Chapter 9

Population Growth and Economic Development

9.1. Introduction

The world is populated today as it has never been before. Although rates of population growth have fallen and will continue to fall, we currently add about a million people every *four* days to the world population, *net* of deaths. According to projections carried out by the United Nations, annual additions to the population are likely to remain close to the ninety million mark until the year 2015.

It took 123 years for world population to increase from one billion (1804) to two billion (1927). The next billion took 33 years. The following two billions took 14 years and 13 years, respectively. The next billion is expected to take only 11 years and will be achieved by 1998, at which time we will arrive at the staggering figure of six billion. Such is the power of exponential growth.

However, more than just exponential growth is hidden in this story. Population growth through the millennia has not proceeded at an even exponential pace. The growth rate of population has itself increased, and the trend has reversed only in the last few years. Part of our purpose in this chapter is to tell this complex and interesting story.

Yet a description of trends is not our only purpose, because this is a book about economics, not demographic statistics. We are interested primarily in how the process of development has spurred (or retarded) population growth and, more important, we want to know how population growth in turn affects economic development. As with the evolution of now-familiar variables such as per capita income and economic inequality, population and development are intertwined, and we seek to understand both strands of the relationship.

The question of how population growth affects development runs into an immediate difficulty. How do we value the lives of the people yet unborn? Is a small population living in luxury better off than a large population living under moderate circumstances? How do we compare the fact that a larger number of people are around to enjoy the "moderate circumstances" with the alternative in which luxuries are available to a smaller number, simply because the births of the rest were somehow prevented?

This is a difficult question and we do not pretend to provide an easy answer. Indeed, we simply sidestep this issue by using per capita welfare (and its distribution) as our yardstick. The implicit ethical judgment, then, is that we are "neutral" toward population: once someone is born, we include that someone as worthy of all the rights and privileges of existing humanity. At the same time, our focus on per capita welfare means that we are indifferent to the unborn and are even biased toward keeping population growth down if it affects per capita welfare adversely.

This ethical judgment is implicit in the dire warnings that we see all around us, especially in developed countries where population growth in the "Third World" appears most frightening. Population growth cannot be good. It eats into resources and into production. There is less per head to go around.

That is fair enough. We adopt the per capita perspective as well. However, this does not imply that we need be averse to population growth from a *functional* viewpoint. The existence of a population of nontrivial size may have been essential to many important advances to the world. It is unclear how much Robinson Crusoe would have accomplished on his own, even with the help of his man Friday. For one thing, there are limits to what one or two brains can think up. For another, necessity is the mother of invention, and without the pressure of population on resources, there may be no necessity and consequently no invention. Just how large population needs to be for the full realization of these salubrious effects is open to debate, but the point remains that the total quantity of available resources may itself be positively affected through population growth.

The doomsday predictions associated with population growth also have a particular slant to them. On the heels of the (perhaps defensible) feeling that population growth is unambiguously bad for humanity, there is also the observation, sometimes made with a great deal of sophistication, that unless we do something about population growth in developing countries, the world will somehow be unbalanced in favor of the peoples of these countries. That would be "unfair."

Both of these misconceptions are, to some extent, unfounded. Moreover, taken to extremes, they can be dangerous. However, clearing up misconceptions is not our main goal. These statements are corollaries of more serious questions regarding the interaction of population growth and economic development that we shall address in this chapter.

(1) What are the observed patterns of population growth across different countries and how do these patterns correlate with other features of development in these countries? Specifically, is there a close relationship between

what the now-developed countries have demographically experienced in the past and what is currently being experienced by developing countries? This will take us into a discussion of the *demographic transition*, a phenomenon you were introduced to briefly in Chapter 3.

(2) What connects these societywide patterns in population growth to the decisions made by individual households regarding fertility? What features of the social and economic environment affect these household-level decisions? In particular, how does economic development affect fertility choices?

(3) Can observed household decisions regarding the number of children be "rationalized" by the environment in which they find themselves? Alternatively, do households have more children than is good for them? This is a difficult question that we must address at two levels. The first level is what might be called the "internal level": given some economically rational level of fertility *at the level of the couple*, do couples systematically depart from this level, either because of miscalculation or because of the absence of effective contraception? The second level is "external" and comes from pondering the meaning of the italicized phrase in the previous sentence. Are there reasons to believe that a couple's decisions regarding family size have a social impact that is *not* fully internalized by them?

(4) Finally, reversing the causality from economics to demography, is it unambiguously true that population growth is harmful to the economic development of a country? What explains the interesting dichotomy between the belief that world population growth is "bad" and the belief, so widespread in developed countries, that population growth will make "them" powerful at "our" expense?

We do not pretend to have comprehensive answers to all these questions, but you will certainly find some of the issues that we discuss very provocative and worthy of further study. However, before we begin a serious discussion, it will be useful to review some basic concepts and terminology that are used by demographers. This is the task of the next section.

9.2. Population: Some basic concepts

9.2.1. Birth and death rates

To conduct a useful analysis of population and its interaction with economic development, it is necessary to understand a few basic concepts and terms. Most of what we study in this section are just definitions, and with a little patience, they are very easy to understand. These definitions set down the language in which we discuss demographic issues.

Fundamental to the study of population is the notion of *birth rates* and *death rates*. These are normally expressed as numbers per thousand of the population. Thus, if we say that the birth rate of Sri Lanka is 20 per 1,000, this means that in each year, Sri Lanka adds 20 newborn babies for every thousand members of the population. Likewise, a death rate of 14 per 1,000 means that in each year, an average of 14 people die for every 1,000 members of the population.

The *population growth rate* is the birth rate minus the death rate. Even though this works out as a number per 1,000 (6 in our example above), it is customary to express population growth rates in percentages. Thus, the population growth rate is 0.6% per annum in our example.

Table 9.1 provides us with data on birth rates, death rates, and population growth rates for selected low-income, middle-income, and high-income countries. There is a cross-sectional pattern here that we will take up in more detail when we study the demographic transition, but certain features come to mind.

First, *very* poor countries such as Malawi and Guinea-Bissau appear to have both high birth rates *and* high death rates, ranging around 50 per 1,000 for births and 20 per 1,000 for deaths. This is Group I in the table. Countries in Group II are not as poor: their death rates are much lower relative to the Group I countries, but their birth rates are still high. This isn't uniformly true of all poor countries though: some, such as India and Bangladesh (Group III), seem to have begun a fall in birth rates that is gathering momentum. Other relatively poor countries, such as China and Sri Lanka (Group IV) have already taken significant strides in this direction: both birth and death rates are low and getting lower. Group V lists some Latin American countries, where the experience is mixed: countries such as Guatemala and Nicaragua have (like the Group II countries) benefited from the drop in death rates, but the accompanying fall in birth rates has not yet occurred. Countries such as Brazil and Colombia are well into the process, as are East and much of Southeast Asia (Group VI): countries such as Korea and Thailand have very low birth and death rates (others, such as Malaysia, *have not completed this process*).

Table 9.1 is constructed very roughly in ascending order of per capita income. The following broad trend appears: at very low levels of per capita income, both birth and death rates are high. Indeed, this is probably an understatement: *age-specific* death rates are probably higher still (see following text). Then death rates fall. This is finally followed by a fall in the birth rates. We will see this much more clearly when we track a single country over its history.

Now for a different concept. It is worth understanding that aggregative figures such as birth rates and death rates, and especially population growth

Table 9.1. Birth and death rates (1992) and population growth rates for selected countries.

Country	Per capita income	Birth rate	Death rate	Population growth rate (%)
I.				
Mali	520	51	20	3.1
Malawi	690	51	20	3.1
Sierra Leone	750	49	25	2.4
Guinea-Bissau	840	43	21	2.2
II.				
Kenya	1,290	45	12	3.3
Nigeria	1,400	45	15	3.0
Ghana	1,970	42	12	3.0
Pakistan	2,170	41	9	3.2
III.				
India	1,220	29	10	1.9
Bangladesh	1,290	36	12	2.4
IV.				
China	2,330	18	7	1.1
Sri Lanka	2,990	21	6	1.5
V.				
Nicaragua	1,900	41	7	3.4
Peru	3,220	27	7	2.0
Guatemala	3,350	39	8	3.1
Brazil	5,370	25	8	1.7
Colombia	5,490	24	6	1.8
VI.				
Thailand	6,260	19	6	1.3
Malaysia	7,930	29	5	2.4
Republic of Korea	9,630	16	6	1.0

Source: *World Development Report* (World Bank [1995]) and *Human Development Report* (United Nations Development Programme [1995]).

rates, hide significant information about the underlying "demographic structure" of the country.

For instance, two countries with the same population growth rates may have dramatically different age structures. This is because one of the two countries (call it *A*) may have a significantly higher birth rate *and* a significantly higher death rate than the other country (*B*) (so that the two cancel out in the comparison of net population growth rates). At the same time, it is true that country *A* is adding more young people to its population than country *B*. Unless the higher death rates in country *A* are entirely concentrated among the young, which is unlikely, there will be more young people in *A* than in *B*. We might then say that country *A* has a "younger age distribution" than country *B*. As we will soon see, age distribution plays an important role in determining overall birth and death rates.

9.2.2. *Age distributions*

The age distribution of a population is given by a list of proportions of that population in different age groups. Table 9.2 gives us the age distribution of populations in different parts of the world, as of 1995. It is apparent from the table that the age distribution of developing countries is significantly younger than in their developed counterparts. I have never met a person who failed to be amazed by these figures when seeing them for the first time, and you will be too. The developing world is very young.

Just as birth rates and death rates affect age distributions, these rates are in turn affected by the age distribution prevailing at any particular moment in time. An aggregate birth rate is the outcome of the age distribution in a country, the age-specific fertility rates of women in that country, and the fraction of the population in different age groups. Similarly, the aggregate death rate is a composite that comes from age-specific death rates in a particular country, as well as the overall age distribution in that country.

These observations have important implications, as we will see. At the moment, let's pursue the more disaggregated view a bit further. An *age-specific fertility rate* is the average number of children per year born to women in a particular age group. The *total fertility rate* is found by adding up all the age-specific fertility rates over different age groups: it is the total number of children a woman is expected to have over her lifetime. In developing countries, this number can be as high as 7 or 8, and often higher. In the typical developed country, this number is 2, perhaps lower.

Of course, high total fertility rates contribute to a high birth rate, but from our discussion, it should be clear that the total fertility rate is not the *only* factor that determines the overall birth rate. In a country with a young age distribution, the birth rate can be significantly high, even if the total fertility rate is not. This is simply because the younger country has a larger percentage of the population in their reproductive years.

A parallel observation holds for death rates. Young populations are biased toward low death rates, and this is true even if age-specific death rates are high. It is worth noticing that even though most developing countries have higher death rates in each age group relative to their developed counterparts, these differences are not adequately reflected in the overall death rates, which lie far closer together. Indeed, it is perfectly possible for country *A* to have higher age-specific death rates at *every* age group than country

Table 9.2. Age distribution of the world population.

Region	*Population (millions)*	*0–15 %*	*15–64 %*	*65+ %*
World total	**5,716**	**32**	**62**	**6**
Africa	**728**	**44**	**53**	**3**
Eastern Africa	227	46	51	3
Middle Africa	82	46	51	3
Northern Africa	160	39	57	4
Southern Africa	47	37	58	5
Western Africa	210	46	51	3
Latin America	**482**	**34**	**61**	**5**
Caribbean	35.0	29	63	6
Central America	126.0	38	58	4
South America	319.0	33	62	5
Asia	**3,457**	**32**	**63**	**5**
Eastern Asia	1424	25	68	7
South-Central Asia	1381	37	59	4
Southeast Asia	484	35	61	4
Western Asia	168	38	58	4
North America	**292**	**22**	**65**	**13**
Europe	**726**	**19**	**67**	**14**
Eastern Europe	308	21	67	12
Northern Europe	93	20	65	15
Southern Europe	143	17	69	14
Western Europe	180	17	68	15
Oceania	**29**	**26**	**64**	**10**
Australia and N. Zealand	21.6	22	67	11
Melanesia	5.8	39	58	3
Micronesia	0.4	—	—	—
Polynesia	0.5	—	—	—

Source: *Demographic Yearbook* (United Nations [1995]).
Note: Individual figures may not add to total because of rounding error.

B, and yet have a lower death rate overall. This is the effect of a young age distribution at work.

Thus high rates of population growth lead to a younger population, and then on to high birth rates and low death rates. This creates an "echo effect" that keeps population growth high.

One important consequence of this observation is that population growth possesses an enormous degree of inertia. Imagine that a country that has had high population growth rates implements a policy to bring down total fertility rates. The point is that *even if this policy were to be successful, population size would probably overshoot the desired limits before settling down at an acceptable* *level*. The reason is simple. High population growth rates in the past lead to a young age distribution. A relatively large fraction of the population continues to be at the age where they are just about to marry and have families. Even if the total fertility rates were reduced the sheer numbers of young people would lead to a large number of births, viewed as a fraction of the *entire* population. This is the grim inertia of population growth, and more than one country has found, to their dismay, that even with the best intentions and implementation, bringing population growth to a halt is a bit like bringing an express train to an emergency stop.

9.3. From economic development to population growth

9.3.1. The demographic transition

Like economic growth, population growth is a modern phenomenon. Indeed, even if we were to know very little about the world, we could deduce this very quickly by regression in time. The world population today stands at around six billion. Let's go backward and *decrease* this number by 2% per year. This exercise would yield a population of 250,000 around 500 years ago, or a population of 10 around 1,000 years ago! This is obviously ludicrous, as the data at the beginning of this chapter indicate. This proves that population growth at around 2% per year is a phenomenon of recent vintage.

The first point to note is that the "carrying capacity" of the world was enormously different in the Stone Age than in the era of agriculture, and considerably lower than it is now. With shallow digging implements and imperfect acumen in the art of agriculture, people were confined to river basins. Starvation was common, as was early death due to a myriad of causes. The advent of agriculture changed all that, or much of that at any rate. With an increase in the carrying capacity of Mother Earth came an increase in population, but net growth was still minimal, because death rates were high and persistent. Famine continued to be commonplace, as were episodes of plague, pestilence, and war. As late as in the eighteenth century Malthus [1798] wrote of God's checks and balances to the sexual energies of women and men. A spontaneously high rate of reproduction was countered with all manner of disasters, such as regular outbreaks of plague, pestilence, and famine. So although birth rates were high, death rates were sufficiently high to keep growth rates down to a crawl. We may think of this as the *first phase* of demographic history.

A major change, however, was taking place, possibly even as Malthus was recording the grim retributions of Nature. With the advent of sanitation methods and increases in agricultural productivity, death rates began to fall around 1700, and the rise in industrial productivity sent Europe into a veritable population explosion. Table 9.3 gives you some idea of this.

The population explosion would not have taken place, of course, had birth rates simply followed death rates on their downward course without any time lag. However, this did not happen, and for two reasons. First, the very forces that caused death rates to decline also caused economic productivity to increase. For instance, the rise in agricultural productivity meant not only that there was a lower incidence of famine (thus bringing down death rates), but also that the overall carrying capacity of the economy in normal times increased. With room for a larger population, the Malthusian restraints were loosened and the urgency to bring down the birth rate therefore dissipated. Second, even if the forgoing scenario had not been the case, birth rates would probably still have been high because of the inertia that characterizes fertility choices made by households. This inertia is so important in our understanding of population trends that we will devote a fair amount of space to it in the next section. For now, we merely note that birth rates remained high even as death rates fell. This meant that population growth rates rose in this epoch, which we dub the *second phase* of demographic history.

Finally, birth rates fell as time overcame inertia, and as the population of the world rose to fill newly created carrying capacity. Population growth rates declined, until they fell to their present level in the developed world, which is around 0.7% per year. This is the *third and final phase* of demographic history.

These three phases jointly make up what is known as the *demographic transition*. Together, they paint a picture that almost all European and North American regions have seen: an increase and then a decline in the rate of population growth, changing the regime from one of high birth and death rates to one of low birth and death rates. Developing countries are going through the very same three phases, and doing so at an accelerated pace, as we will see. Almost all the countries of the world can be described as currently either in the second or the third phase of the transition.

9.3.2. *Historical trends in developed and developing countries*

It is of the *utmost* importance to understand that starting from around 1700 until well into this century, the populations of Europe and North America (most of the modern developed world) grew not only in absolute terms, but also relative to the peoples of those regions we know today as the developing world. To see this shift in population, it is useful to take a long-term view. Table 9.3 is taken from a revised estimate of the world population over the last few centuries (Carr-Saunders [1936]). We append to this table the 1995 estimates from the United Nations *Demographic Yearbook*. The results are very interesting.

Table 9.3. Geographical distribution of the world population.

	1650	*1750*	*1800*	*1850*	*1900*	***1933***	***1995***
World population (millions)	**545**	728	906	1,171	1,608	**2,057**	**5,716**
Percentages							
Europe	**18.3**	19.2	20.7	22.7	24.9	**25.2**	**12.7**
North America	**0.2**	0.1	0.7	2.3	5.1	**6.7**	**5.1**
Oceania	**0.4**	0.3	0.2	0.2	0.4	**0.5**	**0.5**
Latin America	**2.2**	1.5	2.1	2.8	3.9	**6.1**	**8.4**
Africa	**18.3**	13.1	9.9	8.1	7.4	**7.0**	**12.8**
Asia	**60.6**	65.8	66.4	63.9	58.3	**54.5**	**60.5**

Source: Carr-Saunders [1936, Fig. 8] and *Demographic Yearbook* (United Nations [1995]).

The table is constructed to emphasize the earlier centuries. Neglect the last column for the moment. What we have then is an array of population percentages running all the way from 1650 to 1933. Note the decline of Africa, in significant part due to outmigration, and the rise of North America, in large part due to immigration. At the same time, despite outmigration from Europe, her share of the world's population rose steadily over this period. Focus on the first column and the second to last column (both in boldface type) to see how the situation altered over the period 1650–1933. What we see here is the period when Europe began its demographic transition, while large parts of the present developing world still lay dormant in the first phase of demographic history.[1] In 1650, the population of Europe was about 100 million. In 1933, even allowing for emigration (which was large), it had swelled to over 500 million.

Now look at the last column of Table 9.3, which pertains to 1995. It is clear that we are in the throes of a reverse swing. Asia, which lost around six percentage points over the period 1650–1933, has returned to almost exactly the 1650 share. Africa has come back as well, but is still significantly below the 1650 share. The two gainers have been North America and Latin America. It is also instructive to add up what approximately accounts for the developing world. The population share of Asia, Africa, and Latin America combined was 81.1 in 1650. In 1933 it had dropped to 67.6. The share was 81.7 in 1995. We have come full circle.

Without this historical perspective it is easy enough to be alarmist about population expansion in developing countries. No one doubts that such expansions may be harmful, but it is certainly not the case that these countries have grown more than their "fair share." What alarms many governments in the developed world is not population growth, but *relative* population

[1] This description is a bit simplistic. The populations of Japan and China were also in a state of significant increase over the last half of the seventeenth century. China's expansion continued through the eighteenth century. The demographic rise of Europe is even more impressive against this moving background.

growth. A large population means greater poverty and smaller per capita access to resources, but on the international scene, it stands for greater political and economic power. The very same governments that stand for population control in the developing world are perfectly capable of pursuing pronatalist policies at home.[2]

Attitudes to Population[3]

Most individuals and governments, if polled, would agree that world population trends pose a problem. When it comes to judging trends in one's own country, however, matters are often quite different. We may deplore an action as being harmful to the interests of society, yet be tied into taking that very same action, simply because others are. Recent changes in attitudes to population, however, show a welcome transition.

At the International Conference on Population and Development, held in Cairo in September 1994, many governments clarified their stand on the population question. It was clear that many governments were actively pursuing demographic policies to limit population trends, and indeed, population growth in many developing countries has significantly declined. These outcomes are correlated with some changes in government perceptions of population growth. Although the percentage of countries that consider their rates of population growth to be too *low* has steadily declined, the number of governments that view population growth as too *high* has declined somewhat as well. Developing countries take the lead in this change of attitude. Among the developed countries, there has been little change. Indeed, an increasing number of such countries consider their rate of population growth to be too low and are concerned about declining fertility and population aging.

In Africa, we see an increasing number of countries joining the war against population: Namibia, the Sudan, and Tanzania officially inaugurated policies to reduce population growth. The Tunisian government now declares itself satisfied with the declining trend of its rate of population growth. Likewise, in Asia, more governments have declared themselves satisfied with demographic trends, although many still consider their population growth rates to be too high. China and Korea both view their current situations as satisfactory.

In contrast, in Europe, more countries are concerned with aging and population decline. Portugal and Romania now consider their population growth rates to be too low, and Croatia inaugurated a policy to promote fertility rates.

In Latin America, as in Asia, an increasing number of countries consider their population growth rates to be satisfactory. The exceptions lie in the densely populated areas of the Caribbean, and in Central America.

Little change occurred elsewhere. In North America, the United States and Canada remain satisfied with their population growth rates, as do Australia and

[2] On these and related matters, see Teitelbaum and Winter [1985].

[3] The account here relies heavily on a report of the Secretary General of the United Nations, presented to the 28th session of the Population Commission, 1995.

New Zealand in Oceania. The majority of developing countries in Oceania consider their rates of population growth to be too high (Tonga is an exception because of high rates of emigration). In Eastern Europe, four countries (Belarus, Bulgaria, Hungary, and Ukraine) consider their population growth rates to be too low. In the former Soviet Union, a majority of the governments appear to be satisfied with their current demographic regime.

At the same time, we cannot help being concerned about future trends. Look again at Table 9.3. It took Europe and North America a good 300 years to realize their population gains; it took around 50 years to lose them. If we extrapolate these trends, are we not in danger of an enormous population explosion, with a rising majority in the developing countries?

In examining this important question, we take note of a radical difference between the demographic transitions of the developed and developing worlds. The latter is being played out at a pace that is many times faster than that of the former. The second phase of demographic history in developing countries displays an intensity that is unmatched by the experiences of the now-developed world.

In developed countries, the fall in the death rate was relatively gradual, limited by the trial and error of innovation. The improved production of food, the institution of sanitation methods, and the greater understanding and control over disease yielded by medical advances all had to be discovered or invented, rather than transplanted from a pre-existing stock of knowledge.

For several reasons, including norms of late marriage in many European societies, birth rates never attained the same heights that we see in developing countries today. At the same time, birth rates fell slowly, in part due to a greater carrying capacity made possible by technical progress. Thus the second phase of demographic history was protracted, and the time span (centuries) more than compensated for the (relatively) low net growth rate. Population growth in these countries was more of a slow burn than a violent explosion, and its enormous impact was felt over centuries.

Contrast this picture with what has happened to developing countries. The decline of mortality was widespread and sudden. Antibiotics were available for a variety of illnesses; they did not have to be reinvented. The use of insecticides such as DDT provided a cheap way to bring down malaria to manageable proportions. Public health organizations began to pop up all over the developing landscape, some of them funded by international institutions such as the World Health Organization. Last, but not the least, there was widespread application of elementary methods of sanitation and hygiene. These are all blessings, because they brought to people a longer, healthier life.

The fact remains, however, that the easy and universal application of these new techniques led to a precipitous decline in death rates. The speed of decline surpassed anything experienced by Northern and Western Europeans. Everything, then, hangs on the birth rate. How quickly does it follow the death rate on its downward course? On this question hangs the future of the world's population, and certainly the economic future of many developing countries.

9.3.3. The adjustment of birth rates

Macro- and micro-inertia

The preceding story of the demographic transition relies on an enormously important feature—the well-documented failure of the birth rate to instantly chase the death rate downward. Recall from the previous section that the main impetus to the rise in population growth rates comes from the fact that death rates decline rapidly, while at the same time, birth rates hold firm. There are several reasons for this.

It is certainly true that over the past centuries, the factors that precipitated the fall in death rates were also linked with an increase in the carrying capacity of the earth. The leading example of this was a rise in agricultural productivity. This is one explanation for why birth rates did not fall (in those times). Unfortunately, this argument cannot be applied today. Many of the factors bringing down death rates in developing countries today are sanitation and health related: they do not go toward increasing carrying capacity.

We study in this section the various factors that keep the birth rate high. At the outset, it will be useful to distinguish between two forms of inertia in the birth rates: one at the level of the overall population (macro-inertia) and one at the level of the family (micro-inertia).

As discussed already, the distribution of the population by age plays an important role. The fact that both birth and death rates are initially high in a poor country makes the net population growth rate low, just as in rich countries, but there is a second implication that is quite different: the populations of the former type of countries will be very young, on average. This feature tends to keep overall birth rates high even if fertility rates are reduced at different age groups. The sheer inertia of the age distribution guarantees that young people of reproductive age continue to enter the population. One might think of this as macro-inertia—inertia at the aggregate level.

Macro-inertia is not the only form of inertia keeping birth rates high. There is also what we might call micro-inertia—inertia at the household level, perhaps in conjunction with the operation of societal norms regarding children and other socioeconomic factors. This will be our focus of attention for the rest of this section.

Fertility choice and missing markets

The angle that we explore in this section is that offspring are generally a substitute for various missing institutions and markets, notably the institution of social security in old age. This absence often compels a couple to make fertility choices based on the recognition that some of their children will die. These potential deaths must be compensated for by a larger number of births.

Of course, children bring enjoyment to their parents, as they undoubtedly do in all societies, but this is not the only reason why they are produced. On top of this "consumption-good" aspect of children is their role as an "investment good"; that is, as a source of support to the family in old age, and more generally as a form of insurance. If it were possible to obtain insurance or old-age security from a more efficient source, these effects would go away. As we have seen already and will see on several occasions again, the fact that there is a missing market somewhere spills over into other seemingly disparate aspects of economic life.

To begin, then, let us get a sense of what markets are missing in this context. If you live and work in a developed country, you pay a good fraction of your earned income into a government fund that often goes under the name of a social security fund. When you retire, this fund pays you a retirement pension. It is necessary to contribute to this fund to receive benefits from it, although in many countries the pension is progressive (larger contributors do not get back all their payments). A second source of old-age funds is an employer-subsidized retirement plan (where both you and your employer makes contributions). Finally, you can save for your own retirement, not necessarily under the umbrella of any retirement plan.

Next, there are various forms of insurance that are available to you, both in your working life and in your old age. Perhaps the most important of these is medical care, but there are also other forms of insurance. Life insurance is among the most important of these. If you die, your spouse receives a payout from the insurance company that helps to support him or her. There is also insurance that you can buy to protect you from sudden loss of employment, or from disability, or from natural disasters, or from theft. This is not the case that in developing countries: these markets are completely missing. By and large, these institutions are only available to people who work in the formal sector. In the informal sector, where employment is largely casual and wages are abysmally low, there is little or no incentive to set up a retirement scheme between employer and employee, and even if the law says that this should happen, it is impossible to implement. Likewise, appropriate contributions to a government-run social security system are difficult, if not impossible, to assess. Large sections of the population live in rural areas or work in informal urban areas. For the same reasons of limited information, it is very difficult for an insurance company to assess the validity of claims, such as a crop failure or a sudden drop in the income of a streetside hawker. Agriculture is particularly hampered by the fact that income shocks may be highly correlated across policy holders, which necessitates large payouts for insurance companies. Of course, these correlations can be avoided by companies that operate at a national level, but such companies may lack the local expertise to collect adequate information. Therefore, insurance markets in the agricultural and urban formal sectors are often missing.

What about life insurance or personal savings? Both these avenues are somewhat more viable. With reasonable banking systems, individuals can save for their own retirement. It may be impossible to verify a crop failure, but it is certainly easier to verify death. Thus these routes to old-age support are often available.

At the same time, people often do not avail themselves of these routes. The reason why this is so has to do with low incomes. Consumption needs today are often so pressing that there is little left over to save. People therefore often hold on to assets that they might have inherited, such as land or jewelry, and sell these assets only under conditions of extreme duress. These assets form their security in old age.

Note that the more difficult it is to sell an asset for current consumption, the easier it is to save using that asset. You might respond that if it is difficult to sell that asset now, why should it be easier to sell when funds are truly needed? The answer has to do with the nature of the difficulty. Society sprouts norms around the sale of assets such as land and jewelry. It is all right to sell them under severe duress, but the sale of these assets in "normal times" might be frowned upon or regarded as a signal that the family is completely indigent. Thus the difficulty of marketing these assets is created by the emergence of social norms that protect savings in some form.

In this context, take a fresh look at children. Children are assets par excellence. They do not need to be bought, although there are costs to child rearing (see subsequent text) and they embody income-earning possibilities, both now and in the future. Because slavery is banned (and socially unacceptable anyway), it is generally not possible (though unfortunately, not impossible) to market them for cash. At the same time, when children grow up they can convert their labor power into income, both for themselves and their parents. Little wonder, then, that individuals who lack insurance and old-age security, choose to invest in the future in the form of children. This is the background against which we investigate theories of fertility choice.

Mortality and fertility

Consider the probability that a child will grow up to look after its parents. This probability is given by several factors. The child may die young; infancy is the biggest hurdle. As we know from Chapter 2, infant mortality may be close to 150 or 200 per 1,000 in several developing countries, which translates into a 15% probability of death by the end of the first year of existence. After this barrier, there are still the diseases of childhood, which are still a significant killer in many developing countries up to the age of five or so.

Third, there is the possibility that the child may not be an adequate income earner. The poorer the economy, the greater this fear.

Fourth, a child may not look after its parents in their old age. This is an interesting social factor that may cut in either direction. In societies where the norm of looking after one's parents has practically vanished or is relatively nonexistent to start with, the mental calculations that we are going to talk about may have no relevance at all for fertility decisions. For instance, economic historians such as Williamson [1985] have argued that fertility reductions in nineteenth century United Kingdom can be explained by the increasing emigration rates of adult children. If emigrees send limited remittances, this reduces the present value of children (as investment goods) considerably.

At the same time, in societies where it is accepted practice to care for one's parents, the *limited* possibility that some child might not do so may have the opposite effect on fertility: instead of lowering it, it may increase it as parents attempt to compensate for this contingency.

Finally, there is the possibility that the *parents* themselves might not anticipate being around in their old age. This is certainly a possibility in very high-mortality societies, but in general it is of second-order importance. At the stage in their lives when individuals are making their fertility decisions, they have already lived through the bulk of the (non-old-age) high-mortality phase.

Summarize the overall probability of having a given child grow up to look after you by p. This includes, then, infant and child mortality, the eventuality that the child survives but is not an adequate income earner, and the possibility that the child earns adequate income but nevertheless does not look after you. What value might p take? It is hard to tell without detailed data on each of these possibilities, but child mortality by itself might be responsible for raising p to well above 1/5. With the other factors accounted for, p may well be higher than 1/3, and the possibility that parents regard p as a one-half (or close to it) is certainly not unreasonable.

Now contrast this with the probability—call it q—that a couple finds *acceptable* as a threshold probability of receiving support from at least one child. This is a matter of attitudes toward risk and varies greatly from couple to couple. Try introspection: what probability would you find acceptable to be without any form of old-age support? If you could honestly tolerate a probability that is significantly greater than 1/10, you are an unusual person. We may therefore think of q as having values above 9/10—perhaps even as high as 95/100—certainly greater than p. The rest is a matter of simple arithmetic: how many children do you need to have—each child looking after you with probability p—so that the overall probability of having at least *one* child look after you is at least q?

This is easy to calculate (or it should be!). Suppose you have n children. Then the probability that *none* of them will look after you is $(1-p)^n$. Consequently, your rule would be to choose n—the number of your offspring—just large enough so that

$$1 - (1 - p)^n > q. \tag{9.1}$$

Let us check this out with some numbers. Say that $p = 1/2$ and $q = 9/10$. Then it is easy enough to see, using (9.1), that n must be at least 4! If you are more risk-averse than that, so that your acceptable q is 95/100, then you will need five children, and that, too, brings you barely to your acceptable threshold, as you can check by direct calculation.

Gender bias

In this context, gender bias can be immensely costly. Suppose that for some reason, a couple wishes to receive support from a son. Households will then see n as their desired number of *male* offspring. Quite simply and devastatingly, it doubles the expected number of children that the household will have.

For instance, if $p = 1/2$, if $q = 9/10$, and if the couple desire support from a male child in their offspring, then that couple will have, on average, eight children, all for the sake of ensuring *just* one son!

In many societies, the provision of old-age support is thought to be exclusively the task of male offspring. Although support (especially in nonmonetary form) from female children is just as valuable, there may be a stigma associated with receiving support from daughters as opposed to sons. This bias is, of course, a source of discrimination in favor of male children.

To be sure, this argument does not explain the rationale behind such a bias, and there may be many reasons. For instance, Cain's [1981, 1983] study of Bangladesh illustrates the importance of sons as support for widows: *the ability of widows to hold on to land depends on whether they have able-bodied sons.* This will be especially true in situations where property rights are either not well-defined or difficult to enforce by the law.

Information, income, and fertility

Let us summarize the discussion so far. Individuals choose the number of their offspring with the intention of receiving support in their old age. This support may not be forthcoming from a child for several reasons: (1) the child may die, as an infant or later in life, (2) the child may not earn enough income to support the parents, and (3) the child may break parental ties and deliberately not support its parents, even though it has the economic capability to do so. The probability of these uncertain events taken together we denoted by p.

The uncertainty described in the preceding paragraph has to be compared with the tolerance threshold of the parents, which is the minimum probability that they need old-age support, and this threshold changes with the degree of risk aversion of the parents. The degree of risk aversion, in turn, depends in part on the economic security of the parents. A higher level of security generally implies a lower degree of aversion to risk.

These factors help us to uncover, to some extent, the reasons behind a sticky fertility rate in the face of rapidly falling death rates. The first element is *information*. How is the social phenomenon of a falling death rate translated to the level of individual decision making? We have already commented on the rapidity of the decline in death rates in developing countries. For twenty-one developing countries during the period from just before World War II until 1950, the death rate dropped on average by seven per thousand population every ten years (Coale and Hoover [1958, p. 14]). For a historical demographer and indeed for any social scientist, this is a remarkable change indeed and is unparalleled in history for its rapidity. As Coale and Hoover [1958, p. 14] observed, "this rate of improvement surpasses anything from the records of areas inhabited by northern and western Europeans." Yet it would be wondrous indeed if these changes made the newspapers at the time! The fact of the matter is that individuals must often go by their own experience, by which I mean their vision of the experiences of their *parents* and the siblings and friends of their parents. It is the preceding generation that provides the only direct experience that is relevant in this context.

Thus the fall of death rates may not instantly translate into a revised estimate of mortality (see box, Three Generations).

Three Generations[4]

The village of Rampur in India was surveyed by Lewis [1952] and then resurveyed by Das Gupta [1994]. The story of Umed Singh comes from them. Umed Singh's father was Siri Chand, who was born around 1900. Epidemics of plague and cholera decimated his family, including his father and mother. Siri Chand was brought up by his uncle. As a farmer, he faced the kind of uncertainty that is difficult for us to even contemplate: consecutive crop failures, famine, the occasional bumper

[4] This box is based on Das Gupta [1994].

crop, the loss of *six* out of nine live births: two girls and one boy survived to adulthood. The life of the boy, Umed Singh (who was born around 1935), stands in sharp contrast to that of his father.

Umed Singh completed secondary school and became a policeman. He earned a regular salary and also received income from his land (left to him by his father). However, the uncertainties of his father's life never ceased to haunt him. He was the sole surviving son in a family that had given birth to nine children. With no objective reasoning to back his insecurity, Umed Singh worried and then worried some more. His first two children were girls. Because he wanted a boy, he insisted on having more children. He had three more children, and two of them were boys, but he continued to worry that his children would die, and this fear did not leave him until his third son was born. *All* his children survived.

As Umed Singh relived the anxieties of his father, people around him were already changing. His wife, when interviewed, felt that they should have stopped having children much earlier. So did Umed Singh's cousins and his colleagues in the police force.

Das Gupta ends the story thus: "The second generation of people who lead a secure, ordered life do not experience the anxieties left over from past insecurities. Umed Singh's oldest daughter has completed a course in teacher training and will be married shortly. She says she has no intention of childbearing in the way her mother had; three children were the maximum she would have. She is a relaxed, confident woman, who is inclined to be a little amused by her father's anxieties on behalf of his family."

Although falling death rates are central to the fertility decline, there are other factors in the construction of p that have little to do with the fall of death rates. These are the previously mentioned items (2) and (3), which may well go the other way even as death rates fall. These depend on the economic conditions of the region. The poorer the region, the greater the anticipated probability that a single child will not earn enough in adulthood to support parents; hence, the greater the incentive to have more children to compensate for this possibility. Likewise, falling death rates cannot in any way affect the social possibilities of fulfilling parental obligations. These are independent phenomena that continue to leave their mark even as death rates fall, and they might keep birth rates high.

Finally, there are the additional complications introduced by gender bias. Again, there is no guarantee that a fall in the death rates will have any impact on the degree of bias. In making this statement, we actually distinguish between two types of bias. One is what might be called observable bias; that is, measurable indicators of differential treatment of boys and girls. With development, such bias indeed lessens as resource constraints loosen. A second sort of bias has to do with the intrinsic valuation of women in society and it feeds into the perception of women as sources of old-age support. This bias actually increases with economic progress, at least to a certain extent. One important reason for the potential increase is that economic progress is associated with a decline in the importance of agriculture. To the extent that women play an important role in agriculture, they may now be perceived as relatively less capable of providing old-age support on their own. We have already seen that such biases, apart from their obvious intrinsic shamefulness, can brutally affect fertility decisions.

Hoarding versus targeting

Our discussion so far contains an implicit assumption: that parents must make their fertility decisions about later children without being able to use information about the fate of their earlier children. Is this reasonable? Again, the answer depends on just which components of p are dominant in parental psychology. For instance, if an individual worries that the child may not earn enough in adulthood to support his aging parent, this is not an outcome that lends itself to a wait-and-see strategy. By that time, it will not be possible to have a new child! If the source of the uncertainty resides in such features, all the tickets will have to be bought in advance, as it were. We may refer to this phenomenon as one of *hoarding*: children have to be stockpiled in advance, before we know which (if any) among them will provide the requisite support.

Contrast this with a situation where infant mortality (death before the age of one) is the dominant form of uncertainty. In such a situation a wait-and-see strategy acquires greater feasibility. A couple can have a child and condition its next fertility decision on the survival of this child. The desired number of children can be attained *sequentially*; this strategy is called *targeting*. Obviously targeting generally is associated with lower fertility rates, because the total number of children do not have to be created "in advance."

A change in the demographic regime from hoarding to targeting can lead to a drastic lowering of the fertility rate. Again, the rate at which this switch of regime occurs depends critically on the kinds of uncertainties that the couple is most concerned about. It is true, however, that a fall in the death rate can only assist in bringing about this change of regime.

The costs of children

So far we have neglected the costs of child rearing. These costs take two forms. First, there are what might be called the direct costs of children: they need to be fed, clothed, kept in good health, looked after, and schooled. Second, there are the indirect or *opportunity* costs of children that are measured by the amount of income foregone in the process of bringing up the child. Time spent at home with the child is time not spent earning income, so the opportunity cost of children is roughly proportional to the going wage rate multiplied by the number of hours spent in parenting.

In societies where this opportunity cost is low, fertility rates tend to be high. Gender bias plays a role in this as well. In many societies (including many developed countries), it is presumed that women must allocate the bulk of their time to the upbringing of children. In such societies wages for women's work are low as well. This brings down the opportunity cost of having children and keeps birth rates high.

Similarly, if there are high rates of unemployment, the opportunity costs of children comes down. Again, this can push fertility upward.

This cost–benefit approach to fertility choice is natural to economists. Becker [1960] introduced this approach to other social scientists. Often, the methodology is not very useful: simply stating that parents have children up to the point where marginal benefit equals marginal cost may be an impressive piece of jargon, but does not convey much information. To make the cost–benefit approach useful, we must either discuss benefits, or costs, or both in a way that is relevant to the situation at hand. This is what we have done so far with the notion of benefits. Instead of stating that parents derive "utility" out of children, we describe it specifically as old-age support, and this description allows us to draw the specific conclusions that we have arrived at so far. So it is with costs. We need to understand how different *kinds* of costs have different sorts of demographic implications. In the discussion that follows, we illustrate this point by considering a specific case: the effect of income improvements on fertility.

Figure 9.1 considers the preferences of a couple over the number of children it wishes to have and "other goods," denominated in terms of money. Children are on the horizontal axis; other goods are on the vertical axis. In

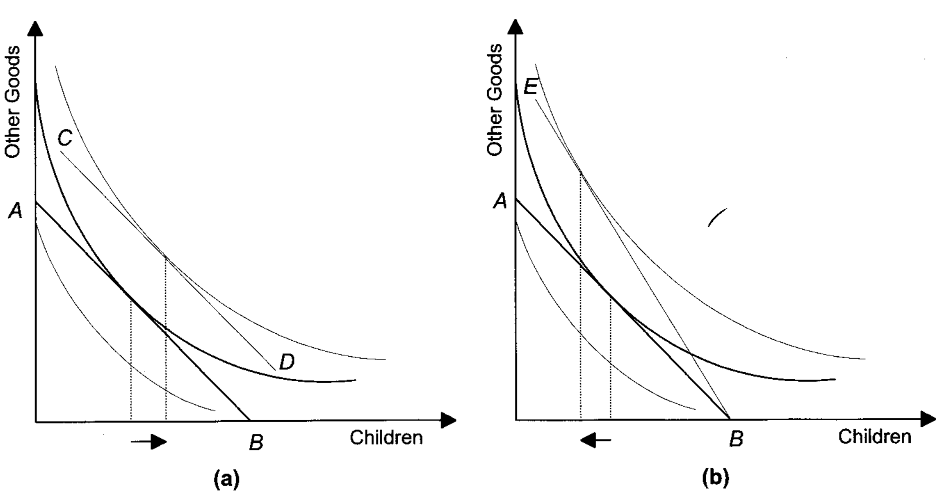

Figure 9.1. Income improvements and fertility.

what follows, we do not pay much attention to the exact form of preferences, which are represented by indifference curves in Figure 9.1. For instance, these preferences may be a reduced form of the desire for old-age support or may simply arise from the intrinsic pleasure of having children. Our focus is on the costs.

Consider, first, the total *potential* income of the couple if they were to have no children at all. Income may be wage labor or some other form of income, such as land rent. If it is the former, potential income includes all income earned by each spouse, under the scenario that they have no children to look after. This amount is represented by the height of the point *A* in Figure 9.1.

Now it should be clear that as the number of children begins to rise, the income left for "other goods" will begin to fall. It will fall for two reasons. First, there are the direct expenses of the children. Next, earned income falls as well, because one or both parents spend some time at home to look after the child. This trade-off traces out the "budget line" *AB*. With very large numbers of children, residual income available to the parents may drop to zero: this is the point *B* at which the budget line cuts the horizontal axis. Panel (a) of the figure incorporates this description.

The *slope* of the budget line (see the indicated angle in Figure 9.1) is a measure of the unit cost of having children. If income is earned, it will be the wage rate multiplied by hours foregone per child. In addition, there are the direct costs per child.

The exercise that we now conduct has to do with an increase in family income. To fix ideas, suppose first that the source of income increase is not wage income. For instance, the individual in question may be a landowner who receives all income from leasing land to tenants. Rents have gone up. In this case, the budget line will undergo a parallel shift, moving to the position *CD* [see panel (a)]. What effect does this have on fertility? Well, if children are "normal goods," the income effect must raise the demand for children, so that fertility rates go up as a result of the income increase.

Contrast this change with a change in *wage* income. In this case, the budget line will not only shift outward, it will *swivel* as well. This is because the opportunity cost of children has gone up. In Figure 9.1(b), we show this by shifting the budget line out and rotating around the point *B* at the same time, so that we have the new budget line *EB*. Potential income has gone up, but *at the same time the opportunity cost of children has gone up as well.* This creates a substitution effect away from children as well as an income effect. *The substitution effect lowers fertility; the income effect raises it.* The net effect is ambiguous.

Despite the ambiguity, one thing is clear from Figure 9.1: fertility certainly does not rise by as much as in the case where the income increase can be traced to "nonearned" sources. The intuition is straightforward. Wage income imposes an opportunity cost of having an extra child, whereas rental income does not. Thus wage income increases have a stronger impact on fertility reductions than rental income. This illustrates the usefulness of the cost–benefit approach, at least up to a certain point.

We can easily extend this argument to the case of gender bias. Suppose that only women look after children. Then an increase in rental income has the same effect as an increase in male wages: both lead to a parallel shift of the budget line, as in the move from *AB* to *CD*. Male wages impose no opportunity cost on childbearing if men play no part in raising children. In contrast, the swiveling of the budget line is characteristic of an increase in *female* wages. The opportunity cost of having children will go up. It follows that a society with gender bias is more likely to exhibit a reduction in fertility when female wages go up, as opposed to the case in which male wages rise. This argument was examined in a paper by Galor and Weil [1996] and by many others (see also the following box).

Women's Wages and Fertility Decline: A Study of Sweden

Over the last century or more, there has been an increase in the wages of women relative to men. This is certainly the case in currently developed countries. Along with this increase, we see a concurrent reduction of fertility. Is this clear evidence of a *causal* relationship between women's wages and fertility? It may not be. It is conceivable, for instance, that a reduction in fertility occurred for some other reason, and this reduction was associated with larger investments by women in education, which raised their wages. In this hypothetical situation, fertility and female wages are *correlated*, but no evidence of *causality* is established. What we need is *separate* evidence, quite apart from actions that may have been taken by women themselves "on the supply side," that female employment is more in demand. Then we can relate this piece of evidence to the fertility decline.

Schultz [1985] raised precisely this question and addressed it in an interesting way using Sweden as an example. In the second half of the nineteenth century, the world grain market went through a declining phase of major proportions. The exports of Swedish grain collapsed. Faced with this decline in grain demand, there was a substantial reallocation of resources in agriculture. Animal husbandry was the benefactor. Swedish exports of butter soared.

Now, dairying and the processing of milk employed a larger fraction of women than did grain farming. As a result of this reallocation, the demand for female labor went up significantly and so did the wages paid to women.

The usefulness of focusing on the butter boom is that it effectively captures a pure demand effect on female wages, rather than an effect that could have been created by supply decisions. Did fertility drop in response to the butter boom?

It did. Schultz shows that in regions where the price of butter relative to rye (the basic food grain in Sweden) is high, the ratio of female to male wages was high as well and fertility rates were lower. Indeed, following up on the link between butter prices and female wages, Schultz estimated that about a quarter of the decline in the Swedish total fertility rate from 1860 to 1910 can be explained by the rise in relative female wages. The conclusion is that "the appreciating value of women's time relative to men's played an important role in the Swedish fertility transition."

9.3.4. *Is fertility too high?*

So far we have tried to provide an account of why fertility rates may be high in the face of falling death rates, but "high" does not necessarily mean "suboptimal": if a family *chooses* to have a large number of children, then why should social considerations dictate anything different? There are three answers to this question.[5]

Information and uncertainty

The first answer relies on the incompleteness of *information*, which is an issue that we have already discussed. People simply may not internalize the general observation that death rates have undergone a decline, as in the example of the Rampur policeman Umed Singh (see box). In such a case, the number of children that couples have may not be socially optimal. Faced with fresh information regarding the environment that influences their fertility choices, the couple would typically revise their fertility decisions.

The second answer relies on the distinction between decisions that are made *ex ante* and their *ex post* consequences. Consider the family that wants one child, but has five, in the hope of increasing the chances of old-age support. As we have already seen, such decisions are based both on the probability of a child dying and on the degree of aversion to risk of the family. Thus it is not unlikely (and this will be true especially for poor families that are highly risk-averse) that, in fact, a large percentage of the children do survive *ex post*. Such families will have too many children and they will suffer because these children will have to be looked after and fed. The evaluation of optimality becomes problematic in this case. If a family with a large number of children is asked if they are happy with this situation, they may reply that they are not, but if asked whether they would have made exactly those fertility choices all over again (in the face of the uncertainty that shrouded

[5] There is a fourth as well, which is that families (especially illiterate and poor families) do not know what is best for them. In particular, they procreate without thinking or being aware of effective contraceptive methods. According to this view, an expansion in the supply of contraceptive devices and a good lecture will take care of the problem. I do not discuss this viewpoint here, but see the subsection on social norms at the end of Section 9.3.

survival), they may well say that they would have. There is no contradiction between these seemingly contradictory answers.[6]

Externalities

The third and most important answer is based on the existence of *externalities*. That is, the fertility decisions made by an individual or a couple may have implications for other members of a family or indeed for other families. To the extent that such effects are not internalized by the decision maker(s), fertility choices that are privately optimal may not be socially optimal.

As the following cases suggest, fertility-related externalities are typically *negative* (though this need not always be the case). Thus private fertility choices generally lead to overly large numbers of children.

Let us begin by studying some effects *across* families. These externalities are particularly pervasive in situations where infrastructure is provided by the government at little or no cost to users. In such cases, it is not possible for individual families to value these resources at their true social cost, because that is not the cost they (or their children) pay. This is not to say that such services should always be provided at market prices (often they represent the only way to redistribute income in an unequal society), but they do enlarge the number of situations in which an externality may be present.

Consider, for instance, the provision of free public education in an urban area. If a benevolent social planner could dictate the number of children that all families should have in that area, she would take the marginal social cost of providing educational resources into account. However, if education is provided free of charge, the *private* cost to the family typically is lower than the social cost, which therefore will not be properly internalized. It follows that the number of children that people choose to have will exceed the social optimum.

The same is true of other publicly provided services that are not valued at their true marginal cost, such as subsidized housing or health services. As I have already mentioned, these may often be the only feasible way to target the poor in a society where direct information on economic characteristics is hard to get hold of. These services have the same effect as the provision of education: they reduce private marginal costs below the social marginal costs and push fertility beyond the social optimum.

A similar set of observations applies to resources that are not properly priced, such as the environment. Such resources can be depleted even if

[6] Even the ex ante choices may be suboptimal, because there are missing markets. Specifically, in the absence of a missing market for insurance, in general, and old-age security, in particular, families tend to overinvest in children. If these options were provided, the number of children per poor family would surely decline. The point is, however, that the choices *are* ex ante optimal given that the markets are missing.

they are renewable: they include fisheries, groundwater, forests, soil quality, and of course the ozone layer. The main characteristic of such resources is that they are generally underpriced, so that financial incentives bias their use in the direction of overexploitation. To the extent that such underpricing reduces the cost of child rearing, fertility is biased upward.

All of these effects can be summarized in one very general framework. In Figure 9.2, we show the costs and benefits of having children (say, for a single family). For simplicity, we take the cost curve to be a straight line (so that each new child costs the same additional amount), even though there are diminishing returns to having more children. This means that the benefit function has the familiar concave shape.

Focus now on the costs. The thick straight line shows the *private* costs of an additional child and the thinner line shows the *social* costs of an additional child. The preceding discussion indicated that in many situations, the private costs may be less than the social costs. Diagrammatically, this is captured by the fact that the "social cost" line is steeper than the "private cost" line. The socially optimum number of children is found by maximizing the vertical distance between the benefit line and the social cost line. This point is found by setting marginal benefit equal to marginal social cost, which occurs at the point A and yields a number of children n^*. In contrast, the privately optimal number of children is found by maximizing the vertical distance

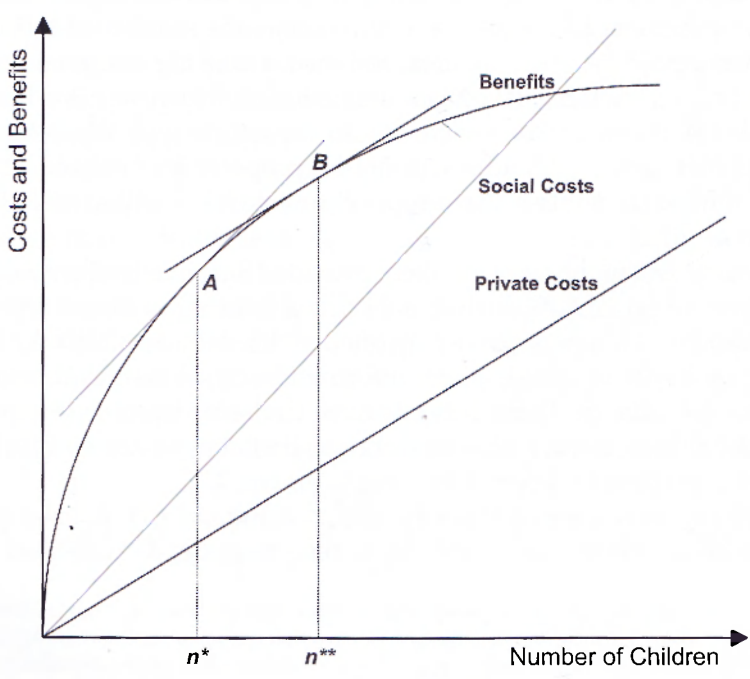

Figure 9.2. Private and social costs and fertility decisions.

between the benefit line and the private cost line. This occurs at the point B with associated number n^{**}. Note that $n^{**} > n^{*}$.

This sort of analysis summarizes all the various situations listed previously and tells us how to think about others. Here are two examples that serve as extensions of the analysis.

First, all situations may not involve a divergence between social and private costs. There may be a divergence between social and private benefits as well. Suppose that prized jobs are available for a high pay, say, $1,000 per month, but that there is a queue for such jobs. Imagine that each family sends its grown-up children to look for such jobs and that, for each child, the probability of getting the job is simply the total number of such jobs available divided by the total number of job seekers. Now having an additional child is just like buying an additional lottery ticket—like having a second shot at the prize. To the family, the probability of getting at least one job offer doubles. However, we must be careful here: the number of job seekers goes up too. This effect is minuscule at the level of the family in question, but the combined effect of *many* families buying their two lottery tickets each on *other* families is significant and negative. In the end, each family has, say, two tickets each and nobody's chances of getting the job are really increased. Worse still, that second ticket is costly: it is a child who has to be clothed and fed.

This kind of situation is easy enough to analyze in the general framework that we have set out. You can easily check that in this example there is no divergence between private and social costs, but there is a divergence between private and social *benefits*. The social benefit of an additional child is the private expected gain *plus* the losses inflicted on all other families by swelling the ranks of the job seekers by one. This is an externality.

Our second example is designed to show that externalities can occur *within* the family as well. This is especially true if there are family members in the household other than the couple making fertility decisions. Consider, for instance, a joint family: typically one in which two or more brothers pool resources to live under a common roof. I do not know if you have ever experienced the wonders of a joint family; I have friends who have. At first glance it is impossible to tell parent from aunt or uncle, because aunt and uncle participate significantly in the upbringing of children. The effect is two-way, of course: my cousins will likewise be looked after by my parents. Now this looks like a happy state of affairs (and often it is and often it isn't), but the point I wish to focus on is the observation that joint families naturally create an *intra*family externality. Knowing that one's brother and sister-in-law will bear part of the costs of child rearing lowers the private costs of having children and raises fertility!

Now something looks suspicious in this argument. There must be a "law of conservation of costs." Everybody's costs cannot be simultaneously lowered. For instance, the brother and sister-in-law are surely passing on some of the costs of child rearing to *my* parents, so why does it all not cancel out, leaving fertility decisions unaltered relative to those which would have been made in a nuclear family? The answer is simple. It is true that my parents are bearing part of the costs of rearing their nephews and nieces, but this is a cost that they cannot control, because the fertility decisions regarding nieces and nephews are being made by *someone else*. Thus these costs are fixed costs as far as my parents are concerned, whereas the costs of their own children that they in turn pass on are *variable*, because they make the decisions regarding their own offspring, and only the *variable* costs count in the fertility decision. This is what Figure 9.2 implicitly teaches us. The *slopes* of the private and social costs, and not their levels, are the key determinants of fertility. This is not easily seen in that figure, so Figure 9.3 provides an appropriate variant. The thin line in Figure 9.3 represents the cost of one couple's children to the entire (joint) family. Because part of this cost is passed on to the hapless brother and sister-in-law, the variable cost to the *couple* is given by the flatter thick line passing through the origin of the diagram. Now, as we said, the same kind of cost transfer is faced by the couple in question, which raises their total costs, but only shifts their cost line in a parallel way (see Figure 9.3). This shift of levels does nothing to affect their fertility choice, which is n^{**}, above the level that is optimal for the joint family as a whole (or for the couple had they been nuclear), which is n^*.

The same kind of argument holds if there are grandparents to look after children. If the grandparents' costs are not fully internalized by the couple,

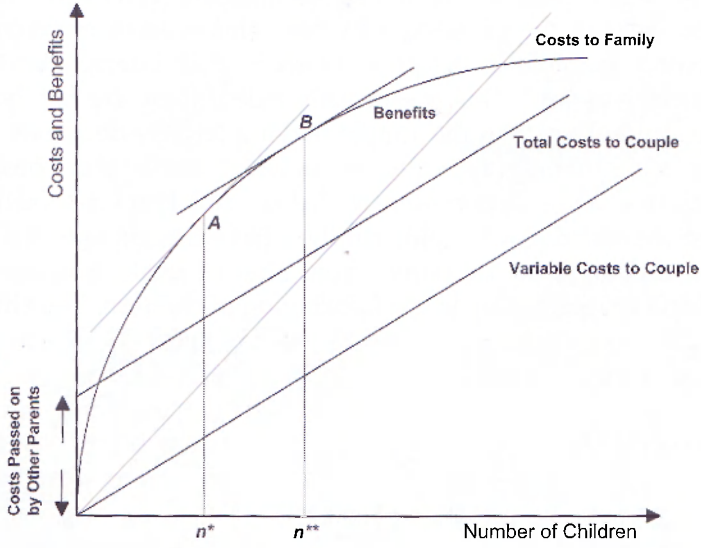

Figure 9.3. Fertility choices in joint families.

they may have too many children relative to what is optimal for *their* family, leave alone society as a whole.

Thus family structure is very important in creating externalities that lead to excessive fertility. As such structure changes from joint or extended families to nuclear families, the costs of children are more directly borne by the couple, which leads to a decline in fertility.[7]

In all the preceding cases there are negative external effects of fertility decision, so that fertility choices are typically high relative to the social optimum. There are situations in which there might be positive externalities as well, especially if the optimum for that society includes the pursuit of pronatalist policies to gain economic or military power. Such concerns may also be felt in societies in which a long history of low population growth has shifted the age distribution uncomfortably in the direction of high age groups, which places immense burdens on social security systems. To the extent that an individual family does not internalize these goals, the state may actually reward child bearing in an attempt to provide appropriate incentives. As we have already seen, such pronatalist policies are the exception rather than the rule, although they do exist.[8]

Social norms

Often, people do what other people do. The glue of conformism is what holds social relationships and societies together. Conformism assures stability and limits the need for law enforcement, and indeed it is the expression of a shared conformism that we know as culture. We have already seen a discussion of social norms in Chapter 5.

The very strength of such norms becomes a weakness when the environment of the society begins to change. Accepted, appropriate practice over many centuries may now become inappropriate, but once this happens, social practice is often slow to alter. It becomes necessary to coordinate on some *new* norm, but such coordination requires many people to move in unison. In Chapter 5, we saw how difficult this is when there are multiple equilibria involving large numbers of people, such as in technology adoption. Here read "norm" for "technology."

Norms do ultimately change and chase what is appropriate, but it may be a long time coming. Consider a poor society with high rates of infant

[7] It is even possible to extend this sort of argument to externalities created at the level of the couple. To the extent that men and women are disproportionately engaged in child rearing, one of the parties (typically the man) may ignore these costs and not take adequate steps to reduce fertility or the risks of pregnancy.

[8] As an interesting twist on this argument, note that pronatalist policies may themselves be self-defeating at the world level, to the extent that they inflict negative externalities on other countries. The analysis of this second layer of externalities can proceed exactly in the same way as the negative externalities discussed in the text.

mortality and intensive use of child labor in farming, as well as for old-age support. It is not surprising to find such a society celebrating the birth of many children (especially sons). Such societies develop certain attitudes toward the "appropriate" age of marriage, the role of women, the importance of contraception, the desirability of primary education, ancestor worship, and even practices such as breast-feeding. Now imagine that advances in sanitation and medicine dramatically bring down infant mortality rates. Suppose that dependence on agriculture is on the wane (or mechanization is increasing, so that child labor is less important). Suppose that institutional forms of old-age security are becoming available. Will fertility change overnight?

We have already seen that it will not, but an additional reason for this is that people still want to conform to the old practices of having and celebrating children, to early age at marriage, and so on, *simply because everyone around them is doing the same.*[9]

These conformist tendencies may be bolstered by social and religious practices such as ancestor worship, that require the continuation of every lineage, often through males. Polygyny might also keep fertility rates high, as might the social importance of community over family (which brings down the private costs of child bearing in a way that we have already described). Even property rights might play a role. For instance, if land is held communally, it might be difficult to internalize the consequent costs of fertility in terms of the fragmentation of land holdings.

Jolting such a society into a "new equilibrium" is not easy. It requires *coordinated* change. An example of such a change is one in which ancestor worship is permitted through adopted children. If everybody thinks this is acceptable, then it's acceptable. It is in this sense that programs such as family-planning programs play a very important role. Quite apart from spreading important information regarding the cost, availability and effectiveness of different methods of contraception, these programs serve as a form of *social legitimization.* Consider the family-planning experiment known as the Matlab project in Bangladesh, in which seventy "treatment villages" were served by a birth control/family-planning program in 1977, while another seventy-nine "control villages" offered no such service. Contraceptive use in the treatment villages jumped from 7 to 33% in eighteen months. By 1980, the fertility rate in the treatment villages had declined to two-thirds that of the control villages.

What does the Matlab experiment teach us? One answer is that contraception was an unknown concept. People *wanted* to have two-thirds the

[9] The desire for this sort of conformity can have surprising consequences. At first glance, we might think that at the margin there will be *some* movement away from accepted practice, as people trade off their desire to conform with the desire to do what is best for them, but even such marginal movements may be blocked in conformist equilibria (on this point, see Bernheim [1994]).

number of children they were having, but could not do so. Perhaps, but at face value, this is unlikely. It is far more likely that the programs sent a strong signal that a lower *desired* fertility rate is actually a good thing: it is tolerated and indeed encouraged by society at large. People responded to this by adopting contraceptive devices to lower fertility. Thus it is possible that the program served two functions simultaneously: first, contraceptives were made widely available; second, and perhaps more important, the program signaled the advent of a new social norm in which lower fertility is actually a "good thing." Thus Phillips et al. [1988] wrote of the Matlab experiment, "An intensive service program can compensate for weak or ambivalent reproductive motives and create demand for services, leading to contraceptive adoption where it might otherwise not occur."

Social norms can be altered in other ways as well. The media is immensely powerful in this regard and can "transmit" norms from one community to another. The use of television and film to suggest that small families are successful can be of great value.

Social Norms and a New Fertility Decline[10]

According to the 1994 revision of the official United Nations world population estimates and projections, a fertility transition is underway in several sub-Saharan African and South-Central Asian countries. Fertility levels have traditionally been very high in these countries.

Total fertility rates have declined in Madagascar (from 6.6 in 1980–85 to 6.1 in 1994), Rwanda (from 8.1 to 6.5), United Republic of Tanzania (from 6.7 to 5.9), Namibia (from 5.8 to 5.3), South Africa (from 4.8 to 4.1), and Mauritania (from 6.1 to 5.4). Fertility declines are also evident in Zambia, Zimbabwe, and Gambia. If we add to this list Kenya and Botswana, where evidence of a fertility decline already exists, we see the beginnings of an overall fertility decline in sub-Saharan Africa.

South-Central Asia shows continued fertility decline: fertility has fallen in Iran (6.8 to 5.0) and continues its downward course in Bangladesh (6.2 to 4.4), India (from 4.5 to 3.7), and Nepal (from 6.3 to 5.4).

As discussed in the text, a widespread change in social norms may be playing a central role. Fertility declines everywhere appear to be accompanied by a significant increase in contraceptive use. We must be careful here to not infer any sort of causal link, but the increased recourse to contraception is indicative of an accompanying social transformation. Huge jumps in contraceptive use have been seen in Kenya (up from 7% of couples in 1977–78 to 33% in 1993), Rwanda (from 10% in 1983 to 21% in 1992), Bangladesh (from 19% in 1981 to 40% in 1991), and Iran (from 36% in 1977 to 65% in 1992).

[10] This account summarizes material made available by the United Nations Population Division, Department for Economic and Social Information and Policy Analysis, at http://www.undp.org/popin/.

Norms regarding the age of marriage must play a role as well. In Tanzania, for example, the incidence of contraception is low (10% in 1991–92), but the average age of a woman at marriage has gone up from 19 years in 1978 to 21 years in 1988. This is also the case in countries where contraception has significantly increased.

To be sure, fertility declines are not universal in this region and do remain high in the large countries of Nigeria, Zaire, Ethiopia, and Pakistan, but going by the broader picture, change is on the way.

9.4. From population growth to economic development

Just as economic development has implications for the pace of population growth, so the latter has implications for the rate of economic development. In large part, this relationship is thought to be negative. A large population means that there is less to go around per person, so that per capita income is depressed. However, this argument is somewhat more subtle than might appear at first glance. More people not only consume more, they *produce* more as well. The net effect must depend on whether the gain in production is outweighed by the increase in consumption. In the next two subsections, we clarify how the negative argument works and then follow this argument with some qualifications that suggest possible gains from population growth.

9.4.1. Some negative effects

The Malthusian view

Beginning with Thomas Malthus, a standard view on population growth is that its effects on per capita welfare are negative. Malthus was particularly gloomy on this score. According to him, whenever wages rise above subsistence, they are eaten away in an orgy of procreation: people marry earlier and have more children, which depresses the wage to its biological minimum. Thus in the long run, the endogeneity of population keeps per capita income at some stagnant subsistence level.

This is not a completely bizarre view of human progress. It probably fit the fourteenth to the eighteenth centuries pretty well. Blips in productivity, such as those in agriculture, increased the carrying capacity of the planet, but population did rise to fill the gap. It is difficult to evaluate this scenario from a normative standpoint. Over time, it was possible to sustain human life on a larger scale, even though on a per capita scale, the Malthusian view predicted unchanged minimum subsistence. Evaluation of this prediction depends on how we compare the prospect of not being born to the prospect of living at minimum subsistence. As I already stated, we sidestep this issue to some extent and concentrate on *per capita* welfare alone. By this yardstick, the Malthusian view is neutral in its long-run view of population growth.

A central ingredient of the Malthusian argument deserves critical scrutiny. Do human beings react to economic progress by spontaneously having more children? Modern experience suggests just the opposite. Individuals do understand that having children is costly, and it is perhaps true that the costs increase with economic development, while the (economic) benefits decline. For instance, we argued in previous sections that economic development is associated with greater provision of organized old-age social security. We have seen that such institutions probably are more effective than any other in bringing down rates of fertility in developing countries. People have children for a reason, not just because it is *feasible* to have them.

Likewise, economic progress may shift societies from an extended family system to a nuclear family system. As labor force participation increases, it becomes progressively more unlikely that individuals in an extended family all find jobs in the same locality. At the same time, the insurance motives that underlie the joint family setup probably decline. With nuclear families, the costs of child rearing are internalized to a greater degree, which brings down fertility.

There are other aspects that we have discussed as well, such as an increase in female wages or reductions in infant mortality with development. All these have a moderating impact on fertility. Thus it is absurd to entertain the notion that people react to any surplus in their incomes by automatically having more children. It is true that the Malthusian theory doesn't do a bad job for fourteenth century Europe, but in poor societies it is very difficult to separate the various determinants of fertility: fertility may have been high enough (for other reasons) relative to per capita income so that the Malthusian checks and balances applied better.

So as a first pass, it may not be a bad idea to think of population growth as an exogenous variable that is driven by features other than per capita income. At any rate, in societies that are not overwhelmingly poor, it is probably the case that if population growth is endogenous, it is a *decreasing* function of per capita income,[11] and not increasing as Malthus suggested. Data such as those presented in Table 9.1 certainly support this hypothesis better than the alternative.

Using growth models

The growth models of Chapter 3 represent a good starting point in this respect. Recall the ingredients of the standard growth model: people make consumption and savings decisions. Savings are translated into investment,

[11] To be more precise, this is true if per capita income is a good proxy for other features of development. On these matters, see the discussion in Chapter 2.

and the capital stock of the economy grows over time. Meanwhile, the population of the economy is growing too.

We know already how to figure out the net effect of all this. The rate of savings determines, via investment, the growth rate of the capital stock. The latter determines, via the capital–output ratio, the growth rate of national income. Does all this growth translate into an increase *per person*? Not necessarily. Population is growing too, and this increase surely eats away (so far as *per capita* growth is concerned) at some of the increase in national output. In Chapter 3, we did the simple algebra that puts these features together. Our first pass at this brought us to equation (3.6), which is reproduced here:

$$s/\theta = (1 + g^*)(1 + n) - (1 + \delta), \tag{9.2}$$

where s is the rate of savings, n is the rate of population growth, δ is the rate of depreciation of the capital stock, and g^* is the rate of growth of *per capita* income. This is the Harrod–Domar model, and the implications are crystal clear. According to this model, population growth has an unambiguously negative effect on the rate of growth. To see this, simply stare at (9.2) and note that if all parameters remain constant while the rate of population growth n increases, the per capita growth rate g^* *must* fall.

Nonetheless, we can criticize this prediction. The Harrod–Domar model, on which (9.2) is based, treats the capital–output ratio as *exogenous*, and therefore makes no allowance for the fact that an increased population raises output. After all, if the capital–output ratio is assumed to be constant, this is tantamount to assuming that an increased population has no effect on output at all. Would it not be the case that a higher rate of population growth would bring down the amount of capital needed to produce each unit of output, now that there is more labor as an input in production?

We have walked this road before; it leads us to the Solow model. In Solow's world, a *production function* relates capital and labor to the production of output. In addition, there is technical change at some constant rate. We obtained the remarkable answer in that model that *once the change in the capital–output ratio is taken into account*, the steady-state rate of growth is *independent* of the rate of savings and the rate of population growth (see our analysis in Chapter 3). All that matters for long-run growth is the rate of technological progress!

This is odd, because the Solow model now seems to tilt us to the other extreme. It suggests that population growth has no effect at all! However, this is not true: what we have shown so far is that population growth has no effect on the long-run rate of per capita income *growth*. There is a level effect, however. We briefly recall the discussion from Chapter 3.

Recall why population growth rates have no growth effect. In the Harrod–Domar model, there is an implicit assumption that labor and capital

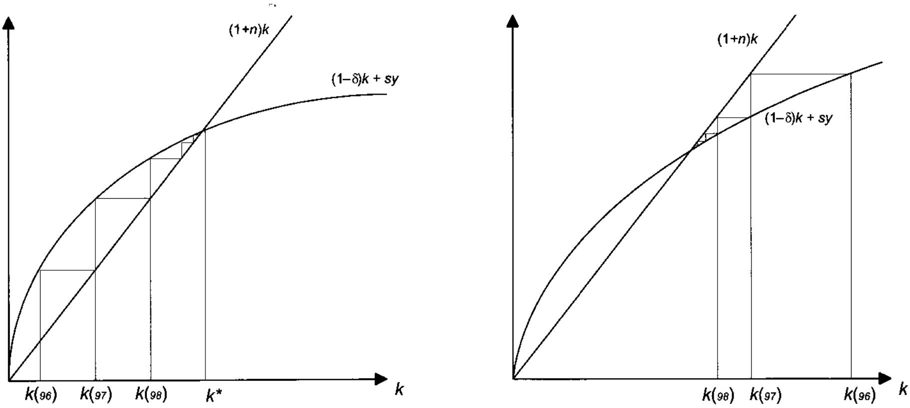

Figure 9.4. The steady state in the Solow model.

are not substitutable in production. Thus added population growth exerts a drag on per capita growth, while contributing nothing of substance via the production process. In the Solow model, on the other hand, population growth, while continuing to have the first effect, contributes to productive potential as the extra labor force is absorbed into productive activity through a change in the capital–labor ratio. Indeed, implicit in the Solow model is the assumption that capital and labor can be substituted for each other indefinitely, although the process of substitution may become more and more costly.[12] Because of this, population growth has no ultimate effect on the rate of growth in the Solow model.

This does *not* mean that an increase in the rate of population growth has *no* effect at all in the Solow model. It lowers the steady-state *level* of the per capita capital stock, expressed in units of capital per effective unit of labor, and in this way affects the level of per capita income, expressed again in units of effective labor. The easiest way to see this is to recall Figure 3.4, which we reproduce here as Figure 9.4.

Recall that the steady state k^*, expressed in terms of effective units of labor, is found as the intersection of two graphs. These are, respectively, the left- and right-hand sides of the equation that describes the evolution of capital stocks in the Solow model with technical progress:

$$(1+n)(1+g)\hat{k}(t+1) = (1-\delta)\hat{k}(t) + s\hat{y}(t). \tag{9.3}$$

It's now easy to see that if n goes up, this "swivels" the left-hand side of (9.3) upward and brings down the steady-state level of the capital stock, expressed as a ratio of effective labor. This means that although the long-run rate of growth is unaffected by a change in the rate of population growth, the

[12] The cost, of course, is expressed by the marginal rate of substitution between the two inputs of production and is captured by the degree of curvature of the production isoquants.

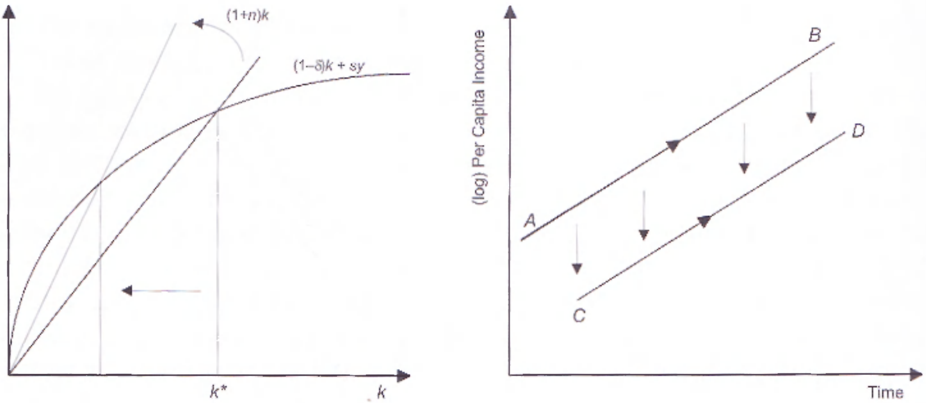

Figure 9.5. Growth rates are unaffected, but the levels shift down.

entire *trajectory* of growth is shifted downward. See Figure 9.5 for a depiction of this scenario.

Thus increased population growth has negative level effects in the standard growth models of Chapter 3. These effects are perfectly intuitive, although as we have seen, they may manifest themselves differently in different models. Population growth means that a given level of output must be divided among an increasing number of people, so that an increase in population growth rates brings down the size of the per capita cake. In the Harrod–Domar model, the effect is resoundingly negative, because population growth has no redeeming features, such as increasing the productivity of capital when more labor is available. In the Solow model, this redeeming feature is partially present. An increase in the population growth rate *both* increases the demands on the national cake and expands the ability of capital to produce the national cake. The net effect on long-run per capita growth rates is zero. Nevertheless, the level of per capita income at any given point in time is lowered.[13] This comes from the assumption in the Solow model that there are diminishing returns to every input, so that an increase in the labor intensity of production (necessitated by increased population growth) reduces the long-run per capita level of output relative to efficiency units of labor.

The growth models teach us that once the Malthusian assumption of unrestrained procreation is dropped, population growth certainly does not condemn society to everlasting subsistence. Growth in per capita income is still possible. At the same time, increased population growth does adversely affect this growth rate (as in the Harrod–Domar model), and even if it keeps the long-run growth rate unchanged, it affects the level of the trajectory (as in the Solow model).

[13] These last two statements are perfectly consistent, as Figure 9.5 shows.

Population and savings

There is yet another negative effect of population growth that is not considered in the growth models already presented, but is easy enough to incorporate. Faster population growth lowers the aggregate rate of savings. This happens simply because population growth eats into aggregate income. If it is true that the rich save a higher fraction of their income, savings rates may be adversely affected. More importantly, faster population growth shifts the age structure of the population toward the very young and in so doing increases the dependency ratio in families. Because children consume more than they produce, this tends to lower savings rates as well. This is one of the effects emphasized by demographers Coale and Hoover [1958] in their classic work on the subject.

The savings effect works in very much the same way as the direct population growth rate effects. In the Harrod–Domar model, it exacerbates the reduction in growth rates [allow s to fall as well in equation (9.2)]. In the Solow model, there continue to be no growth effects, but the long-run time trajectory of per capita income is shifted down.

Population, inequality, and poverty

A high rate of population growth will exacerbate the poverty problem, if the arguments in the previous section are valid. It will also worsen inequality if population growth among the poor is disproportionately larger.

Do the poor have more children? From the discussion in the previous sections of this chapter, that would appear to be the case, although the connections are not unambiguous by any means. It is more likely that the poor need children for old-age support. It is more likely that infant mortality rates are higher for the poor, so having a larger number of children to compensate is more likely to occur for the poor. We already know that this will translate into a higher expected number of *surviving* children (because risk-averse couples generally overcompensate for these risks).

It is somewhat harder to compare the relative costs of child bearing. Poor families are likely to have a higher degree of labor force participation by females, simply because additional income is of greater importance. This raises the opportunity costs of having children. However, it is also true that growth in income creates a quantity–quality trade-off in children. Richer households may want to invest proportionately greater sums in the education of their children. Consequently, the costs of an additional child (given the quality considerations) are proportionately much higher, which brings down the total number of children desired.

These considerations suggest that the poor may have higher fertility rates than the rich. To the extent that this is true, a high overall rate of population growth will have a disproportionately heavy impact on those who are already poor, or on the threshold of poverty.

Population growth and the environment

Recall the discussion on whether fertility is too high. In that discussion, one of the most important features is the underpricing of infrastructural resources. Government-provided education, health, and public transportation may all be subsidized. We also discussed why they are subsidized: it may be a second-best way to transfer resources to the poor. (Direct transfers may be infeasible because it may be impossible to credibly identify the poor.)

This observation has two corollaries. First, these resources must be consumed largely by the poor. Second, the inability of individuals to internalize the costs of these resources leads to higher fertility and consequent increased pressure on those very resources.

Under pricing arguments are not restricted to infrastructure alone. They apply to resources such as the commons (grazing land, fish stocks, groundwater) and the environment (forest cover, pollution, the ozone layer). Population growth places additional pressure on these scarce resources. Moreover, growth theory cannot be profitably applied to many of these resources: having more people around does not "produce" more forests, fish, water, or ozone. The effects are therefore stronger and more immediate.

9.4.2. Some positive effects

In the previous section, we began with the naive argument that all that population growth does is eat into available production. This is implicit in the Harrod–Domar model, for instance, but we know that population growth means a larger labor force, which contributes to additional production. Thus, at the very least, we have a tussle between the productive capabilities of a growing population and its consumption demands. The Solow model captured this well. Long-run growth of per capita income is unchanged in the Solow model because these two forces balance each other. We did note the existence of a level effect: there is more labor relative to capital on the long-run growth path. This brings down the *level* of income measured per unit of (effective) labor. This is an example of diminishing returns to labor at work. A higher ratio of labor to capital reduces its average product.

However, is that all labor is good for: production? In some broad sense, the answer is yes, but it is useful to return to a distinction between two notions of production: production using the same set of techniques, as embodied by the production function or technical know-how at any one point of time, and the production, invention, or application of *new* methods; in short, technical progress. Put another way, the pace of technical progress may be endogenous in the sense that it is affected by population size. Although we have discussed the endogeneity of technical progress before (see Chapter 4), the demographic effect on population growth merits additional attention.

The effect of population growth on technical progress can in turn be divided into two parts. First, population growth may spur technical progress out of the pressures created by high population density. This is the "demand-driven" view explored by Boserup [1981]. Second, population growth creates a larger pool of potential innovators and therefore a larger stock of ideas and innovations that can be put to economic use. This is the "supply-driven" view taken by Simon [1977] and Kuznets [1960].[14]

Population, necessity, and innovation

Necessity is the mother of invention, and population pressure has historically created necessity. Nowhere is this more true than in agriculture, where increasing populations have historically placed tremendous pressure on the supply of food. It is certainly the case that such pressure was often relieved by the Malthusian weapons of famine and disease that wiped out large sections of the population. However, it is also true that scarcity drove man to innovate, to create, or to apply methods of production that accommodated the increased population by a quantum jump in food output.

Several indicators permit us to see evidence of this even in today's world. Boserup [1981] classified countries into different grades by population density: *very sparse*, between 0 and 4 people per square kilometer; *sparse*, between 4 and 16 people per square kilometer; *medium*, between 16 and 64 persons per square kilometer; *dense*, between 64 and 256 persons per square kilometer; and *very dense*, 256 persons per square kilometer and upwards.[15]

Now consider an indicator such as irrigation. Which countries have more of it? Not surprisingly, the high-density countries do: in 1970, all the countries in Boserup's sample (of fifty-six) with more than 40% of the arable land under irrigation were dense or very dense countries, in the sense defined in the previous paragraph. Alternatively, consider the use of chemical fertilizer: it increases systematically with population density. In addition, study multicropping: four out of five very dense countries (in a sample of twenty-four) had more than 50% of the land devoted to multiple cropping; no other country in the sample exhibited this sort of ratio. More generally, Boserup suggested the pairing of population densities and food supply systems shown in Table 9.4 as a summary of her overall observations.

The point to be made is simple, perhaps obvious. At least in agriculture, high population densities go hand in hand with technologically more inten-

[14] The demand-driven story was studied in the context of a formal model by Lee [1988]. The supply-driven story was similarly explored by Kremer [1993].

[15] This scale, which is actually a coarsening of a finer division used by Boserup, is logarithmic, like the Richter scale for earthquake intensities. Each higher category used by Boserup has twice the density of the category immediately preceding it.

Table 9.4. Population densities and food supply systems.

Very sparse	→	Hunting and gathering, pastoralism, and forest fallow (one or two crops followed by a fallow period of around two decades)
Sparse and medium	→	Bush fallow (one or two crops followed by a fallow period of around one decade)
Medium	→	Short fallow (one or two crops followed by one or two years of fallow) with domestic animals
Dense	→	Annual cropping with intensive animal husbandry.
Very dense	→	Multicropping

sive forms of farming. This by itself isn't proof that such techniques were actually *invented* in high-density societies, although they almost surely were, but it does suggest that these methods, even if they were universally known, were applied more frequently in high-density societies.[16]

Agriculture is a leading example of how population growth stirs up technical progress, but it is not the only example. Here is Boserup again [1981, p. 102]:

> The increasing population density in Europe facilitated development of specialized crafts and manufactured goods. In areas of dense population, a large number of customers lived within a relatively small territory. Direct contact with customers was possible and transport costs for products could be kept at a minimum. Manufacturing industries... required skilled workers and traders as well as the financial services and administrative skills which were concentrated in urbanized areas. Therefore, the areas in Europe which first developed manufacturing industries were those with the highest population densities—Tuscany and the Low Countries . . . . Such concentration occurred only later in France and England.

The argument thus far is quite clear, but what is unclear is how we evaluate it. The first major problem is that much of what is attributed to population growth can also be attributed to increased per capita income. Income creates demand just as population might, and it is a combination of the two that is likely to drive innovation, or at least the sort of innovation that is motivated by the desire to make economic profit. Put another way, an increased population might correspond to a greater social need, but that need must be manifested in *economic* demand through the marketplace for innovators to respond. The income aspect of the phenomenon possibly acquires greater relative importance as basic needs (such as food) cease to pose a threat: it is

[16] More detailed and careful analysis of this theme must correct for the simultaneity of population and technique observations: it is possible, though unlikely, that some other set of forces (such as exogenous invention) first drove the adoption of certain methods of farming, which then increased the carrying capacity of that society. This alternative cannot be logically ruled out in the way in which Boserup presents the data.

hard to imagine how a larger population per se could spur innovations pertaining to more sophisticated products unless there is additional income to spend on such products.

The second problem with the demand-driven story is that it predicts some degree of cyclicity in per capita incomes: innovations raise per capita income as production levels kick up following the innovation, but a long hiatus should follow as population swells to bridge the newly created gap, with per capita incomes falling once again until the pressure of resources triggers another bout of innovation. As we shall see in the next section, this is not the sort of long-run pattern that we observe.

Both these points are connected to the observation that population *alone* is unlikely to be a major force on the side of demand for innovation unless we are in a world in which the innovator is himself directly affected by the population pressure. Early agriculture, in which the farmer and the innovator were often the same person, is probably the only persuasive example of such a phenomenon. Once the innovator is separated or relatively insulated from the overall pressures of population, it takes market forces to trigger innovative activity, and population growth by itself is not enough.

Population, diversity, and innovation

A large population is a diverse population, and the chances are higher that someone will be lucky enough or smart enough to come up with an idea that benefits everybody else.[17] This is the gist of the supply-driven argument. The easiest way to appreciate this line of reasoning is to imagine that everybody has an *independent* chance of coming up with a idea that will benefit the rest of the human race. It is immediate how things progress in this situation: the larger the population, the larger would be the number of people that have useful ideas, and so the higher is the rate of technical change. There are several senses in which we might want to qualify this statement, and we will, but let us stick with it for now.

We combine this statement, which is about how technical progress reacts to population, with a statement about how population might respond to technical progress. Specifically, let us suppose that population growth increases with per capita income up to a point and then falls. This is a crude version of the demographic transition that we have already used in a different context (see Chapter 3). The left-hand panel of Figure 9.6 depicts a typical curve that might relate per capita income to population growth.

Now let us begin the analysis by considering an initial level of per capita income that is so low that population growth increases with per capita income. This means that we are currently on the upward-sloping segment of

[17] The discussion here follows Kremer [1993].

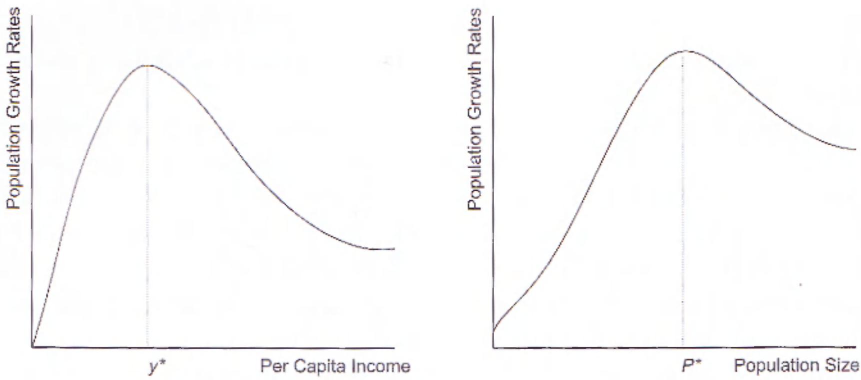

Figure 9.6. Population growth, per capita income, and population levels.

the curve in the left-hand panel of Figure 9.6. Population is growing, and it follows from our first postulate that the pace of technical progress must accelerate. Observe that per capita income *cannot* be stagnant during this phase. To prove this, suppose by way of contradiction that per capita income is unchanging. Then, after some time passes, population is higher, so that technical progress is higher. Consequently, the population growth rate now required to keep per capita income stagnant must be even higher. However, the only way to accomplish this is to increase per capita income,[18] which contradicts our supposition that per capita income was stagnant all this time.[19]

Thus as long as we are on the upward-sloping part of the curve, per capita income must rise and so must the *rate* of population growth. Thus during this phase we obtain the prediction that the population growth *rate* is increasing with the size of the population. This state of affairs continues until we reach the point at which population growth rates begin to decline in income. As long as growth rates are positive, however, the population will still grow, so that technical progress will continue to accelerate. Coupled with a diminishing pace of population growth, this implies an acceleration in the long-run rate of growth of per capita income. Thus population growth rates decline even faster. This period is therefore associated with a leveling-off and consequent decline in the rate of growth of the population. No longer will population growth rates increase with population: they should decline.

To summarize, then, if technical progress is "supply-driven" by the population, then population growth should initially be an *increasing function* of population itself, but this trend should reverse itself after some stage. The right-hand panel of Figure 9.6 puts these observations together diagrammatically: P^* is the threshold level of population that permits technical progress at a rate such that the threshold per capita income of y^* (see the left-hand

[18] This is because we are on the upward part of the curve in Figure 9.6.

[19] This does not rule out the possibility that there might be an initial phase in which per capita income can actually fall, but this phase must be temporary: see Kremer [1993] for a rigorous analysis.

panel) is just reached: after this point population growth rates turn down as per capita income climbs even further.

Is there any empirical truth to this assertion? Figure 9.7, which is taken from Kremer [1993], puts together various estimates of world population and the implied annual growth rates from 1 million B.C. to 1990. Observed population growth rates are plotted against the baseline values of population. Clearly for much of recorded history, population growth rates appear to be increasing with population size. The trend reversed itself only after population passed the three billion mark, which is circa 1960.

Thus the simple model of "supply-driven" innovation works surprisingly well. It does predict the same qualitative shape that we observe in the data, but we need to be careful about taking this model too literally. For instance, as set out now it also predicts that countries with large populations should exhibit a high rate of technical progress. That is, the time-series prediction passes over into a cross-section prediction, which is far more dubious, to put it mildly. However, a simple extension of the model can be used to account for this seeming discrepancy: simply allow technical progress to be a function not just of population size, but also of the per capita income of the society. After all, it takes brains coupled with economic resources to carry out useful scientific research. With this modification in place it is evident that the strong (but wrong) cross-section prediction disappears, but the time-series prediction survives unscathed. After all, if per capita income increases over time (as we argued that it will in this model), then this extension cannot overturn the qualitative features discussed earlier.

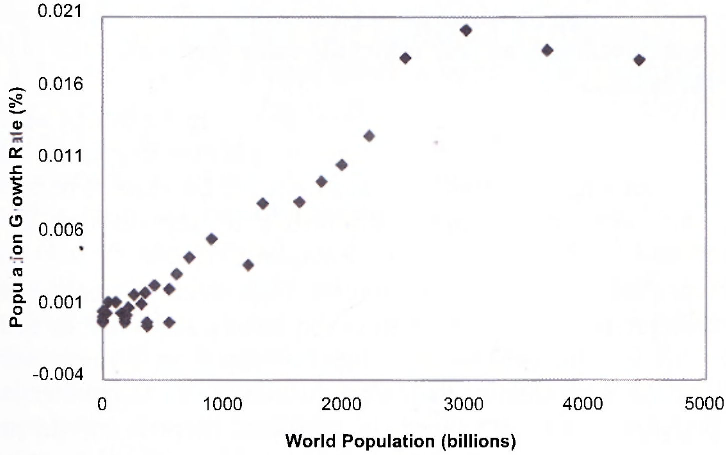

Figure 9.7. Population growth in world history. Source: Kremer [1993, Table 1 and Fig. 1].

9.5. Summary

In this chapter, we studied population growth and its interaction with economic development. Our goals were as follows:

1. To describe and understand the observed patterns of population growth in different countries; in particular, the phenomenon of the *demographic transition;*

2. To analyze the social and economic factors that affect fertility decisions at the level of the household;

3. To take note of possible features that create a systematic bias between levels of fertility that are *privately* optimal for the household, as opposed to those levels that are *socially* optimal; and

4. To understand the impact of population growth on economic development.

We began with terminology—the concepts of *birth rates* and *death rates*—and took a look at data on these rates for various countries. In the poorest of countries, both birth and death rates appear to be relatively high. Then death rates fall, while birth rates still remain high. Finally, countries that have higher per capita income or have made systematic efforts to control population growth exhibit birth rates that are also low. Thus population growth rates, which are just the difference between birth rates and death rates, seem to rise and then fall over the course of development.

We observed that the concept of an *age distribution* is important in this regard. Fast growing societies are also young societies, and this feature both reinforces a high birth rate (relatively large numbers of people are entering reproductive age) and keeps death rates somewhat lower than they really are in age-specific terms (because most people are young). Thus a policy that brings down the *total fertility rate* may still cause population to overshoot a desired target because of inertia.

From these various observations comes the central concept of the *demographic transition*, which is a description of three phases. In phase 1, both birth and death rates are high. In phase 2, death rates fall because of improvements in hygiene, sanitation, and medicine, but birth rates remain high. In phase 3, birth rates follow the death rate on its downward course.

The observation that birth rates remain high even as death rates fall is central to understanding the population explosion, not only in today's developing world, but historically as it has happened in Europe. Why don't birth rates decline with death rates? One answer is the *macro-inertia* of birth rates in a population that is young. In addition, there is inertia at the *micro* level as well. If the flow of current demographic information is limited, couples might instinctively use the experience of generations before them to make their decisions regarding fertility. But limited information is not the only factor that causes a slow reduction in fertility. It turns out that missing markets, notably, the missing market for old-age security, are also central to our understanding of this phenomenon. We studied the connection between the desire for old-age security, mortality, and fertility choices. Gender bias plays an important role here: to the extent that a family desires sons, this can greatly increase fertility rates. In general, these are factors that collectively shed some light on the stickiness of birth rates.

We then turned to factors that cause a systematic deviation between decisions that are privately optimal (from the point of view of the family) and decisions that are optimal from the point of view of society. Lack of information and uncertainty play a major role here, as does the presence of *externalities*, both within the household and across households. Externalities arise because there is a divergence between the social and private costs (or benefits) of having children. In this context, the role played by joint families or by externalities that are environmental or employment related are of great importance. Social norms that create a high degree of conformism to exaggerated rates of fertility are relevant as well.

Finally, we turned to an analysis of the effects of population growth on economic development. Both negative and positive effects coexist. The simplest negative effect comes from the observation that population growth eats away at a given level of resources or income, leaving less per head to go around. This sort of prediction naturally emerges from the Harrod–Domar model of economic growth, in which labor is not regarded as an essential input of production (recall that the capital–output ratio is fixed in that model). However, this observation is unduly naive, for the simple reason that an increased population means more labor input, which expands production. In the Solow model, these two effects cancel exactly as far as long-run growth rates are concerned: these rates are unaffected by the pace of population growth. However, there is a level effect: a higher rate of population growth pushes the economy onto a lower trajectory of per capita income (with the same growth rate as before). This summarizes the consensus argument that population growth is unambiguously bad for economic development.

However, there are positive arguments as well—two of them—and it is with these that we finish the chapter. One view is that population growth creates economic necessity, which forces the adoption or creation of new ideas that expand carrying capacity. This is a demand-side argument: population growth fosters spurts of development because the pressures that it creates necessitate bouts of innovation. The second view is a supply-driven argument: population growth fosters development simply because many heads are better than one. If we think of each human being as a repository of ideas then more human beings means more ideas that can be put to use for the economic benefit of mankind. Thus the rate of technical progress should increase with population size. We examine these arguments in some detail at the end of the chapter.

Exercises

■ (1) Review the concepts of birth rates, death rates, and age distributions, and the way in which these notions interact with one another. Construct an example of countries A and B, where A has higher death rates than B in *every* age category, yet has an overall lower death rate.

■ (2) Why does a young age distribution make it more difficult for a country to slow its rate of population growth? If a country suddenly drops its total fertility rate to two (which makes for a stationary long-run population), describe the path that population will take before settling at this long-run level.

■ (3) Discuss factors that have altered the carrying capacity of the planet. Explain how such increases in carrying capacity might ultimately affect fertility decisions at the level of the household.

■ (4) Explain why each country might want to take a pro-natalist stand for military or political reasons, but the combination of *all* countries taking the same pro-natalist stance may make all countries worse off relative to a neutral stance on population.

■ (5) In this chapter, we studied a model where a family wants one surviving child to provide old-age security. Let us say that the probability of any *one* child living to look after its parents in old age is 1/2 (i.e., 50–50). However, the family wants this security level to be higher, say a probability of $p > 1/2$.

(a) Describe the family's fertility choices for different values of p, by first writing down a model that captures this story, and then examining the results for different values of p.

(b) Calculate the *expected* number of surviving children for this family, under various values of p. (For a definition of "expectation" see Appendix 1.)

■ (6) Review the concepts of *targeting* and *hoarding*. Discuss the various components of the "survival probability" p (see exercise 5 as well as the discussion in the text) that will affect the relative prevalence of these two forms of reproductive behavior.

■ (7) In a world in which families use the experience of their parents in determining their own fertility behavior, discuss the role of the media (such as television) in affecting fertility.

■ (8) Why is a well-implemented ban on child labor likely to reduce fertility rates?

■ (9) Organized social security, health care, etc., will lead to a fall in fertility rates but will lead to a fall in savings rates as well. The net effect on per capita income growth rates is ambiguous. Comment.

■ (10) In the land of Oz, there are *three* inputs to production: capital, physical labor, and mental labor. Men in Oz have more physical labor power than women, but both men and women have the same amount of mental labor power.

(a) Who earns more in Oz, men or women? What do these differences depend upon?

(b) Now imagine that the technology is such that more capital raises the marginal product of mental labor faster than it raises physical labor. As the economy of Oz grows over time, its stock of physical capital is steadily increasing. How would you expect the *relative* wage of men to women to change over time? Explain.

(c) Women have one unit of labor time that they can allocate between raising children and being part of the workforce. How would this allocation be affected by the changes over time that you found in your answer to (b)? Discuss the implications for fertility levels in the population.

■ (11) Studies for many countries have shown that labor force participation by women tends to have a U-shaped pattern with respect to growth in per capita income. Explain why this contrasts with your findings in exercise 10 above. Discuss reasons why such a U-shaped curve of female labor force participation may occur. The use of income and substitution effects will help you to formulate your answer.

■ (12) Evaluate the validity of the following statements.

(a) A developing country is likely to have an overall death rate that is lower than that of a developed country.

(b) The populations of Europe and North America grew at a combined rate between 1750 and 1900 that significantly exceeded the population growth rates of developing countries at that time.

(c) If country A has a population growth rate that is lower than country B, then the average woman in country A has less children than her counterpart in country B.

(d) Birth rates may be high even when death rates may be falling.

(e) If total mortality among children remained constant, but the *incidence* of that mortality shifted from late childhood to early childhood, fertility rates should decline.

■ (13) Review the data on the demographic transitions for one developed and one less developed country. For instance, you could study the demographic transitions of England and Sri Lanka (see Gillis, Perkins, Roemer, and Snodgrass [1997, Chapter 7]) and make sure that you understand the main trends in birth rates, death rates, and net population growth rates. Think about and explain the reasons why the picture for Sri Lanka looks more "compressed" in time than the corresponding picture for England.

■ (14) You are gathering demographic data in a village. You suspect that families have a *gender bias*; that is, they have children until a certain *target number of sons* are born, but you don't have direct evidence of this. All you have is information on the sex and birth order of each child born to each family in the village. How would you use the data to test your hypothesis that there is gender bias?

■ (15) Here is more on gender bias. In many Southern Asian countries, there is evidence that the ratio of boys to girls is too high (see also Chapter 8). A ratio of 110 boys to 100 girls is not uncommon. One obvious hypothesis that springs to mind is that girls are treated badly relative to boys (or are perhaps even selectively aborted or killed), so that their mortality rates are higher. There could be much truth in these assertions. Nonetheless, it is worthwhile to investigate alternative possibilities.

(a) Begin by looking at the target rule discussed in exercise 14. Simplify the target rule using the following assumption: families have children until they have just one boy, and then they stop. Then show that on average, larger families will have more daughters.

(b) Now use the insight of part (a) to argue that in poor societies, girls might have a higher death rate than boys even if there is no discrimination.

■ (16) This is a question on joint families, externalities, and fertility choice. Suppose that Ram and Rani are the heads of a nuclear family, making their fertility decisions. For simplicity, assume away gender bias and issues of child survival. The following table details the costs and benefits (in dollars, say) of different numbers of children.

(a) Based on the information in the table, how many children would Ram and Rani have in order to maximize their net benefit?

(b) Now consider two identical nuclear families: Ram and Rani (as above), and Mohan and Mona. Ram and Mohan are brothers and the two couples form a joint family. Both couples have exactly the same costs and benefits of having children as in the table. Now suppose that 50% of the upbringing costs of each child (e.g., child care) can be passed on to the other family.

Number of children	*Total benefit ($)*	*Additional cost*
One	500	100
Two	750	100
Three	840	100
Four	890	100
Five	930	100
Six	950	100
Seven	960	100
Eight	960	100

Each couple makes independent decisions, taking only its own welfare into account. Now how many children will each couple have?

(c) Explain the reason for this seemingly paradoxical result, using the concept of externalities, and try and understand why larger families (either integrated across generations or between siblings in the same generation), will tend to have a larger number of children per couple.

■ (17) Discuss the impact of population growth on per capita income and its growth.