Canadian Studies in Population, Vol. 27(2), 2000, pp. 355-376

How Many People should the Earth Hold? Framing Questions about Values

Thomas H. Wonnacott
Department of Statistical and Actuarial Science
University of Western Ontario
London, Ontario

Abstract
The number of people the earth *can* hold currently is estimated to be anything up to 30 billion or more. Just where within that huge range we should aim is not a question of scientific fact as much as a question about human values, and what criterion will lead to a wise choice. This paper is an invitation to look at that choice thoughtfully, and try to deal with one of the most elementary but vital issues: the questions that probe our values get different answers, depending on how we frame the questions – in terms of gain or loss, in terms of acts of omission or commission. Amid a wide diversity of views, however, one commands strong support: we should aim for a population where the average welfare $\overline{X}$ and the total number of people N are *both* large. It is a mistake to maximise just $\overline{X}$ alone or N alone.

Résumé
On estime que la terre peut soutenir actuellement jusqu'au 30 billion de personnes ou de plus. C'est une grande portée et notre population idéale s'agit plus d'une question de choix avisé et des valeurs humaines qu'un fait scientifique. Il faut réfléchir prudemment à ce choix-ci, de s'adresser aux questions les plus fondamentales et essentielles. Les questions qui sondent nos valeurs reçoivent des réponses diverses; cela dépend de comment on les pose en fonction des pertes et des avantages, ou en fonction de ce que l'on fait ou ne fait

pas. Cependant, au milieu des opinions diverses, il faut soutenir la meilleure politique; on doit avoir pour but une population où la moyenne du bien X et le nombre total des personnes N sont grands, tous les deux. C'est une faute de porter au maximum seulement X ou seulement N.

Key Words: average income, framing, optimum population, population decline

Introduction

This paper will deal primarily with the issue of human values, of how we should choose. But first perhaps we ought to mention the comparatively simple but dynamic issue of how many people *can* the earth hold – its "carrying capacity". This vague but crucial notion needs to be addressed because it determines the upper limit on our choice. There is fortunately a whole book devoted to this (Cohen 1995, especially Chapter 11.12) which details the enormous range of answers that have been worked out, from about one billion to more than 1000 billion, with most falling currently in the more believable range of 4 to 16 billion. This is still a very wide range, because the carrying capacity depends on (a) what level of prosperity we want, (b) what level of development (social and technological) we have, and (c) what level of carelessness we permit ourselves. To take extreme but illuminating examples, the carrying capacity could be far less than one billion if (a) we insisted that every person should have the right to unlimited experience of virgin wilderness, that is, be able to travel extensively where no other human has travelled ,or (b) we had the hunter/gatherer level of technology from 25,000 years ago, or (c) we made no attempt whatsoever to control pollution.

Of course, none of these conditions is widely recommended. Nearly everyone envisages some sort of accommodation to human numbers, to give us a population of several billion. But how much accommodation, how many billion? Should we aim at 10 billion in a hundred years? or 5, or 15 billion? The answer depends as much on the values we hold as on the facts we know.

Even for the facts, there is wide disagreement as to what they mean. To illustrate, in 1980 two protagonists – Paul Ehrlich and Julian Simon – who had the same historical facts on resource scarcity available to them, nevertheless had widely divergent views about how scarce natural resources would become in the future. They were so sure of their views in fact, that they bet on them, Simon betting that the scarcity of a representative group of minerals would decrease. (Incidentally, Simon won: the scarcity did indeed decrease, reflecting a surprising but historically supported trend of increasing availability. Tierney, 1990).

Or to take another example of how even the most carefully weighed views on population facts can differ, consider two successive reports of the National Academy of Science. In 1971 they reported that "rapid population growth had seriously negative economic consequences," while in 1986 they reported that "concern about the impact of rapid population growth on resource exhaustion has often been exaggerated," and cautiously concluded, "On balance, we reach the qualitative conclusion that slower population growth would be beneficial to economic development for most developing countries." Note that this just speaks of slower rather than slow population growth, it does not necessarily apply to all developing countries, and they further caution that "a simple model suggests that the effect is comparatively modest".

Is there any more of a consensus on values? Our values are examined so little that we hardly know what they are, much less how different they may be, or how they should inform our decisions. And one of the most troublesome problems in examining our values is that the framing of the questions influences the answers we get. This paper will therefore develop the framing paradox through several examples, and then suggest implications for how many people the earth should hold, or at least, in which direction our population ought to be heading.

Section 1—The Classical Framing Paradox

Example 1.The classical framing paradox of Kahneman and Tversky: The Flu Epidemic

The values we hold – technically, our utility functions – often depend on how the probing questions are *framed*. The following experiment from Kahneman and Tversky (1984) will illustrate. They asked one group of subjects (N_1=152) to decide between two programs to deal with a hypothetical flu epidemic. Then they put a second group (N_2=155) through the same exercise, except that the same two programs were *described* or *framed* differently, in terms of losses instead of gains. Here is what they were asked in detail (modified slightly in notation), with the proportion responding in favour of each program indicated in brackets:

> Preamble: Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:
>
> Frame 1 (Group 1,N=152):

> *If Program A is adopted, 200 people will be saved. (72%)*
> *If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved. (28%)*

Which of the two programs would you favour?

The second group had the same preamble, but then had the question framed differently:

<u>Frame 2 (Group 2,N=155):</u>

> *If Program A^1 is adopted, 400 people will die. (22%)*
> *If Program B^1 is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die. (78%)*

Which of the two programs would you favour?

You can easily verify that the options A and B in frame 1 are *logically* equivalent to the two options A^1 and B^1 in frame 2; the only difference is in their description or framing. Yet the choices the subjects made show a remarkable difference in frames 1 and 2. (72% preferred option A, while only 22% preferred option A^1).

What lessons can be drawn from this inconsistency? Although both frames showed exactly the same outcomes in their alternatives, the reference point – the point of view of the decision maker – was different. In Frame 1, it was a question of how many could be saved. For gains such as this, people are typically risk averse (they prefer the sure thing). In Frame 2, it was a question of how many lives would be lost. And there is a very human tendency to hate to lose, even a little. People will often pay relatively highly for a chance to completely avoid a loss, even if it means a greater expected loss – which makes them risk seeking (they prefer the gamble) in frame 2. So the decisions people make depend not only on the actual outcomes, but on how the choice is framed.

What implications does this have for decision making? Is there any way to make a decision that is somehow best, independent of the whimsies of framing? In this particular example, it might make sense to use the expected lives saved as the criterion, since all lives are about equally valuable[1]. However, for many decisions, for those involving money for example, this "actuarial" solution is not appropriate since all dollars are not equally valuable (the first million dollars you might win at a lottery would be much more valuable to you, in terms of what you would buy with it, than the second million dollars).

We can offer some implications for decision-making in general: It is important to be aware of how much our decision can be affected by how it is framed, by considering two or more different frames. Then some sort of compromise among these different frames might be taken, to try to make the best possible decision. Different frames are analogous to different replications in an experiment: the replications give us an idea of the uncertainty involved, and the average of the replications provides the best estimate of the parameter.

We shall next give four framing examples of our own, hoping eventually to illuminate some issues about population values. The subjects were 44 undergraduates taking an elementary statistics course, drawn mainly from Commercial Studies at UWO. They were broken up into two groups ($N_1 = 23$, $N_2 = 21$).

Example 2. Our more extensive treatment

This example used exactly the same two frames as the classical example of Kahneman and Tversky above, but asked for more extensive responses.

(i) Instead of a simple preference for A or B, the subjects were asked to answer how strong their preference was on a five point scale: For example, the first group of subjects were asked:

▼	▼	▼	▼	▼
definitely prefer A	slightly prefer A	no preference	slightly prefer B	definitely prefer B

This was compared to the preference (on the same five point scale) of the second group using a second frame.

We next asked the first group of subjects (after a few minutes distraction with other questions) to redo the decision now expressed in the second frame, and, similarly the second group of subjects to redo the decision now expressed in the first frame – a "crossover" design.

(ii) Finally, each subject was asked to look back on their two answers above, and revise them according to whatever insights had been provided by comparing the two frames. This protocol is laid out in Figure 1, with every

Frame 2 boxed to easily distinguish it. (The actual questionnaires handed out are available in Wonnacott 1996.)

Following the protocol of Figure 1, the results of the questionnaire are given in Figure 2, which show:

(i) The very top graphs of Figure 2 are of primary interest, because they show how the two different frames tried on two fresh groups of people elicit quite different answers (recognising that choice A and A^1 are logically the same, as are B and B^1, the only difference being their framing). This is the data directly comparable to what Kahneman and Tversky did.

To summarise the distribution for group 1 on the left, we note that of those who stated some sort of preference (all except the middle 6 subjects), the proportion who preferred A (slightly or definitely) was

$$(6+8)/ (6+8+2+1)= 82\%.$$

To summarise group 2 on the right, the comparable proportion who preferred A^1 (the logically equivalent choice, with the certain outcome rather than the uncertain one), dropped to a mere

$$(2+4)/(2+4+9+3) = 33\%.$$

These figures agree closely with the roughly comparable figures of 72% and 22% of Kahneman and Tversky quoted earlier, and thus confirm that people's ordinary risk-aversion changes to risk-seeking when outcomes are felt as losses.

(ii) Within each column (group), note that the last two (revised) distributions are quite similar to each other, indicating that the subjects by and large tried to reconcile the two frames consistently. And this final pair of distributions is quite similar to the initial frame distribution at the top, indicating that first impressions are hard to shake. That is, later choices tend to be "anchored" near the initial choice.

Section 2: Framing Three Population Scenarios

Having tested our protocol, and found it consistent with the classical results of Tversky and Kahneman, we proceeded to use it for three questions on population: These same subjects now were given three population scenarios – some far-out science fiction – to help focus on what criteria are appropriate for judging population size. In each scenario, the two frames were no longer exactly equivalent logically: They differed in that Frame 1 is about *saving* the

Figure 1
Protocol of Successive Frames given to the Two Groups of Subjects

PURPOSE	GROUP 1 SUBJECTS ($N_1 = 23$)	GROUP 2 SUBJECTS ($N_1 = 21$)
(i) Best comparison of two frames, using two fresh groups	FRAME 1	$\boxed{\text{FRAME 2}}$
To see to what extent subjects will give a different answer in a different frame	$\boxed{\text{FRAME 2}}$	FRAME 1
(ii) To see how subjects will finally decide after reconsidering both frames	FRAME 1 answer revised	$\boxed{\text{FRAME 2}}$ answer revised
	$\boxed{\text{FRAME 2}}$ answer revised	FRAME 1 answer revised

Figure 2. Results of the Asian Flu Decision, Using the 5 Point Scale

GROUP 1

frame 1, in terms of lives saved

prefer certainty ← → prefer gamble

	Row 1	Row 2 (boxed)	Row 3	Row 4 (boxed)
definitely prefer A	6 ******	3 ***	5 *****	5 *****
slightly prefer A	8 ********	7 *******	8 ********	6 ******
no preference	6 ******	3 ***	6 ******	6 ******
slightly prefer B	2 **	8 ********	3 ***	4 ****
definitely prefer B	1 *	2 **	1 *	2 **

GROUP 2

frame 2 (BOXED), in terms of lives lost

	Row 1 (boxed)	Row 2	Row 3 (boxed)	Row 4
definitely prefer A'	2 **	1 *	1 *	0
slightly prefer A'	4 ****	5 *****	5 *****	3 ***
no preference	3 ***	2 **	2 **	4 ****
slightly prefer B'	9 *********	10 **********	10 **********	12 ************
definitely prefer B'	3 ***	3 ***	3 ***	2 **

Note: Following the crossover protocol of Figure 1, the second row shows the frames reversed. Then the last two rows show the reconsidered vote on the first two rows.

right planet, while Frame 2 is about *creating* the right planet. But they were equivalent in terms of their consequences, that is, in terms of the planets that would be left in the universe after our choices are made.

In each scenario, furthermore, the subjects were told that after 10 years, an exploding star would wipe out all the alternatives. Thus they could express their preference for the lives as they were being currently lived, rather than having to speculate about how they might turn out in the long run. It kept the choices clearer.

Example 3: The voyage of the S.S. Endymion

Here is what the subjects were asked:

Frame 1 As captain of our Federation's magnificent new starship, Endymion, you have been commissioned to "seek out all forms of life, in the farthest reaches of the galaxy". Within two brief hours of setting out, you are astonished to find a tiny blue-green planet that is an exact 1999 clone of the western hemisphere – North America (NA) plus Latin America (LA).

As if that isn't remarkable enough, ten minutes later you come across a second blue-green planet that is a clone of North America, that is, exactly like the first planet except that Latin America is completely devoid of human population. Your fascination with it, however, is suddenly interrupted. Your first mate reports that a cosmic tornado is bearing this way to wipe out both planets, and your starship is not powerful enough to stop it. You can deflect it, however, and manage to save one of these planets:

Planet A.	*1999 populations of NA and LA*
Planet B.	*1999 population of NA alone.*

Which would you save?

After the subjects answered this, and had a few minutes distraction doing other questions in order to give them a fresh perspective, their questionnaire continued:

Now you are asked to make this decision all over again – in a second framework that looks quite familiar. But there are enough subtle differences (in **boldface**) to make it worthwhile answering afresh.

Frame 2. As captain of our Federation's magnificent new starship, Endymion, you have been commissioned to "seek out all forms of life, in the farthest reaches of the galaxy". Within two brief hours of setting out, you are astonished to find a tiny blue-green planet remarkably like earth. **Its world government is planning their future, and has narrowed the options to two possible plans for how the world might turn out:**

Plan A [1]	*1999 populations of NA and LA*
Plan B [1]	*1999 population of NA alone (leaving LA untouched and unpopulated).*

Being duly impressed by your spaceship, and your galactic reputation for wisdom, they insist on taking your advice about which plan to implement. Making the heroic assumption that plans work out unfailingly, and that both plans are equally costly to achieve, which would you recommend?

The results are given in Figure 3, which show:

(i) The distributions at the very top of Figure 3 once more indicate what a difference framing makes. The proportion preferring the larger population is 18/21 = 86% in frame 1, and drops to 7/15 = 47% in frame 2.

(ii) Again we note the anchoring to the first impression: in each column (group), the last two distributions are quite similar to the initial distribution at the top.

Example 4: The second voyage of the Endymion

The next two examples are much simpler, and will demonstrate that the preferences for choice A can be so strong that framing differences are relatively weak. Here is what the subjects were asked in Example 4:

Frame 1: You are again the captain of the Endymion, and must choose one of these two planets to save:

Planet A is a 1999 clone of Canada, a prosperous country of 28 million.
Planet B is a little larger country of 30 million, but ruled so badly that its level of misery is phenomenal: half the population are political prisoners in a forced labour camp, the other half are terrorised and impoverished.

Figure 3. The Endymion's First Voyage Choosing NA + LA or LA Alone

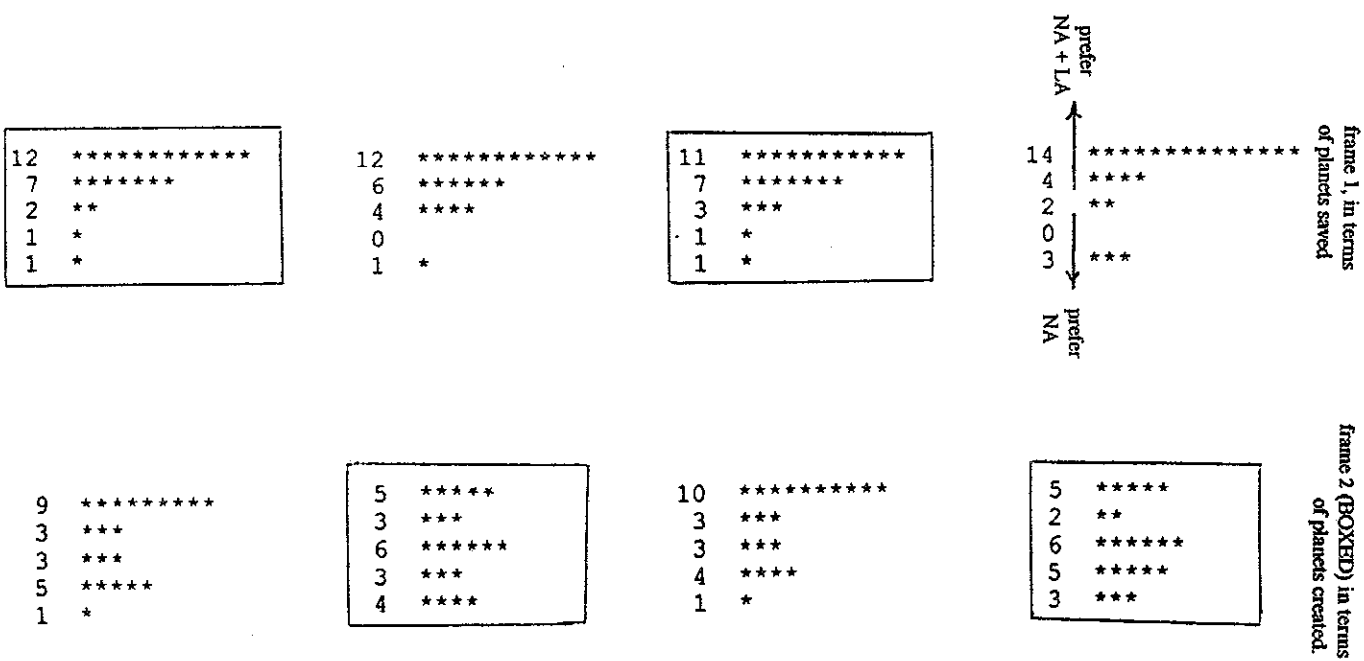

Note: Following the same protocol of Figure 2, the second row shows the frames reversed. Then the last two rows show the reconsidered vote on the first two rows.

Which would you save?

Frame 2 You are again the captain of the Endymion, **and must choose one of these two plans to be implemented:**

> ***Plan A^1*** *would produce a 1999 clone of Canada, a prosperous country of 28 million.*
>
> ***Plan B^1*** *would produce a little larger country of 30 million, but ruled so badly that its level of misery would be phenomenal: half the population being political prisoners in a forced labour camp, the other half terrorised and impoverished.*

Making the heroic assumption that plans work out unfailingly, and that both plans are equally costly to achieve, which would you recommend?

The results are given in Figure 4, where we see that choice A (or A^1) is preferred so strongly that the framing effect is relatively weak. In Figure 4 (and Figure 5 later) we show only the first row. The remaining three rows that appeared in earlier figures were omitted, because they were completely unsurprising: they showed the usual anchoring to the initial position in the first row. (Wonnacott 1996).

Example 5: The Third and Final Voyage of the Endymion

Here is what the subjects were asked:

Frame 1: Once again you are the captain of the Endymion. But now this last voyage occurs much earlier, when technology is very primitive. As captain, your decision is badly constrained by a lack of information. You don't have the amazing technology that allows you to perceive in detail what the individual lives on these planets were exactly like. Instead, you have a "statistical analyser" that only registers averages – while the number and character of the individual lives remain unknown. On the basis of these averages alone – since no other information at all is available – you must choose which planet to save:

<u>AVERAGE LEVELS OF WELFARE</u>

	average income (per capita GNP)	*average length of life (life expectancy)*
<u>Planet A</u>	*$24,000*	*76 years*
<u>Planet B</u>	*$11,000*	*72 years*

Which planet would you save?

<u>Frame 2:</u> Once again you are the captain of the Endymion, **and must choose one of two plans to be implemented. This last voyage occurs much earlier, when the world government's plans are much more vague. They no longer envisage their world in exact detail, down to a specifying the individual lives. Instead, they just have target averages for a couple of crucial welfare indicators:**

Figure 4
The Endymion's Second Voyage
Choosing the Smaller Population N or Larger N

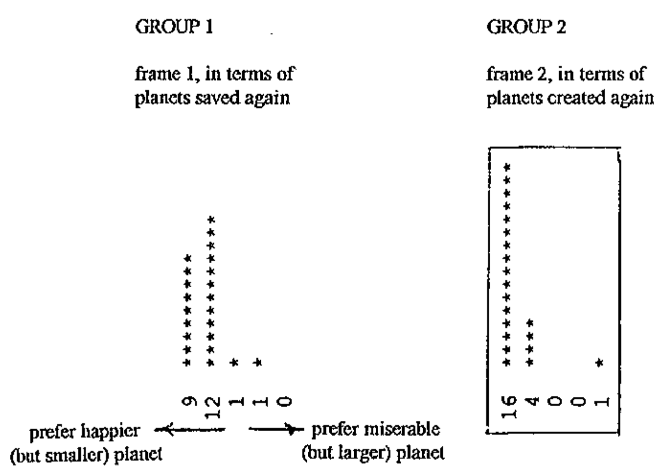

	AVERAGE LEVELS OF WELFARE	
	average income (per capita GNP)	*average length of life (life expectancy)*
Planet A	*$24,000*	*76 years*
Planet B	*$11,000*	*72 years*

Making the heroic assumption that plans work out unfailingly, and that both plans are equally costly to achieve, which would you recommend?

The results are given in Figure 5, and are very similar to Figure 4: Choice A (or A^1) is preferred so strongly that the framing effect is relatively weak. But if we look at the *degree* of preference (definite versus slight), there is still a discernible framing effect there.

Next we shall draw some important conclusions from these examples.

Section 3: Aversion to Loss, and to Acts of Commission

Loss aversion, Background

Tversky and Kahneman (1978) have shown how strong is our human tendency to hate to lose, what they call *loss aversion*. We are too averse to loss for our own good. As Robyn Dawes (1988) explains, "people will get themselves into a mess to avoid a serious loss." As a financial example, people too often hold on to an asset – a house or a stock – trying to recover the price they paid for it, looking back to avoid losses instead of conducting a rational analysis of future gains and losses.

Thaler (1983) has analysed an aspect of loss aversion that affects public policy, what he calls the "endowment effect … once people have something, it is very hard to take it away. Residents of communities with declining school populations know how that has created problems for their school boards. People who would be unwilling to pay for a tax increase to add a school in their neighbourhood nevertheless become incensed if an existing school in their neighbourhood is closed in order to avoid a tax increase."

For readers interested in the sociobiological source of loss aversion, some clues are given in Dawkins (1976) and Wright (1994).

Examples 1 and 2: Risk - Seeking Behaviour for Losses Related to Utility Functions

Loss aversion is closely related to utility functions and risk-seeking in statistical decision theory. Let us summarise the argument in Kahneman and Tversky (1984). Typically people are risk averse in trying to gain assets in risky situations such as stock markets. This is reflected in their utility function or "value curve" being concave, as the curve "a" shows in Figure 6 (explained fully in Raiffa 1968, for example).

One would expect the typical utility function to continue left smoothly on the loss side as in curve a^1. However, humans hate to lose, and seem to suffer an extra psychological loss shown as the vertical stripes. This brings the "value curve" down from a^1 to the curve b, which is convex now – reflecting risk-seeking behaviour on the left of the origin.

Example 3: Acts of Omission are Less Serious

Loss aversion has repercussions for our social conscience as well. The destruction of something (act or sin of commission) tends to weigh more heavily on our conscience than the mere failure to create it (the sin of omission). Eight of the ten commandments of Moses in the Bible, for example, are commandments that forbid various acts or sins of commission: "Thou shalt not kill, Thou shalt not steal…".

To see how this is related to Example 3, let us compare the very top distributions in Figure 3. On the left, a strong majority (18 to 3, ignoring the 2 middle votes) voted to save the populous planet (avoid the sin of commission). On the right, however, it was not quite a majority (7 to 8, ignoring the 6 middle votes) who voted to create the populous planet (avoid the sin of omission). So the sin of omission is perceived as less serious.

Section 4: Some Consequences for Demographic Values

Review of the three demographic examples

<u>Example 3 (First voyage of the Endymion).</u>

In Figure 3, let us look at just the 23 subjects in Group 1, in the left hand column. A large majority preferred the populous planet, whether expressed in frame 1 at the top (sins of commission) or in frame 2 just below it (sins of omission). And when they reconsidered their views in the last two panels on the left, they become a stronger and even more consistent majority preferring the populous planet.

It is from their view, what we will call the "majority view", that we will shortly pursue the ethical implications. We recognise they do not represent everybody, so our arguments will not be of universal relevance. But this majority is an important segment who make consistent choices, from which we shall next draw some very clear and interesting conclusions.

Example 4.

This was intentionally set up to be straightforward. There is practically unanimous agreement shown in Figure 4: It is preferable to have a happier planet, even if the population is smaller.

Example 5.

In Figure 5 we find almost as strong agreement as in Figure 4: It is preferable to save the more prosperous planet (with high average income and length of life).

Figure 5
The Endymion's Third and Final Voyage
Choosing the Larger or Smaller Average Level of Welfare

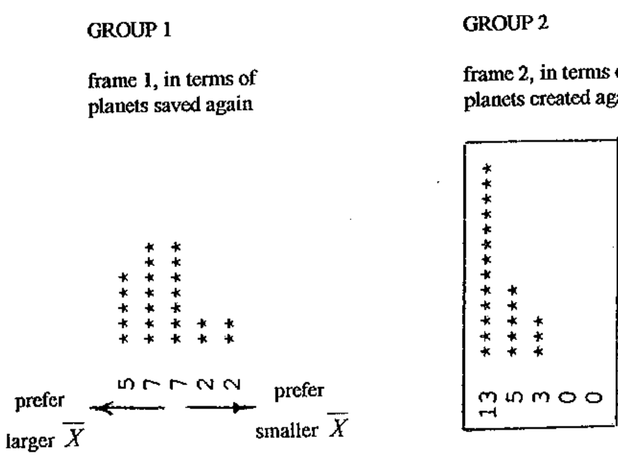

Two basic principles

Now we shall draw some implications, some principles for guiding our decision about how much or how little earth's population ought to grow.

Let us first get something obvious out of the way. The lesson from Example 4 is that to decide which population outcome is better, maximising the sheer number of people N just won't do; we *must* consider how happy they are – of course. So we state this as our first and very obvious principle, in deciding which population is best:

$$\boxed{\text{Considering N alone is not enough}} \qquad \text{(P1)}$$

In Example 3, what lessons can we draw from the "majority view" (that LA was worth including)?

Figure 6
Loss Aversion makes the Value
or Utility Function Convex on the Left of the Origin

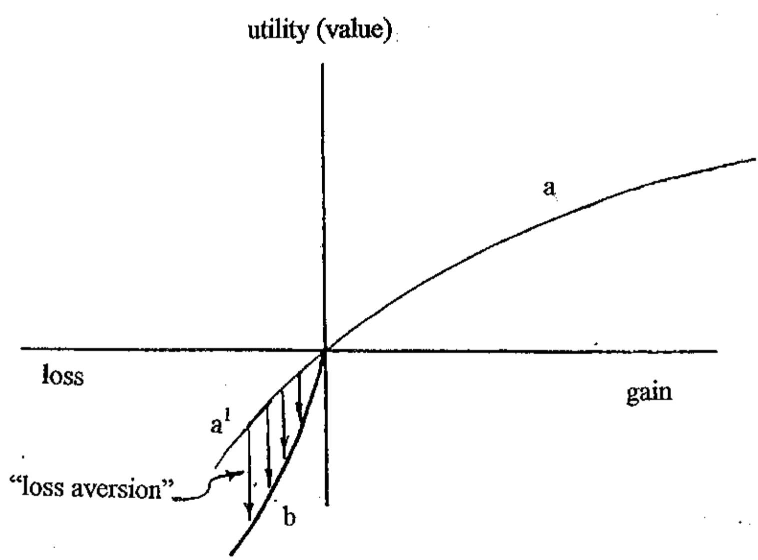

In the many discussions I have enjoyed with students, the typical reason for saving (or creating) the larger population ran something like this: " Save the larger population, because it has exactly the same North America that the other planet has, *and in addition*, you will be saving a Latin America too. Why throw Latin America away? (Of course, if the Latin American lives on the whole were so miserable they wished they were dead, this decision would be reversed. But from everything we know of Latin America, only a dedicated pessimist could assume that).

This answer might be stated more generally in what we will call ...

The <u>"simple addition principle":</u> If the only difference is to add on something that is itself good on the whole (that you would choose as worthy in its own right, like Latin America), then choose it.	(P2)

A Tragic Disclosure in Example 5

Now we shall reveal something in Example 5 (the last voyage) that was not known to the captain – because he was in a relatively primitive starship that could only disclose averages, and was not capable of picking up the detailed lives of the inhabitants.

Imagine we are back on the bridge of the Endymion on its last voyage. The captain (like the great majority of our subjects) has just made the decision to save planet A because it is more prosperous. After the switch is thrown to veer the cosmic cyclone towards planet B however, the first mate finds he can get a closer view of the two planets after all. In fact he can see the same detail as in the first voyage, and he discovers to his horror that planet B is a clone of North America plus Latin America. (In fact it was the inclusion of LA that brought down its averages – average length of life and income). And this of course is the planet that should have been saved (assuming, as we said, the "majority view").

It is a tragedy caused by ignorance. When all the captain had to go on was averages $\bar{X}$, he chose the best he could – to save the most prosperous planet. But full information, if only it had been provided in time, would have correctly lead him to save the other planet, because it included LA too.

In other words, all those fine looking averages were just not enough to make the wise choice. We might generalise this as...

<u>The "Integrity Principle"</u> To make a wise decision, it is important to get the whole picture.	(P3)

We might state it even more forcefully, in a form appropriate to teach to every class in demographic statistics:

Half the picture, even supported by the best statistics in the world, can lead to a tragically wrong decision.

Or, in a very brief and specific form:

Considering $\bar{X}$ alone is not enough.

(P4)

The crucial missing information was the population size N. Instead, we were just given the measures related to *average* happiness such as average income or average length of life. In the feast of life that each planet provides, so to speak, it matters how many people (N) are enjoying it, as well as the size of the average plate ($\bar{X}$).

Of course, the second voyage showed the mirror image of this: considering the population size N alone is not enough either. So N and $\bar{X}$ are both crucial, and we might state this as our final principle. Combining (P1) and (P4),

N and $\bar{X}$ are both crucial for a wise decision on population size.

(P5)

Demographic Implications

Of course, N and $\bar{X}$ are seldom independent: a policy to increase one will likely change the other too. Ideal policies that increase both N and $\bar{X}$ may be next to impossible to find. What we can realistically look for, however, are policies that increase one without decreasing the other as much. For example, population growth reduced to moderate levels may increase $\bar{X}$ more than it decreases population size N. This overall benefit is what drives many population policies in the least developed countries. In nearly all developed countries, however, we find an immanent decrease in population N that may not produce any corresponding increase in average welfare $\bar{X}$. So this would seem to be a decidedly suboptimal "policy".

Principle 5 has another surprising implication. We might ask, for example, what criterion is used by those working in the field of development and population policy, to decide which outcomes are better for a country? The vast majority consider only *average* welfare or happiness (measuring it with the same available proxies that we used – average income usually. See National Academy of Sciences 1986, p. 7.) They are forgetting that population size is not only a control variable, but a criterion variable as well.

Now the starship Endymion had every excuse for using averages as the sole criterion in its third voyage. Population size N was completely unknown, which lead them to a tragic choice – through no fault of theirs. But what excuse do we social scientists have for ignoring N, when it is about the most readily available statistic for every country in the world? Why have we uncritically accepted the common statistical emphasis on averages, and failed to think through the implications of ignoring N?

One of the reasons may be that our ability to control population size is a relatively recent phenomenon and we have not yet recognised how revolutionary are its implications. When population size was fixed or exogenous, then whatever policies increased average welfare would automatically increase total welfare, and vice versa; that is, average and total welfare were equivalent criteria. Now however, with population size coming more and more under our control, using average welfare as our criterion would lead to a population size smaller than using total welfare – a population much too small.

Checking it Out.

We have just reached an astounding conclusion: the almost universally accepted criterion of progress – increasing average welfare alone – may not be appropriate after all; in fact, it can produce terrible decisions (like the tragic disclosure showed in Example 5). The implications for population policy are so far reaching that we surely ought to check it out. So we shall give a small sample of the philosophical arguments that can be made for our conclusion (P4) – referring the reader to Barry and Sikora 1978, or Ng 1986, for a much more complete discussion:

(i) Since looking forward to the future is so terribly uncertain, let us look backwards and rerun history for some insights. What would have happened if the world had taken Zero Population Growth (ZPG) seriously 200 years ago in Malthus' day, so that the world's population levelled out at 1 billion, instead of rising to 6 billion? There would now be only one sixth as many of us enjoying life[2]– only one sixth the laughter and one sixth the tears. But for those of us who feel that the laughter outweighs the tears, so that life on the whole is good, would that levelling off have been a triumph or tragedy?

(ii) To take a more extreme case, suppose the world's population growth over the past 200 years was not just ZPG, but a drastic decline to a mere hundred people today, and suppose their level of happiness was a little higher than the average level of the six billion currently alive, say at the 67^{h} percentile rather than the 50th.

By the criterion of maximising average happiness, this tiny group would have to be considered preferable to the current world's population of six billion. But by the "majority view" (that regards sins of omission and commission to both count), this tiny population would be a disaster. After all, the happiest 2 billion (top third) of the current world's population would be missing. And many of the other 4 billion who are missing would still have lives worth living. It would be a terrible violation of the simple additional principle P(2).

Conclusion

People hold a wide diversity of views, and even the same person could arrive at either of two different views depending on how the issue is first framed. Among all this diversity, however, on some issues there is a majority view (a majority among the participating subjects). This majority view implies the principles (P1) to (P5) that lead us to seriously question the usual criterion of maximising *average* happiness alone, without regard to the *numbers* who are enjoying that happiness.

Acknowledgements:

For their inspiration and discussion, many thanks to the participants at the Canadian Population Society meeting at Brock University in June 1996, and to a referee, and especially to three colleagues at UWO: Rod Beaujot, Tom Burch,
and Kevin McQuillan.

End Notes:

(ii) Kahneman and Tversky suggested calculating the expected number of lives saved (or lost – in either formulation, we will get exactly the same conclusion). That criterion is not only eminently sensible, it gives us a unique answer free of the framing paradox: Programs A and B (or A^1 and B^1) are *equally* good, having 200 lives expected to be saved (or 400 expected to be lost).

(iii) And we would probably be enjoying life at roughly the same level of prosperity. As we mentioned at the beginning of the paper, the review by the National Academy of Sciences (1986) showed that there is very little hard evidence one way of another about the relation of prosperity (average income, length of life, etc.) to population growth. Extremes of very high growth, or negative growth, seem to be harmful. But the kind of moderate growth the world has experienced over the past 200 years – averaging about 1% -- seems likely beneficial. After all, life in terms of average income, or reduced mortality, is vastly better today than 200 years ago. (In this ultimate "historical" or "observational" study, however, we have to recognize how soft our conclusions must be.)

References:

Barry, B. and R.I. Sikora (eds.) 1978. *Obligations to Future Generations.* Philadelphia: Temple University Press.

Cohen, Joel. 1995. *How Many People Can the Earth Support?* New York: Norton.

Dawes, Robyn. 1988. *Rational Choice in an Uncertain World.* New York: Harcourt Brace Janovitch.

Dawkins, Richard. 1976. *The Selfish Gene.* New York: Oxford University Press.

Kahneman, Daniel and Amos Tversky. 1984. "Choices, Values and Frames," *American Psychologist* 39: 341-350

National Academy of Sciences, Committee on Population. 1986. *Population Growth and Economic Development: Policy Questions.* Washington D C: National Academy Press.

Ng, Y-K. 1986. "On the Welfare Economics of Population Control," *Population Development and Review* 12: 247-266.

Raiffa, Howard. 1968. *Decision Analysis.* Reading, Mass: Addison-Wesley.

Sumner, L.W. 1978. "Utilitarianism and the Population Optimum" in Barry and Sikora, (Eds.). *Obligations to Future Generation.* Philadelphia: Temple University Press. Pp. 91-114.

Tierney, J.1990. "Betting the Planet", *New York Times Magazine,* 2 December 1990.

Thaler, R.H. 1983. "Illusions and Mirages in Public Policy," *The Public Interest* 73: 60-74.

Wonnacott, T.H. 1996. "How Many People Should the Earth Hold? Framing Questions about Values," Research Paper TR-96-4, Department of Statistical and Actuarial Science, University of Western Ontario.

Wright, Robert. 1994. *The Moral Animal.* New York: Random House.

Received July 1999; Revised July 2000