Ricerche Economiche (1995) **49**, 179–205

Dynamics of demographic development and its impact on personal saving: case of Japan†

ALBERT ANDO‡, ANDREA MORO, JUAN PABLO CORDOBA AND GONZALO GARLAND

Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 1904-6297, U.S.A.

Summary

A dynamic model of the demographic structure of Japan is summarized. It is capable of tracing the dynamic development of the Japanese population, including the distribution of families by age, sex, and marital status of the head, as well as by the number and age of children and other dependents. This model is combined with specification of the processes generating family income and consumption, and then used to generate the pattern of aggregate income, saving and asset accumulation for the period 1985–2050 under alternative fertility assumptions. The results suggest that the saving–income ratio for Japan will increase slightly in the immediate future as the number of children per family declines sharply, and then falls moderately as the proportion of older persons in the population increases. Qualitative results depend critically on the labour force participation rate of older persons and on the probability of older persons merging into younger households.

J.E.L. Classification: D12, E21, J10.
Keywords: Saving, demography, aggregation, Japan.

1. Introduction

For a national economy to follow a path of orderly economic progress, one of the essential requirements is that its needs for

† Part of the data used in this analysis is cohort means computed from individual returns from the National Survey of Family Income and Expenditure conducted by the Statistics Bureau of the Japanese Government, to which Ando had access as a member of a team headed by Fumio Hayashi, at the Department of Economics, Osaka University. While we would have preferred to work with somewhat differently organized data, we no longer have individual data and results reported here are restricted to those producible based on the original summary tabulations of the data.

‡ Research also done by A. Ando at National Bureau of Economic Research.

0035–5054/95/030179 + 28 $12.00/0

capital accumulation are more or less matched by the saving generated by society when resources in the economy, especially labour, are fully employed. Since most economists would consider the positive relationship between the growth rate of output in the economy and the investment–output ratio to be a fairly natural implication of production technology and the rational behaviour of producers, the above requirement suggests that the saving–income ratio and the rate of growth of income should be positively related in a well functioning economy.

The above proposition for producers applies to individual firms more or less uniformly since firms do not have obvious phases of a life cycle except at the starting point and those related to the nature of vintage capital owned by them. The positive relationship between the investment–output ratio and the rate of growth of output is therefore thought to hold, at least qualitatively, for the aggregate of all firms in an industry or an economy as well.

The relationship between the rate of growth of income and the saving–income ratio for households is more complex. Other things being equal, the higher the expected rate of growth of income in the future, the smaller we expect the current saving–income ratio to be. On the other hand, given the current and expected future level of income, the higher the rate of growth of income has been in the past, the higher the current saving–income ratio is likely to be. This is because, if the past rate of growth of income is higher given the current level of income, that is, if the past level of income is lower, it may be presumed that the accumulation of wealth has been lower, and the lower level of wealth reduces the level of current consumption.†

Thus, there is little uniformity among households in the relationship between the growth rate of income and the saving–income ratio. Furthermore, each household in society is in a particular phase of its life cycle, and this and other demographic characteristics of each household, such as its marital status, number of children, other membership in the family, have major effects on its saving–income ratio. The aggregate relationship between the saving–income ratio and the rate of growth of income, therefore, is a result of the aggregation process rather than a reflection of a uniform feature of individual households.

Perhaps the best known theoretical construct explaining the positive correlation between the growth rate of aggregate income and the saving–income ratio was offered some years ago by Modigliani and Brumberg (1980) as an implication of their original formulation of the life cycle theory.

For our purpose, we find it convenient to generalize their proposition somewhat by dividing all households in society into cohorts

† There may be an element of habit persistence as well.

defined not only by the age of the head but also by other demographic characteristics such as the marital status and sex of the head and the number of children, and note the following definitional relationship:

$$\frac{S(t)}{Y(t)} = \sum_a \sum_f w(t, a, f) \frac{y(t, a, f)}{Y(t)} \cdot \frac{s(t, a, f)}{y(t, a, f)} \qquad (1)\dagger$$

where $S(t)$ and $Y(t)$ are aggregate saving and disposable income of the household sector, $w(t, a, f)$ is the weight (the total number of families in the cohort defined by the age of the head being a and its demographic characteristics being f), and $y(t, a, f)$ and $s(t, a, f)$ are the mean values of disposable income and saving for the cohort (a, f) in period t. Note that f is a vector, so that the summation sign above f is in principle multiple summations. This definitional relationship makes it clear that movement of the aggregate saving–income ratio over time can be decomposed into three factors: movements of the relative size of cohorts (w), the relative level of mean income among cohorts (y/Y), and the relative size of the saving–income ratio among cohorts (s/y).

In terms of equation (1), the Modigliani proposition is that s/y is positive for younger households and very small or negative for older ones, and hence a shift of w from younger cohorts to older ones (due to slower growth of population) would reduce saving, and so does a shift of income from younger cohorts to older cohorts (due to a slow down of generation specific productivity increase).‡

† This formula does not apply when $y(t, a, f)$ is zero. For such cohorts, the cohort saving must be expressed directly as:

$$w(t, a, f)s(t, a, f)$$

and then divided by $Y(t)$ to be included in the summation on the right hand side of equation (1). Since $y(t, a, f)$ is a cohort mean and not income of an individual household, it is very rare that $y(t, a, f)$ is in fact zero.

‡ A substantial part of the current literature on aggregate consumption and saving behaviour is based on the assumption of a representative consumer. In such a model, the only productivity increase that can be accommodated must result in a faster rate of growth of income for the representative consumer himself. We refer to this type of productivity increase as "calendar year specific" productivity growth. Modigliani, on the other hand, had in mind a generalized overlapping generations model in which each family lives not two but a multiple number of periods, and when he talked of productivity increase, he supposed that each generation had a fixed pattern of productivity over its life that was not subject to change once the family began its working life, but that the younger generation was always more productive than the preceding one. We refer to this type of productivity increase as "generation specific" productivity growth. We believe that, theoretically, both types of productivity increase may coexist. However, we seldom observe significant dissaving by young families in societies where the aggregate productivity increase is very rapid, such as Japan and Italy from 1960 to 1975. See Ando, Guiso and Terlizzese (1994*a*). Hence, in this paper, we proceed assuming that the aggregate impact of the negative correlation between the saving–income ratio and income growth within a family is not dominant.

Let us digress to the U.S. case where some familiar attempts somewhat similar to this inquiry already exist. In their recent papers Bosworth, Burtless and Sabelhaus (1991) and Cannari (1994) have investigated whether or not the decline in the saving–income ratio in the U.S. and in Italy during the 30 years ending in 1990 could be attributed to shifts in weight due to a changing demographic pattern of these countries. They focused their attention almost exclusively on the age distribution of the head of households and did not find that shifts in weights can account for the decline in the aggregate saving–income ratio.

Gokhale, Kotlikoff and Sabelhaus (1994) find that the recent decline in the saving–income in the U.S. can be attributed to a shift of resources from younger families to older households, and to an increase in the consumption–income ratio of these older households when their income and consumption is defined appropriately, due in large part to a rise in medical costs. We are inclined to agree with Gokhale *et al.* (1994) that the earlier papers were too narrowly focused.

In parallel to Gokhale *et al.* (1994) we present in Appendix 1 a table suggesting that at least a part of the decline in the saving–income ratio in the U.S. during the past 30 years is due not so much to the ageing of the population, but to a shift in the relative weights among family types, from normal families (headed by a head and his/her spouse) to single parent families and to single individuals.

As the birth rate in most OECD countries declines, and the size of the older population becomes larger relative to the younger, working population, economists and the public in general are increasingly concerned that the saving–income ratio may decline sharply and cause serious disequilibrium in these countries. This is perceived as the basic prediction of the life cycle theory of saving. But such a major change in demographic structure is likely to be accompanied by other shifts, such as a decline in the number of children per family, changes in the social convention of work, for example, higher labour force participation of older individuals and perhaps women, counteracting some effects of the ageing of the whole population. In other words, shifts in weights in equation (1) may be accompanied by changes in the distribution of income and the saving–income ratios of various cohorts. In this paper, we propose a framework for dealing with their simultaneous movements in response to changes in demographic patterns, and present some results of analysing a relatively simple case for Japan.

We begin our discussion by briefly looking at what the distribution of weights, income, and the saving–income ratio over cohorts in the base year, 1985, looked like in Japan. We then summarize our model of demographic dynamics to determine the future values of the weights, and our hypotheses for the determination of the distribution of income among cohorts and of the saving–income ratio for individual cohorts. Finally, we report the results of simulating these models together for a sufficiently long period of time to see the consequences of alternative assumptions concerning demographic dynamics.

2. A preliminary view of the data

The basic data sets for our analysis are the National Survey of Family Income and Expenditure for Japan, 1984, and the Final Report of the Population Census of Japan, 1985. The former is a very detailed survey of income and expenditures of families based on a sample of some 54 000 households, and it includes fairly detailed information on the demographic characteristics of households as well as summary information on their assets and liabilities. For the analysis presented here, we have used a combination of estimation results using individual returns obtained in our earlier project and some published tabulations.

In Table 1, we show income, the saving–income ratio, and the net worth–income ratio for Japanese families in 1985 by age and family type. By normal family, we mean all those families headed by a married couple, including nuclear families (those consisting of a married couple and their children 18 years of age and younger) and extended families (nuclear families plus additional dependents, such as grown children, parents of the couple, etc.). We see that those families who belong to the "normal family" group continue to earn sizeable income after they reach the age of 60 and 70, and they continue to save. Looking at these families, one may be tempted to conclude that the life cycle hypothesis does not hold in Japan.

A careful investigation of older households in Japan reveals, however, that those older individuals who have in fact retired tend to merge into younger households, and disappear as independent units. Note that in this survey the person who is earning the highest income in the family is designated as the head of household. A large fraction of those older individuals who do not merge into younger households continue to work. We show in Table 2 the ratio of older individuals who have merged into younger households to the total number of older individuals as of 1979 and in Table 3 the labour force participation rate by age for 1984. We would like to call the reader's attention to major differences in the labour force participation rate between the U.S. and Japan. For males aged 70–74, for instance, the participation rate in the U.S. is 0·144, whereas it is 0·403 for Japan. This very high participation rate of older individuals in

TABLE 1
Saving–income and net worth–income ratios for detailed demographic group, 1985

Category	*Weight*	*Income*†	*Sav/inc*	*N.W./inc*
All people	38 318 998	453	0·131	5·274
Families, totals	28 001 322	524	0·128	5·665
Families, by age groups:				
<=29	1 587 121	316	0·041	2·201
30–39	7 647 673	455	0·108	4·361
40–49	7 974 612	569	0·119	5·294
50–59	6 222 498	629	0·136	6·276
60–69	3 316 032	516	0·176	8·053
>=70	1 253 385	404	0·194	8·610
Single head of household, by sex				
Males	457 418	468	0·154	6·747
Females	1 109 434	343	−0·076	6·606
Single head of household, by sex and age groups				
Males <=29	75 554	334	−0·061	6·337
30–39	144 523	456	0·144	5·732
40–49	52 620	534	0·194	5·747
50–59	63 517	595	0·225	6·611
60–69	54 144	516	0·248	6·829
>=70	67 061	433	0·137	10·471
Females <=29	34 750	206	−0·185	3·082
30–39	164 222	253	−0·134	4·117
40–49	319 567	344	−0·101	5·215
50–59	300 820	375	−0·149	7·311
60–69	213 451	382	0·025	8·350
>=70	76 624	354	0·148	8·820
Single, by sex				
Males	4 898 274	297	0·276	1·660
Females	3 852 550	164	−0·018	3·211
Singles, by sex and age groups				
Males <=29	3 473 326	259	0·218	1·252
30–39	757 023	398	0·370	1·806
40–49	229 651	476	0·415	2·329
50–59	192 884	483	0·421	2·925
60–69	114 622	280	0·303	3·667
>=70	130 769	163	0·005	4·847
Females <=29	1 771 643	155	0·028	0·977
30–39	209 050	183	0·031	2·220
40–49	191 133	192	−0·032	3·961
50–59	439 030	187	−0·062	4·708
60–69	741 567	167	−0·075	5·866
>=70	500 125	154	−0·061	5·432

† In ¥ 10 000

Table 2. *Number of persons living in younger households as a percent of total number of persons, (%) by age, 1979.*

Age	63	64	65	66	67	68	69	70	71
%	39	45	47	46	46	51	57	63	57
Age	72	73	74	75	76	77	78	79	80
%	66	68	71	71	79	72	78	80	77

Japan has a major consequence for the nature of the impact of the ageing of the population. An inspection of Tables 1, 2, and 3 makes clear the difficulty of interpreting the pattern shown in Table 1 as a simple saving and asset accumulation pattern over life prevailing in Japan.

The saving pattern of households headed by unmarried males is by and large similar to the normal family, and in any event they constitute a very small group of families. We may note that households headed by an unmarried male have a relatively large income and large asset–income ratio. For relatively younger cohorts, this arrangement appears to be the result of a young unmarried male with relatively high income living with his parents, and because of his high income, he is designated as the head of the household. Older ones, on the other hand, result from widowhood or from divorce, and again the single male in

Table 3 *Labour force participation rate by age and sex for Japan and the U.S., 1985.*

	United States		*Japan*	
Age class	*Male*	*Female*	*Male*	*Female*
15–19	—	—	0·173	0·163
20–24	0·753	0·641	0·708	0·693
25–29	0·871	0·657	0·938	0·509
30–34	0·898	0·654	0·953	0·478
35–39	0·905	0·676	0·960	0·564
40–44	0·902	0·682	0·962	0·644
45–49	0·888	0·644	0·957	0·648
50–54	0·846	0·581	0·944	0·585
55–59	0·761	0·479	0·880	0·490
60–64	0·532	0·322	0·699	0·366
65–69	0·236	0·130	0·567	0·260
70–74	0·144	0·074	0·403	0·151
75 and over	0·068	0·021	0·213	0·055

question becomes designated as the head only if his income is high. Others, for example, may have grown children whose income is higher and they become designated as heads. For male singles, we note that there are a significant number of them until about 40 years of age, and they save heavily, presumably in preparation for their marriage. In another paper, Ando *et al.* (1994*a*) show that similar male individuals living with their parents earning similar levels of income save even more.

Households headed by unmarried females, on the other hand, have much lower income and they dissave. The critical feature to be noticed here is that, in spite of negative saving, these households on average possess significant amounts of assets, and the asset-income ratio increases with age throughout all ages. This is a consequence of a complex process in which female headed households are created by divorce or widowhood, and the older the newly created unmarried female headed household, the larger their starting net worth. At the same time, existing unmarried female headed households disappear, partly because of remarriages if they are relatively young, partly through their merger into younger households especially if they are older and their net worth is small.

3. Income, consumption and asset accumulation

Income received by an individual household whose head is less than 63 years old is described by the following equation:†

$$\ln y(t, a, g) = \alpha_0 + \sum_i \alpha_{a(i)} a(i) + \sum_i \alpha_{g(j)} g(j) + \sum_i \sum_j \alpha_{a(i)g(j)} a(i) g(j) + \gamma \cdot t \qquad (2)$$

where y is income before income taxes, $a(i)$ represents a set of one–zero dummies indicating whether or not a household falls into the age class i, $g(j)$ also represents a set of one–zero dummies indicating whether or not a household possesses the j-th characteristic. In particular, we take into account the number of additional income earners in the family. α represents the estimated coefficients, t represents time, and γ is the rate of growth of income. The interaction terms represent primarily the fact that the age pattern of income depends on the occupation of the worker.

We first apply this equation to individual households in the sample and estimate the coefficients. Since the sample refers to a single year, for estimation purposes, we omit the term $\gamma \cdot t$ from the above equation, and γ is separately estimated. We must comment on two potential pitfalls of this estimation.

† Numerical results of the estimation of this equation are given in Ando and Moro (1995).

First, this equation ignores the response of the labour supply to the real wage rate. This seems justified for the main income earner who is primarily responsible for supporting the family, but for secondary workers, this may be a questionable assumption. Our effort to check on this question using two rounds of surveys separated by 5 years did not indicate that this is a serious problem. Second, if income and life expectancy are correlated, coefficients of variable a may be biased. In Ando, Guiso and Terlizzese (1994*b*) it is shown that the correlation between income and life expectancy does not appear to be present any more. In Japan, however, elderly individuals with relatively low income have a much higher probability of merging into younger households than high income ones. Thus, we must be careful to interpret the result of estimating the above equation as applying to those older individuals maintaining independent households. Income of those who merge into younger households but who continue to work is estimated through the coefficient of one of the g values.

For the purpose of estimating the consumption function, the prediction of equation (2) serves as the measure of current income. Future expected income of the household is constructed by applying equation (2) with the value of t increased successively, with appropriate choice of $a(i)$ taking on the value "1", and some of the g values taking on "1" or "0" multiplied by a probability. This process is continued until the head of the household reaches 63, and then the annualized present value of the sum of expected income thus estimated is defined as the expected future income.

Since the value of γ is not estimated, we have used several alternative assumptions in our analysis. The result reported here is based on the assumption that the value of γ is equal to the real interest rate implicit in the assumed value of the discount factor. In the remainder of this paper, we denote the prediction of equation (2) by a plain y, and the future expected income constructed as described above by ye.

For those families whose heads are 63 or older, we have estimated alternative equations for two distinct groups. For those families whose head is still working, we have estimated an equation similar to equation (2) above. For those families whose heads are fully retired, we have assumed that their pension income will remain the same in the future in real terms. For those who are newly retiring, we have assumed their starting pension to be the same as the level received by currently retired persons one year older, increased by the growth in average productivity per capita. If these fully retired families have income from capital, we have assumed that the same level of capital income in real terms will continue to be received until they merge into younger households. This assumption seems reasonable because those families who exhaust their wealth will merge into younger households, while those who remain independent appear to maintain their wealth without reducing it much.

We now come to the determination of the level of consumption. For those families whose heads are 63 or younger, we have estimated the following type of equation:†

$$\frac{c}{y}(i,j) = \beta_y + \beta_e \frac{ye}{y}(i,j) + \sum_i \beta_{a(i)} a(i) + \sum_i \beta_{a(i)v(i,j)} a(i) \frac{v}{y}(i,j) + \sum_j \beta_{f(j)} f(j) \tag{3}$$

where $a(i)$ are, as before, dummies indicating age class; $(c/y)(i,j)$ and $(v/y)(i,j)$ are, respectively, the mean consumption–income ratio and the mean net worth–income ratio for the subgroup of the population defined by the age class, $a(i)$, and having characteristic $f(j)$, while $f(j)$s are vectors of demographic and other characteristics of families. $g(j)$ and $f(j)$ are not the same set, but they may include common elements. βs are estimated coefficients.

Since equation (3) is not explicitly derived as a consequence of an optimization process from a well specified objective function, parameter estimates of a decision rule such as equation (3) are subject to doubt that they may not be well identified corresponding to parameters of the objective function, and that they may be subject to serious bias if they are interpreted as indicating the marginal effect of a change in the value of variables for which they are coefficients. On the other hand, it is doubtful that we can write a uniform objective function that applies to all households in a society, and that an estimate of parameters of such an assumed uniform objective function is meaningful (Kirman, 1992). For the purpose of our paper, we assume that a decision rule in the *form* of equation (3) is applicable to all individual households, although parameter values may vary from one household to another. Our estimates are then meant to be the weighted average of parameter values for individual households belonging to appropriate subgroups. We discuss briefly below the nature of some possible biases in our estimates of coefficients from a more practical point of view.

We first note that, because equation (3) is in a ratio form and does not contain a term in the form of $1/y$, it assumes homogeneity of degree one for consumption in y, ye, and v. This is an important point because we would be using this function for simulations lasting as long as 60 years, and even a very small deviation from homogeneity matters importantly in simulations over such a long period. We believe that this is a reasonable assumption on the basis of our survey of existing evidence, though the judgment may vary among students of the

† Numerical results of the estimation of this equation are given in Ando and Moro (1995).

subject. We offer one more piece of evidence in support of the homeogeneity which we have obtained from the data at hand.

If we introduce a term in the form of $1/y$ into equation (3) and reestimate the equation, this term acquires a marginally significant positive coefficient. The question then is whether this is a genuine indication of the presence of a non-homogeneity, or it is an evidence for some biases in our estimate, for example, resulting from errors of measurement for independent variables.

To answer this question at least partially, Ando, Yamashita and Murayama (1986) have estimated this equation using data from the 1974 and 1979 surveys, and Hayashi, Ando and Ferris (1989) did the same for the 1984 survey. Coefficients of equation (1) remained very stable for these surveys, while the coefficient for the term $1/y$ increased from 1974 to 1979 and again from 1979 to 1984 more or less in proportion to the movement of the mean value of c. This means that c, y, ye, and v are all increasing more or less proportionately from one survey to the next, indicating the long run homogeneity of a relationship like equation (3). We have decided to accept its homogeneous form for present purposes.

Given this decision, any biases involved in our estimates of parameters in equation (3) is on the distribution of the coefficient of wealth, v, and the coefficients of income and expected future income, y and ye. While y and ye are separately estimated as described above, in the estimation of equation (3), the sum of the two coefficients for these two variables rather than each separately proved to be quite stable under minor variations of specification, so let us take them together in the assessment of their biases.

The coefficients of the interaction term between v and $a(i)$ may be thought of as indicating the fraction of total resources (the market value of net worth plus the present value of current and future earnings) that the household wishes to consume during the current year. This proportion, of course, varies from one family to another, depending on many circumstances. The only restriction due to the homogeneity assumption discussed above is that it should not depend on the *absolute* level of total resources, although we do not exclude the possibility that it may depend on the relative position of the household in question in the distribution of total resources in a particular age cohort.

A careful review of the implications of a variety of environments faced by individual households indicates that, if the life cycle theory is to retain approximate validity, then the coefficient of net worth must increase with age on average, and its order of magnitude must reach the level of something like 0·10 by the time the head of the household is retired and reaches the age of 70 or so, even allowing for a fairly significant bequest motive.

It can be as small as 0·01 for families the age of whose heads is in their 20s.

Compared with this theoretical expectation, the actual estimates we have obtained for equation (3) using Japanese survey data for coefficients of the interaction terms between v and $a(i)$ seem somewhat too small. This necessarily implies that our estimates of the coefficients of y and ye are biased up to some extent.

This type of bias is likely to be based on one of two possible sources. First, it is possible that there is some simultaneous equations bias in our estimates of the coefficients of y and ye, although y and ye used in the estimation of equation (3) are predictions generated by equation (2) rather than their actual values, so that the possibility of simultaneous equations bias is relatively small. Second, it may be due to error of measurement of v. Such errors of measurement are likely to bias the coefficient of v towards zero, and given the homogeneity restriction, to bias coefficients of y and ye upwards. This second possibility cannot be ruled out because we know that the measurement of v is more seriously defective than the measurement of c and y (Hayashi *et al.* 1989).

In the analysis reported here, we accept the estimates of the parameters of equation (3) obtained through the instrumental variable regression procedure in spite of the potential biases discussed above. In another analysis currently under way, we work with alternative estimates obtained by assuming that the observed mean consumption of a cohort is, in fact, the desired fraction of the total resources of the cohort in question. This alternative estimate appears to make the coefficient of net worth somewhat larger and rising with age, thus conforming to the prediction of the life cycle model more closely than the one used here.

So far, we have been concerned with the consumption behaviour of active working age households 63 years old or younger. For those families with heads older than 63 years who are continuing to work, we can estimate an equation that is separate from but similar to equation (3). For those families who are remaining independent but whose heads are fully retired, we have much less information in our sample, and we assume that they will continue the behaviour exhibited during the current period, namely, they tend to consume almost all of their pension receipts and a small fraction of their net worth.

Those who are older than 63 and merged into households headed by a younger individual disappear from our system as independent households. Their income, consumption, and indirectly their saving and net worth, however, enter our system through their impact on the value of dummies in equation (2) for income and in equation (3) for their consumption. The description of the process determining the critical choice of which households merge into younger households and which younger households accept older individuals is given in the next section.

4. Dynamic model of demographic development

Equations (2) and (3) introduced in Section 3 above and similar equations for older groups generate predictions for the distribution of income and of the saving–income ratio for equation (1). In order to utilize equation (1) to decompose the aggregate saving–income ratio, then, we must have a model of demographic development that can generate the weights of each cohort. Demographic projections are quite common, and indeed, we can obtain a tape from the United Nations (1989) containing not only current demographic data but also models to generate the projection of the future development of population for all member nations.

These population projection models are, however, limited to the age and sex distribution of the population, and as far as we know, there does not exist an operational model of population dynamics which is capable of generating predictions about the distribution of family structure, such as the number and age of children in each family, the marital status of the head and the age of the spouse, presence or absence of other dependents and their sex and age. We need this additional information in order to utilize equation (1). Since this is our first attempt to model the dynamics of population, we have adopted the official model and projection for Japan provided by the Institute of Population Problems, Ministry of Health and Welfare (1992), as the shell for our more detailed demographic model. That is, we have adopted all of their assumptions and added some additional structures, and made sure that our projections match theirs to the extent that their projection exists, adding more details needed for our purpose.

Our starting point is the classification of all families in Japan in 1985 into cohorts, defined by age of the head, age of the spouse, and the number of children. Information required was taken from the 1985 census and the 1984 National Survey of Family Income and Expenditure. To simplify our task, in Japan, if both husband and wife are present in the family, we designate the husband as the head. We recognize individuals as capable of being a head from the age of 19 to 79 and include in the age class 80 all individuals aged 80 and above. There are thus 62 possible age classes for the head. We find very few families headed by persons aged less than 19, and we simply reclassified such families as headed by 19-year-olds. We also recognize a female spouse to be at most 5 years older or 10 years younger than her male spouse, and when we find exceptions, we have reclassified them to eliminate them. The number of children can be zero, one, two, three, and four or more, so there are five possibilities. Thus, the number of cohorts of families headed by married couples would be potentially $(62\times16\times5)=4960$. The number of cohorts for single parent families headed by a male or female is $(62\times2\times5)=620$, and the number of cohorts of male and female single individuals is $(62\times2)=124$. In practice, we found no member in some marginal cohorts, and the probability that someone will move into such cohorts in the future is virtually zero, so that the total number of cohorts in our analysis turned out to be a little less than 4000. To each cohort a weight is assigned, representing the population size taken from Population Census of 1985 except for some details estimated from the National Survey of Family Income and Expenditure of 1984.

For each cohort, we must maintain information on the age and sex distribution of all dependents. A dependent is considered a child if he/she is aged 18 or less, otherwise such a person is considered an adult dependent, and we recognize him/her to be from age zero to 79 and 80 or over. Thus, we carry this set of demographic information for each cohort as 162 variables.†

In addition to demographics, each cohort must carry what may be called "semi economic" information, such as the distribution of occupation and employment status among its heads, its spouses, and among its adult dependents. Finally, one piece of economic information that must be carried by each cohort is the initial value of net worth. In this initial attempt, we have carried only the mean value of net worth for each cohort, although we recognize that it would be very useful to carry at least the second moment assuming that net worth is distributed according to, for example, the log–normal distribution.

Once the cohort structure is fully constructed for the base year, we can specify the detailed procedure for updating this structure from one year to the next. For this purpose, we found that it is best to break up the transition process for a year into several substeps and treat them as though they occur sequentially. The substeps specified are the following:

Phase 1: Death, divorce, and remarriage

† The number of dependents in each age for a particular cohort is the *average* number of dependents of that age for each family in that cohort. Therefore, the number may be fractional, but when the recorded number of dependents for the cohorts are added from age zero to 18, this sum must be equal to the number of children defining the cohort.

Phase 2: Ageing
Phase 3: Birth of new children
Phase 4: New marriage, movement of dependent young adults to independent status, and retirement
Phase 5: Merger of older families and individuals with younger families.

We briefly comment on some of these processes.

The Institute of Population Problems (1992) provides detailed estimates of the death rate by age and sex for current and future periods, and we have simply adopted their estimates with one exception. We have reduced the birth rate and then set the death rates for those aged zero to 18 equal to zero, so that at age 19 our projection matches the official projection in all future periods. This simplification greatly reduces computational requirements in our simulation, and it does not seem to affect our result noticeably given the very low mortality of children in Japan.

In most cases, how weights among cohorts must be adjusted when someone dies is quite clear, except for one situation, namely, when a single head (without a spouse) of a household dies. We then must allocate dependents in this household somewhere. In this case, we have adopted an arbitrary rule by which we designate another adult in the household as the head if such an adult exists, and if not, we moved an adult whose characteristics are similar to the deceased person from a single person category to a single head of household category and assign children of the deceased to this person. The consequence of such an arbitrary rule appears to be negligible in any event because the death rate of single heads of family young enough to have child dependents is very low.

Divorce and remarriage are treated as a net process in this model (together with consequences of death of one spouse of a married couple), and it involves an obvious transfer of weights among cohorts. Probabilities for these events are inferred from information provided in the Final Report of the census.

The handling of the ageing process is reasonably obvious, but we wish to remind the reader that, when an 18-year-old child ages by one year, he is no longer a child, so that the family to which he belongs loses one child, and must move to a cohort with one less child.

As the result of the ageing process, all cohorts have a value of "zero" for the position of children with age zero. Thus, newborn children can be recorded readily in all cohorts. The main complication here is that the fertility rate used in population projections given by the official sources is conditional only on the age of the female, while what we need is the fertility conditional on a female of a particular age being married and having had zero, one, two, three, or four or more children. In the case of Japan, as part of the discussion of the methodology of the projection, the Institute of Population Problems (1992) provides three alternative limiting distributions of the number of children for married females associated with the three alternative fertility assumptions used in their projections. We have used these limiting distributions and calculated the implied fertility for married women of a particular age with a given number of previous births.

We then come to the description of first time marriages. What we need is the probability of marriage for an unmarried male of a specific age, and conditional on his marriage, the probability that he marries a woman of a specific age. We begin with the observed actual distribution of the marriage pattern of males in 1985 and the age distribution of their spouses, infer the probability of the male's marriage at each age conditional on their not yet being married, and modify the result in accordance with the discussion given by the Institute of Population Problems (1992), concerning gradual shifts of the marriage age of a female in Japan over time.

We now come to the last and a more complex demographic transition pattern that is specific to the Japanese case, namely, the retirement process and merger of older families and individuals with younger families. We have indicated how widespread the practice of the merger is in Table 2. For the retirement process, we have assumed that the probability of retirement at each age and occupation remains the same in the future as it was in 1985. Since participation in the labour force is an important factor in determining the saving–income ratio for a family, a more satisfactory explanation of the retirement process is a critical refinement that should be undertaken in our future work. For the merger process, we have adopted an earlier estimate of a probit equation describing this process as a function of age, marital status, sex, and the position in the wealth distribution in the appropriate age group of the individual in question, with some modification since we no longer have access to some of the variables used earlier.

In order to insure that our demographic model is generating patterns that are internally consistent, we have insured that the sum of the male/female population of various types (married heads of households, spouses of heads, single heads of households, independent single persons, and dependents in families headed by others) add up to the total population in each age, and that the total number of married males over all age groups is identical to the total number of married females over all age groups.

These transfers of families and individuals from one category to another inevitably involves a transfer of wealth along with persons. To describe the wealth transfer accurately is difficult because we have no information on parents or children living away from the family in question. We have adopted the following rules:

(1) If one of the spouses dies in a family in which both spouses are present, then one half of the family's net worth goes to the remaining spouse. (a) If there are one or more children living in the family (for this purpose, any young adult living in the family whose age is appropriate is considered a potential recipient of the estate), then the remaining half goes to those children living in the family and distributed among them equally. (b) If there are no children living in the family, then one half of the estate goes to presumed children in younger cohorts. Presumed children are defined by the potential fertility of the female spouse of the family, whether she is actually present or not.
(2) If a single person or single parent dies, the same process as in rule (1) above takes place except that the entire net worth is distributed among children or presumed children instead of one half of it.
(3) When a single person living independently marries, he/she is assumed to bring his/her entire net worth into the marriage.
(4) When a single person living as a dependent in the parent's family marries, he/she is presumed to be entitled to carry with him/her the same amount of net worth as the independently living single person of the same sex and age would have brought with him/her.
(5) When older persons merge into younger households, they bring their entire net worth with them and add it to the net worth of the host household. Since the identity of the older person is known only as a member of specific cohort in our simulation analysis, we do not know the exact net worth being carried by this person. We estimate the expected value of net worth involved assuming that the relative distribution of net worth among members of the cohort remains the same from the starting point, and taking account of the probit equation for determining the probability of the merger in which the relative position of net worth among the age cohort was an independent variable.
(6) The handling of inheritance and gift taxes is discussed in the next section.

5. Results

We begin by looking at changes in demographic structure in Japan between 1985 and 2050. As we have stressed earlier, our

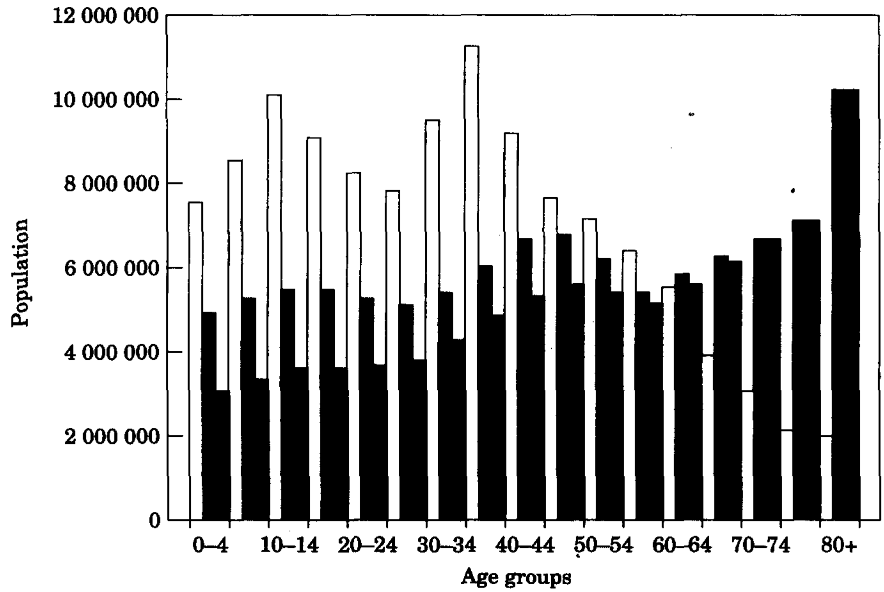

FIGURE A. Population by age, 1985 (□); 2050-middle fertility assumption (■); 2050-low fertility assumption (▩).

demographic model is designed so that the sex and age distribution of the population projection will conform to the one generated by the Institute of Population Problems, so that this aspect of our results is not new, and it is summarized in Figure A. The contrast between the 1985 pattern and the 2050 pattern is quite striking, and almost dwarfs the differences implied by the middle and low fertility assumptions. This is especially so because, in the year 2050, those aged 65 and above are identical under both assumptions because they had already been born before fertility assumptions deviated from each other in 1990.

Behind these simple figures are very different family structures and other patterns, and to illustrate the point, we present just two tables, one for actual patterns for 1985 and one expected to prevail in 2050 under the low fertility assumption, as Tables 4a and 4b.

In Table 1, we presented the actual savings pattern that prevailed in the years 1984–85. In Tables 5a and 5b, we show the corresponding savings patterns that the model predicts for 2020 and 2050 under the low fertility assumption. Perhaps the most surprising feature of these results is that, for the year 2020, the overall saving rate at 15% is actually higher than it was in 1985, then at 13%. This is in spite of the fact that the number of older individuals (70 years and over) as a fraction of the total population has increased from 6% in 1985 to 20%

Table 4 *Demographic structure by age and family type (in thousands)*
(a) 1985: Actual.

MALE

Age	*Total population*	*Married without children*	*Married with children*	*Single head of household*	*Single independent*	*Dependent*
0–19	17 941	0	0	0	370	17 573
20–29	8 114	605	982	76	3 104	3 349
30–39	10 108	783	6 864	145	757	1 558
40–49	551	794	7 181	53	230	293
50–59	6 646	4 063	2 160	64	193	166
60–69	4 191	2 620	696	54	115	706
70 & over	2 940	1 000	254	67	130	1 489
Total	58 491	9 865	18 137	459	4 899	25 135

FEMALE

Age	*Total population*	*Married without children*	*Married with children*	*Single head of household*	*Single independent*	*Dependent*
0–19	17 081	10	5	0	255	16 811
30–39	10 646	685	8 393	164	209	1 196
40–49	8 259	1 633	5 819	320	191	296
50–59	6 889	4 356	1 211	300	439	583
60–69	5 354	1 993	563	213	742	1 069
70 & over	4 320	376	99	77	500	3 268
Total	60 453	9 865	18 135	1 109	3 852	27 493

(b) 2050: Low fertility assumption

MALE

Age	*Total population*	*Married without children*	*Married with children*	*Single head of household*	*Single*	*Dependent*
0–19	7 014	0	0	0	19	6 995
20–29	3 871	387	285	9	919	2 271
30–39	4 620	820	2 223	42	651	884
40–49	5 583	1 037	3 531	35	415	565
50–59	5 275	2 660	1 782	53	327	453
60–69	5 762	4 035	270	46	348	1 013
70 & over	10 500	3 253	46	529	497	6 175
Total	42 625	12 192	8 137	764	3 176	18 356

Age	*Total population*	*Married without children*	*Married with children*	*Single head of household*	*Single*	*Dependent*
0–19	6 739	0	0	0	18	6 721
20–29	3 722	499	466	18	562	2 177
30–39	4 484	860	2 579	78	249	718
40–49	5 498	1 111	3 534	193	170	490
50–59	5 367	3 062	1 377	281	164	483
60–69	6 243	3 932	146	413	248	1 504
70 & over	15 813	2 719	35	1 526	667	10 866
Total	47 866	12 183	8 137	2 509	2 078	22 959

Table 5 *Projected saving–income and net worth–income ratio low fertility assumption*
(a) 2020

Category	*Weight*	*Income*†	*Sav/Inc*	*N.W./Inc*
All population	38 881 656	972	0·151	6·740
Male single head of household	714 747	817	0·077	13·778
Female single head of household	2 598 004	687	0·022	10·681
Male single	4 310 246	742	0·193	6·670
Female single	2 569 600	344	−0·199	8·351
Normal families	28 689 059	1 093	0·167	6·347
Normal families, by age				
<=29	952 630	668	−0·009	3·804
30–39	4 451 028	957	0·110	4·904
40–49	7 463 464	1 213	0·160	4·923
50–59	6 945 548	1 332	0·200	5·350
60–69	5 294 820	1 016	0·225	8·018
>=70	3 581 569	776	0·092	13·866

(b) 2050

Category	*Weight*	*Income*	*Sav/inc*	*Ass/inc*
All population	28 855 609	1 678	0·103	8·725
Male single head of household	763 888	1 389	0·028	17·828
Female single head of household	2 509 038	1 204	−0·016	15·057
Male single	3 174 866	1 232	0·079	8·736
Female single	2 098 282	603	−0·288	0·789
Normal families	20 329 536	1 926	0·129	7·923
Normal families, by age				
<=29	672 217	1 213	−0·047	4·785
30–39	3 042 416	1 743	0·070	6·084
40–49	4 567 603	2 202	0·100	6·392
50–59	4 442 428	2 393	0·149	6·835
60–69	4 305 773	1 835	0·220	8·943
>=70	3 299 099	1 349	0·092	14·937

† In ¥ 10 000

in 2020 under the low fertility assumption. The saving–income ratio for the over 70 years-of-age group in 2020 is 9%, considerably smaller than the average, so that the increase in the size of this group should have reduced the average saving–income ratio.

The most important counteracting factor here is the reduction in the number of children. The number of individuals 19 years old or younger as a proportion of total population declined from 29% in 1985 to 17% in 2020 under the low fertility assumption. The number of children 19 years or younger living with their parents has declined from 1·25 per couple in 1985 to 0·72 per couple in 2020 under the low fertility assumption. This decline in the number of children per couple clearly helped reduce the consumption–income ratio of middle aged families, ranging from the ages of 30 to 69, where most of the weights are concentrated. It is through the higher saving of this group, presumably due to the smaller number of children, that more than offsets the increase in the proportion of the oldest group of individuals and families, thus increasing the aggregate saving–income ratio.

By the year 2050, the overall saving–income ratio drops to 10%. The number of children as a fraction of the total population is no longer declining much (from 17% in 2020 to 15% in 2050), while the number of those aged 70 and above has increased dramatically, to 29% of the total. This increase in the number of older persons and families contributed to the decline in saving, but it is also clear, by comparing the pattern of the saving–income ratio for 2050 with that of 2020, that the saving–income ratio declined slightly for each group.

A careful inspection of equation (3) as it generates the saving–income ratio for each age cohort from the year 2020 to 2050 suggests that this decline in the ratio is the net result of three factors. First, the average number of children further declined slightly, and this has increased the saving–income ratio. Second, the average number of dependent adults per family, especially older adults, has increased, and this contributes to a reduction in the saving–income ratio. Third, the net worth–income ratio has increased for most cohorts, especially for older households, and this has contributed to a further decline in the saving–income ratio for most cohorts. The net effect of these three forces is to reduce the saving–income ratio for most cohorts.

The first two of these are natural consequences of the ageing population, and we believe that the feedback effects are relatively weak and indirect. The effects of income per family on fertility seem quite weak at this level of income, compared to other factors contributing to the fertility rate. The main suspect for potential bias is the influence of the income and wealth distribution on the decision of older individuals and families as to whether to remain independent or to merge with younger families. To the extent it is possible, however, we have endogenized this feedback channel and it is active in our simulation system.

The third one, on the other hand, may still be a remaining source of potential bias in our analysis. As we have noted earlier in our discussion of estimates of parameters for equation (3), it seems plausible that our estimate of the coefficient of net worth for most cohorts, especially for older ones, is underestimated. We have attributed this problem to probable greater errors of measurement for net worth. We have also noted that, given the homogeneity assumption we have imposed on equation (3), this underestimate of the coefficient of net worth probably means a counterbalancing overestimate of the coefficient of current and expected income in the equation.

We have not been able to resolve this issue satisfactorily at this time. The error of measurement of net worth is especially complex. We know that the net worth of households other than the value of the main residence is largely under-reported. We have made a major attempt to impute the market value of land for the main residence of each household, and we believe that we were reasonably successful in our attempt (Hayashi *et al.*, 1989), but here we encounter the most basic question of the saving behaviour by Japanese households. Given that the relative price of land in Japan is so much higher than in other countries, is it reasonable to judge coefficients of an equation like equation (3) estimated for Japan by comparing them with corresponding estimates for a country like the U.S., where the value of land is radically lower? For a Japanese household that owns a small house on a plot of land whose size is only one-thirtieth of an acre but whose market value is 10 to 20 times its annual earnings, should we expect it to save little because its net worth is high or should we consider its net worth to be not so high in some sense and therefore should we expect it to continue to save? We find it very difficult to resolve this question. We intend to investigate this question further by imposing somewhat stronger behavioural assumptions on equation (3) and then running alternative calculations.

In the meantime, there is another source of potential bias which we have investigated. According to the Annual Statistical Report of the National Tax Office of the Japanese Government, the average effective rate of inheritance and gift taxes on assets reported to be transferred from one generation to the next is roughly 18%. Assets *reported to be transferred* must be much smaller than assets *actually transferred,* partly because the value of land transferred, which is more than 67·5% of total assets *reported* to be transferred, is radically under-estimated for this purpose, and partly because a significant portion of net worth is exempt from inheritance and gift taxes. In most of our simulation analysis, we have ignored this question, but we have made one simulation in which we assumed that all intergenerational transfers were subjected to the transfer tax of 18% in order to see the maximum possible effect of this on net worth accumulation. The difference between these two alternative simulations was only 7·7 against 7·9 for the net worth–income ratio for all households in 2050, although the difference tended to be concentrated in younger cohorts so that the effects on them was a little larger. We believe, therefore, that the effect of ignoring this tax was not negligible, but it is not a major one.

Our analysis indicates, then, that the effect of the declining birth rate in Japan on the saving rate of that country over a fairly long period of time is that the saving rate will first rise somewhat and then decline, but not as much as is generally believed. This is partly a result of the positive effect of the reduced number of children per family on the saving–income ratio, and partly due to the fact that older individuals in Japan continue to maintain a high probability of participating in the labour force and earning income, thus dissaving less than expected. The high price of land may contribute to the older households' tendency to maintain a larger value of net worth–income ratio than their counterparts in other countries, but this may not contribute to the higher capital–output ratio in the productive sector of the economy, since land is not a reproducible asset.

There are two major reservations attached to this conclusion. First, when everything is said and done, we believe that the coefficient of net worth in equation (3) for older cohorts is somewhat underestimated, and to this extent, our estimate of the target net worth–income ratio and hence the saving–income ratio towards the end of our simulation is somewhat overestimated. Second, in this set of simulations, we have assumed that the pattern of productivity for individual households remained the same over time. It is probably more reasonable to assume that the rate of growth of productivity will decline from the 1980s to the 2020–2050 period. We believe that such a decline in the rate of growth of income would have reduced the saving rate of most households through habit persistence in consumption.

While it is clear already that the design of cohort simulation analysis which we have used in this paper can be improved substantially, we believe that the analysis of this type provides some important new insights into the complex aggregation process which connects the micro behaviour of individual households and the aggregate pattern of saving and asset accumulation.

The analysis presented here suggests that results can be quite surprising even in familiar, relatively simple cases.

References

Ando, A., Guiso, L. & Terlizzese, D. (1994*a*). Young households' saving and the life cycle of opportunities. Evidence from Japan and Italy. In A. Ando, L. Guiso & I. Visco, Eds. *Saving and the Accumulation of Wealth.* Cambridge: Cambridge University Press.

Ando, A., Guiso, L. & Terlizzese, D. (1994*b*). Dissaving by the elderly, transfer motives, and liquidity constraints. In A. Ando, L. Guiso & I. Visco, Eds. *Saving and the Accumulation of Wealth.* Cambridge: Cambridge University Press.

Ando, A., Yamashita, M. & Murayama, J. (1986). Analysis of consumption and saving based on National Survey of Family Income and Expenditure. *Economic Analysis,* **101,** Tokyo, Japan, Economic Research Institute, Economic Planning Agency (in Japanese).

Ando, A. & Moro, A. (1995). *Dynamics of demographic development and its impact on personal saving: case of Japan.* NBER Working Paper, Cambridge, MA.

Bosworth, B., Burtless, G. & Sabelhaus, J. (1991). The decline in saving: evidence from household surveys. *Brookings Papers on Economic Activity,* **1,** 183–256.

Cannari, L. (1994). Do demographic changes explain the decline in the saving rate of Italian households? In A. Ando, Eds. *Saving and the Accumulation of Wealth.* Cambridge: Cambridge University Press.

Gokhale, J., Kotlikoff, L. & Sabelhaus, J. (1994). *Understanding the post-war decline in United States saving: a cohort analysis.* Mimeo. Department of Economics, Boston University, Boston, MA.

Hayashi, F., Ando, A. & Ferris, R. (1989). Life cycle and bequest savings. In *Saving Behavior, Investment, and Rate of Return on Capital in the U.S. and Japan: Comparative Analysis and Perspectives into 1990s.* Tokyo: National Institute for Research Advancement.

Institute of Population Problems, Ministry of Health and Welfare (1992). *Population Projections for Japan: 1991–2090,* Research Series No. 271. Tokyo, Japan, Ministry of Health and Welfare (in Japanese).

Kirman, A. (1992). Whom or what does the representative individual represent? *The Journal of Economic Perspectives,* **b,** 117–136.

Modigliani, F. & Brumberg, R. (1980). Utility analysis and aggregate consumption functions: an attempt at integration. In A. Abel, Ed. *The Collected Papers of Franco Modigliani, Vol. II.* Cambridge, MA: MIT Press.

National Tax Office, Government of Japan (1986). *Annual Statistical Report,* 1984. Tokyo, Japan.

Statistics Bureau, Management and Coordination Agency, Government of Japan (1987). *1984 National Survey of Family Income and Expenditure.* Tokyo, Japan.

Statistics Bureau, Management and Coordination Agency, Government of Japan (1990). *Final Report of the 1985 Population Census: Population of Japan (Statistical Tables).* Tokyo, Japan.

United Nations (1989). *Global Estimates and Projections of Population by Sex and Age,* the 1988 Revision, New York.

Appendix 1: Family types and saving rates U.S.A.†

PART 1

	1960 (%)	1972–73 (%)	1984–85 (%)	1986–87 (%)	1988–90 (%)
Saving rate‡					
All families			3·47	4·20	5·86
Single parent			−12·56	−12·36	−9·67
Single			−3·12	−7·48	−6·10
Nuclear			6·76	9·15	10·55
Extended			5·64	7·57	8·52

PART 2

	1960 (%)	1972–73 (%)	1984–85 (%)	1986–87 (%)	1988–90 (%)
Household distribution					
All families	100·00	100·00	100·00	100·00	100·00
Single parent	3·34	5·93	11·59	12·19	12·20
Single	15·05	22·11	27·25	28·86	28·57
Nuclear	81·61	60·33	46·54	45·53	45·18
Extended		11·63	14·62	13·42	14·05

PART 3

	1960 (%)	1972–73 (%)	1984–85 (%)	1986–87 (%)	1988–90 (%)
Relative disposable income					
All families			100·0	100·0	100·0
Single parent			67·5	62·9	65·8
Single			57·7	54·2	57·5
Nuclear			121·2	119·3	124·6
Extended			137·0	132·7	137·1

† Prepared from the Public Use Tapes of Survey of Consumer Expenditure, Bureau of Labour Statistics.

‡ Income and Expenditure for these calculations are defined to make them as close as possible to the definition used in National Income and Product Accounts. The resulting estimates, however, still contain significant conceptual differences from the NIPA accounts; the most important difference is that saving here does not include employer contributions to private pension funds. Top-coded entries are adjusted by our estimates of their actual values.

PART 4 *Estimate of Aggregate Saving Rate for 1980s Using 1960 and 1972–3 Weights of Family Types*

	(1) Actual (%)	(2) 1960 Weight (%)	(3) 1972–73 Weights (%)	(4) (2)/(1)	(5) (3)/(1)
1984–85	3·47	6·07	4·94	1·75	1·42
1986–87	4·20	7·91	6·40	1·88	1·52
1988–90	5·86	9·82	8·18	1·68	1·40

Note: Part 1 of the table gives the saving–income ratio for various family types in the 1980s, and Part 3 gives the relative size of income for the same groups. Part 2 provides the relative size of these groups for the 1980s as well for 1960 and 1972–3, the earlier years for which the same survey results are available. We can then ask the question: assuming that the saving–income rate for these groups and the relative size of income were the same for earlier years as for the 1980s, does the shift of weights among these groups explain a significant portion of the change in the aggregate saving–income ratio from 1960 to the 1980s? To answer this question, we recompute the aggregate saving–income ratio taking the group ratio and relative income positions in the 1980s as given but using relative weights for 1960 and 1972–73. The results are given in Part 4. We find that a significant part of the decline in the aggregate saving–income ratio is indeed explained by shifts in weights among these groups.