the total happiness is greater in (B).

Over the last century, this distinction has risen to increasing prominence. Philosophers have come to several conclusions. Firstly, this distinction is not just of interest for utilitarian theories of ethics. Whether or not one accepts utilitarianism, one should care about whether it is better if outcome (A) occurs or if outcome (B) occurs. However, in this general case, we shouldn't just think of the happiness in someone's life, but rather some measure of all the value in their life, which may include more than just happiness. Secondly, there are potentially far more than just two theories for the value of different populations (Parfit 1984; Arrhenius 2000b; Blackorby, Bossert, and Donaldson 2005; Arrhenius forthcoming). The study of these many different theories has come to be known as *population ethics*.

³ There has also been discussion of this question from a related angle in the economics literature, as explained by Tony Atkinson in Chapter 2 of this volume.

For now though, let us focus on these original two theories, that have come to be known as *the Total View*, and *the Average View*. These views give different recommendations in terms of the intrinsic value of people. According to the Total View, the value of a group of people is the value of each of the people put together. If we add a new valuable life to a population without changing anything else, it makes the population more valuable. Of course there will eventually come a point where adding a new person will have more costs for the others than it has benefits, making the other lives worse overall. This is a bad effect, but according to the Total View, the intrinsic value of this life may be enough to make up for the costs of overcrowding for others. The Total View recommends adding new lives until the point at which this would no longer increase the total value (because the combined instrumental and intrinsic value of the life come out to be net negative).

In contrast, the Average View will typically recommend a smaller population size. It says that we should stop at the point where this would no longer increase the average value of a life. This could be because the intrinsic value of the new life would be below the average value (and wouldn't make up for this through instrumental benefits to others) or because the new life would impose enough costs on others to lower the average, or some combination.

It should be stressed that it is very unclear whether we have yet passed either notion of an optimal population. These ethical views explain what the intrinsic value of a population is, and thus what it would mean to have too many people. Much more empirical evidence on quality of life and environmental and economic issues would be needed to actually make use of either standard.

While these are the most well known views on population ethics, they have both come under considerable attack. In 1984 Derek Parfit published a book called *Reasons and Persons*, which launched population ethics as a major strand in moral philosophy. In it he forcefully showed that the Total View leads to 'The Repugnant Conclusion' described in the previous chapter: the supposition that any expansion of a 'sizable' and 'happy' population will be beneficial, even if the value of many lives in utility terms end up approaching zero and become 'barely worth living' (Parfit 1984: 388). For example, if the larger population had a thousand times as many people, and their lives were a hundredth as good, then it would contain more overall value according to the total view, though many of us would judge it as inferior. It is worth noting that the example populations we are considering here are not realistic options for our world. Restricting ourselves to realistic options would not allow us to create a choice as stark as this one, and it is through simple clear thought experiments that philosophers try to judge whether there are problems for an ethical view.

Many philosophers believe that the Repugnant Conclusion makes the Total View implausible as a theory of population ethics. However, others argue that this is premature (Tännsjö 2002; Huemer 2008). For example, the Repugnant Conclusion asks us to imagine a world with trillions or more people and our intuitions about such large numbers are notoriously shaky. When we try to imagine a trillion people, our mental image might not look very different from our image of a billion people, or even just a million people, whereas we find it easy to imagine people with much less value in their lives. This predictable failure of imagination could lead us to undervalue that outcome.

In contrast, the arguments against the Average View are so strong that to my knowledge there are no remaining philosophers at all who advocate the Average View. The first challenge for the Average View is that it implies that the world would be improved by killing people whose lives are worse than average (at least insofar as this doesn't reduce the quality of life for the others). This is widely considered to be an absurd conclusion. However, this would only follow from an *instantaneous* version of the Average View, in which we look at the average quality of life at an instant. The Average View can also be interpreted in a *timeless* manner (Parfit 1984: 420; Broome 1992: 117). In this case, we consider all the lives that were ever lived and will ever be lived and try to increase the average of this vast 'timeless' population. Once someone is born, they are irrevocably part of this timeless population, so the average will not be improved by their death (unless their life is so bad for them that it has negative value, in which case it is much less clear that their death would be a bad thing).

However, there are numerous other problems for the Average View that the move to a timeless population does not solve. One of these is that it makes the ethics of our decisions more interconnected than seems plausible. For example, the question of whether a couple on an isolated island should have children now depends upon the world average, even if no one on that island ever effects the rest of the world, or vice versa. The move to the timeless population actually makes this worse, as we need to consider the timeless population average, which includes people in the distant past and in the distant future. We would need to study how many people have lived in the past and how good their lives were to make a decision, as well as predicting population size and quality in the future. Perhaps we would be compelled to produce more people now to increase the historical average quality.

The strongest argument against the Average View is probably the following hypothetical situation that Parfit presented in *Reasons and Persons* (Parfit 1984: 422):

Hell Three

Most of us have lives that are much worse than nothing. The exceptions are the sadistic tyrants who make us suffer. The rest of us would kill ourselves if we could; but this is made impossible. The tyrants claim truly that, if we have children, they will make these children suffer slightly less.

In this situation, the Average View says that it would be better to produce another hellish generation since this would slightly increase the average quality of life. However, since this would involve creating many new lives of negative value and has no compensating benefits, it seems very counterintuitive that this would be an improvement at all, and that the best thing they could do would instead be to have no more generations.

Examples like this have caused the Average View to lose all support among moral philosophers.⁴ Since we also have reason to doubt the Total View, this has led to a wide proliferation of alternative views on how to value populations. While a full discussion of such views would take us well beyond the scope of this chapter, I shall briefly describe two of the most promising approaches.

Critical level theories are based around the idea that although lives which have negative value for that person are always bad to add to a population, lives with positive value for that person might not always be good to add (Blackorby, Bossert, and Donaldson 1997; Broome 2004). They could be neutral to add, or even bad to add if their value is sufficiently close to zero. This family of theories is similar in some respects to the Total View, but doesn't suffer from the Repugnant Conclusion as lives barely worth living will either have no value or negative value. However, there are other important problems for such theories, not least of which is that they can say that it is better to add lives that are not worth living than to add some larger number of lives which are worth living (Arrhenius 2000a).

Perhaps the most popular in recent years are the various *person-affecting* views. These are theories which are based around the intuition that a state of affairs can't be better (or worse) than another unless there is at least one person for whom it is better (or worse) (Narveson 1967; Glover 1977: 66).⁵ These views say that certain kinds of

⁴ Curiously, it still seems to be used within economics. I suspect this is because the practitioners are not familiar with the critical literature.

⁵ In cases where there are the exact same people in both states of affairs this is akin to the economists' concept of a Pareto improvement. However, if the number or identity of the

merely potential people shouldn't enter into our assessment when comparing two outcomes. They differ in exactly how they spell this out, with some theories saying only presently existing people count, or that only the people who will actually end up existing count, or only the people who would exist in all alternatives count. While tempting, these theories have been shown to be open to a very similar set of devastating objections (Arrhenius 2009; Arrhenius forthcoming: 151–207).

In summary, there is no consensus on exactly how one should measure the intrinsic value of increases in population. The Average View is widely rejected, but there are reasons to be cautious about using the Total View. Other views are still very much under development, but those that have been proposed in the literature appear to have their own grave flaws. Indeed, there are even several compelling arguments that any approach to valuing a population will have to have at least one of a short list of undesirable properties (Arrhenius 2000a).

This does not mean that we can simply ignore the intrinsic value of additional people. As the initial intuitive arguments of this section show, we do have the intuition that subsections of our population can have considerable intrinsic value and failing to count them in an analysis of overpopulation would be morally reckless. For now though, there is no consensus on exactly how this is to be done.

Costs

I have spent considerable time discussing the instrumental and intrinsic benefits of higher populations as these topics are often neglected in discussions of overpopulation. Costs are much more commonly discussed—especially in the form of limits, such as the number of people that could be supported by the world's fresh water, by our food production, or by our mineral and fossil fuel resources. Chapters 6, 7, and 10 in this volume discuss such limits in detail.

Limits like these are really a form of cost. They are not typically binding constraints, but rather points at which we would have to change from business as usual. We might have to develop new forms of food production, or even just eat less meat, both of which would involve costs. Even without changing our food production, we could probably exceed a calculated limit, but only for a finite amount of time, before we paid a cost in terms of deaths due to starvation. There are also costs as we approach a limit. For example, if we were to get close to using all natural fresh water for

people can be different in the two cases it is far more controversial. For example, if we were to add people with hellish lives to a population, we intuitively think that this makes it worse even though there is no-one for whom it is worse.

drinking, we would need to exploit more and more of the wilderness surrounding sources of water, eventually including some of the most beautiful places which we would much rather preserve.

There is an important source of confusion here that can derail many conversations about the ideal population size. Consider the question of whether we are currently near the limit on food production. In one sense we are, for we are already using most of the world's arable land, and produce only slightly more food than is needed to feed all of the world's people. However, as Charles Godfray points out in Chapter 6 of this volume, meat production uses these natural resources very inefficiently. If we really wanted to feed as many people as possible, we could reduce meat production, or even abandon it completely.

So in one sense we are near the limit, while in another we are far from it. This can cause considerable confusion if a person using one sense engages in debate with someone using the other. The best way to see this is that the lower limit is a 'soft limit', where we can no longer pursue business as usual without disaster ensuing, while somewhere above that is a 'hard limit' where even abandoning business as usual and using all technological and policy measure available we still can't exceed that population without disaster.⁶

As we exceed a soft limit and head towards a hard limit, there are two types of cost that we might encounter. One is that departing from business as usual will impose costs (such as the costs of researching and introducing a new technology, or the cost of forgoing meat consumption). The other is potentially much worse. It is that even though we could safely exceed the soft limit were we to adjust our activities, we may well not have the will or the coordination to do so. For example, people in rich countries may refuse to abandon meat, and instead just pay the increasing costs of meat in a period of food scarcity, pricing those in poorer countries out of sufficient food to survive (Sen 1981).

These limits of political or social feasibility are harder to calculate than technical soft or hard limits, but they are no less important. If it really is socially impossible for rich countries to become vegetarian in order to let people in poorer countries survive, then we certainly want to find out and to factor this into our thinking on the ideal population. However, we should be careful in how we describe the problem. In that case, the situation could either be described as overpopulation, or as shocking selfishness of the world's rich. After all, even a population of just two people could

⁶ While I describe this as a 'hard limit' in the sense of being unavoidable, it will be a vague rather than sharp boundary.

count as overpopulated on that definition if the more powerful of the two demanded all the resources for his or her personal use. This issue frequently arises when people claim that with Western resource consumption the world can only sustainably support a fraction of the current population. This may be true, but it could just as well be interpreted as overconsumption instead of overpopulation.

We should also remember that not all limits take the form of global limits. For example, we might be able to support additional people in some regions or countries but not in others, or the costs of having additional people might be lower in some places than in others. Along these lines, people in poorer countries typically consume much less in the way of resources than those in richer countries, so the earth may be able to support more population growth in poor countries than in rich ones. Similarly the benefits created by additional people in some places might be greater than additional people in other places. This means that even if we knew all the answers to the many complex empirical questions concerning overpopulation, the answer is unlikely to be a simple 'yes'/'no', but more a complex set of policy prescriptions that could involve using different approaches in different places.

It is important to note that the costs and benefits of having a larger population are not static, but change greatly over time. For example, Paul Ehrlich (Ehrlich 1968: xi) got it wrong when he opened his treatise on overpopulation by saying:

The battle to feed all of humanity is over. In the 1970s and 1980s hundreds of millions of people will starve to death in spite of any crash programs embarked upon now. At this late date nothing can prevent a substantial increase in the world death rate . . .

Such a disaster may have ensued based on the business as usual limits of food production, but that period saw an extremely rapid rise in food production due to the so-called 'green revolution'. The use of irrigation, fertilizers and pesticides, alongside the introduction of new high-yield cereals saw wheat yields per hectare dramatically increase (Evenson and Gollin 2003; Goldin, and Reinert, 2012, table 5.5). These developments expanded the technical limits on how much food could be produced with the available arable land, and so changed the global picture regarding overpopulation.

It is in the nature of science and technology that it is difficult to predict whether and by how much we could raise other limits, but we should certainly bear this possibility in mind. For example, genetic engineering of crops or expansion of new farming methods such as aquaculture could again expand the available food supply, while any source of clean cheap energy would relax a number of limits

simultaneously. Consideration of technological improvements is important for planning and prediction purposes, and because we can direct funding towards finding ways to relax these limits. Indeed this might be one of the best policy levers for dealing with potential overpopulation. Expanding limits allows us to get the benefits of additional population without some of the costs, and may also be easier to achieve than preventing the people from being born in the first place.

Technology is not the only way in which limits could be relaxed. As mentioned earlier, many of the soft limits we face are due to social limitations. We are near the soft limit on food production because we eat so much meat. We are exceeding our limits of CO₂ production because we aren't prepared to pay more to use the clean technologies that have already been invented. If we could make social progress on convincing people to change their behaviours to be more altruistic on these fronts, we could have as much impact as a technological breakthrough. Even if people were just altruistic enough to allow their governments to pass the appropriate legislation to disincentivise these activities, this could make a dramatic difference. Social change on this scale is not easy, but nor is it impossible, as examination of the rise of vegetarianism, civil rights, or the environmentalist movement shows.

Improvements in science, technology and social change can also tip the balance in favour of larger populations by increasing the benefits of population as well as decreasing the costs. For example, the computer and the digitization of music and film have greatly reduced the costs for reproducing these information goods and have thus increased the benefits of large populations. Social reforms in intellectual property law might further increase the benefits by allowing more people to benefit from each new invention or artistic work.

Finally, it is possible that in the future the global population will trend downwards on its own. This was previously thought likely by the UN (2004: 13), though their updated report now suggests that this is less probable (UN 2011). The further the world's population were to fall, the more likely it would be that it would be underpopulated: that it would be better to have more people. If it fell sufficiently far, say to less than a billion people, then there would be widespread agreement that the world was underpopulated. However, the arguments that I have given as to the possibility of the world being underpopulated do not rely on a falling population. They show that population policy is best thought of not as maintaining the status quo, but as reaching an ideal level. Thus the world might be underpopulated even if the population is increasing.

Conclusions

As I have shown, a mature population policy cannot be constructed just from considerations of resource limits or the costs of additional people. We must consider the instrumental benefits that additional people will bring for the people who would have existed otherwise—especially from the creation of sharable information goods such as art, invention, and science. We must also consider the value of the additional people in and of themselves: their intrinsic value. There is active philosophical debate about exactly how we should take account of the intrinsic value of new people, but this does not mean that we can simply ignore this component of value. Instead, it means that there is considerable uncertainty in how we should assess the intrinsic value of differently sized future populations, and thus in how much we should add to their instrumental value in order to determine their overall value. However, even just properly accounting for their instrumental value alone may be enough to suggest that the planet might be underpopulated.

We have seen that the various limits on population due to resource constraints come in both soft and hard forms, and can be translated into the language of costs, to be weighed against benefits. We should also be aware that even the hard limits can potentially be raised with technological or social changes and we should consider attempting to raise limits as a very important policy option. We must also be aware that it can matter where the additional people are located: it is possible for some areas to have too many people while others have too few.

Most importantly though, we should stop looking at increasing population as just a problem to be managed, any more than our own lives are just problems to be managed. Like us, new people are the springs of great joy, novelty, and prosperity. The resources they consume may outweigh this, or they may not, but we certainly don't know the answer to this yet and would need to investigate much more fully. We should see increasing population as an opportunity as well a challenge.

References

Arrhenius, G. (forthcoming). *Population Ethics*. Oxford: Oxford University Press.

(Page numbers refer to unpublished manuscript).

Arrhenius, G. (2009). 'Can the Person Affecting Restriction Solve the Problems in Population Ethics?' in M. Roberts and D. Wasserman (eds), *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, London and New York: Springer, pp. 291–316.

- Arrhenius, G. (2000a). 'An Impossibility Theorem for Welfarist Axiology', *Economics and Philosophy*, 16: 247–66.
- Arrhenius, G. (2000b). 'Future Generations: A Challenge for Moral Theory,' FD-Diss. Uppsala: University Printers.
- Blackorby, C., Bossert, W., and Donaldson, D. (1997). 'Critical-Level Utilitarianism and the Population-Ethics Dilemma', *Economics and Philosophy*, 13(2): 197–230.
- Blackorby, C., Bossert, W., and Donaldson, D. (2005). *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. Cambridge: Cambridge University Press.
- Boserup, E. (1965). *The Conditions of Agricultural Growth: The Economics of Agrarian Change under Population Pressure*. London: Allen & Unwin.
- Broome, J. (1992). *Counting the Cost of Global Warming*. Cambridge: The White Horse Press.
- Broome, J. (2004). *Weighing Lives*. Oxford: Oxford University Press.
- Ehrlich, P. R. (1968). *The Population Bomb*. New York: Buccaneer Books.
- Evenson, R. E., and Gollin, D. (2003). 'Assessing the Impact of the Green Revolution 1960-2000', *Science*, 300(5620): 758–62.
- Glover, J. (1977). *Causing Death and Saving Lives: The Moral Problems of Abortion, Infanticide, Suicide, Euthanasia, Capital Punishment, War, and Other Life-or-Death Choices*. Harmondsworth: Penguin.
- Goldin, I. and Reinert, K. (2012). *Globalization for Development: Meeting New Challenges*, New Ed. Oxford, UK: Oxford University Press.
- Huemer, M. (2008). 'In Defence of Repugnance,' *Mind*, 117(468): 899–933.
- Narveson, J. (1967). 'Utilitarianism and New Generations', *Mind*, 76(301): 62–72.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Sen, A. K. (1981). *Poverty and Famines: An Essay on Entitlement and Deprivation*. Oxford: Oxford University Press.
- Sidgwick, H. (1907). *The Methods of Ethics* (7th edn). London: Macmillan.
- Simon, S. L. (1981). *The Ultimate Resource*. Princeton, NJ: Princeton University Press.
- Tännsjö, T. (2002). 'Why We Ought to Accept the Repugnant Conclusion', *Utilitas*, 14(3): 339–59.
- UN (2011). *World Population Prospects: The 2010 Revision*. New York: United Nations, Department of Economic and Social Affairs.
- UN (2004). *World population to 2300*. New York: United Nations, Department of Economic and Social Affairs.