

Lindh, Thomas

Working Paper

Medium-Term Forecasts of Potential GDP and Inflation Using Age Structure Information

Sveriges Riksbank Working Paper Series, No. 99

Provided in Cooperation with:

Central Bank of Sweden, Stockholm

Suggested Citation: Lindh, Thomas (1999) : Medium-Term Forecasts of Potential GDP and Inflation Using Age Structure Information, Sveriges Riksbank Working Paper Series, No. 99, Sveriges Riksbank, Stockholm

This Version is available at:

<http://hdl.handle.net/10419/82391>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Medium-Term Forecasts of Potential GDP and Inflation Using Age Structure Information^α

Thomas Lindh^γ

Department of Economics, Uppsala University
Institute for Futures Studies, Stockholm
Submitted as working paper at Sveriges Riksbank

November 28, 1999

Abstract

Economic behavior as well as economic resources of individuals vary with age. Swedish time series show that the age structure of the population contains information correlated to medium-term trends in growth, inflation and other macroeconomic data. GDP gaps estimated by age structure regressions are closely related to conventional measures using filters, time series methods, and structural models. Reliable estimates of potential GDP and other macroeconomic series often become available after (at least) a one-year lag. Monetary policy is believed to affect inflation with a lag of one or two years. Projections of the population's age structure are comparatively reliable several years ahead and provide additional information to improve on 3-5 years-ahead forecasts of potential GDP and inflation.

JEL codes E31, E44, J10, O52

Keywords: Age structure, inflation forecasts, potential GDP

^αI am very grateful to Anders Vredin, Per Jansson, Tomas Lindström, Bo Malmberg and seminar participants at Sveriges Riksbank for helpful comments. I also owe Tomas, Per and Bo as well as Anders Forslund for assistance with data. Research funding by Sveriges Riksbank is gratefully acknowledged. The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank.

^γCorrespondence: Department of Economics, Uppsala University, Box 513, SE-751 20 Uppsala, Sweden, tel +46-18-471 11 03, fax +46-471 14 78, mob +46-70-688 48 79, email thomas.lindh@nek.uu.se, URL <http://www.nek.uu.se/faculty/lindh>

Contents

1	Introduction	1
2	Introductory remarks and caveats	3
2.1	Three points to keep in mind	3
2.2	Age distribution data and methodological issues	6
3	Exploration of age effects on the economy	8
3.1	The Post-War Period	9
3.1.1	The complex economics behind age structure effects	13
3.1.2	What have we learned?	15
3.2	Specification issues	16
3.2.1	Inflation	16
3.2.2	GDP growth	17
4	Medium-Term Forecast Performance	20
4.1	Forecasting the 1990s out-of-sample	22
4.2	Phillips-Okun relations	29
4.3	Two further sensitivity tests	35
5	Conclusions	36
	Data appendix	42
	GDP	42
	CPI	42
	Population	42
	Current Account	43
	Gross investment	43
	Unemployment	43

1 Introduction

A stable correlation between inflation and the age distribution of the population has been uncovered in empirical data for the OECD countries (Lindh and Malmberg, 1998; Lindh and Malmberg, 1999c), for the United States (McMillan and Baesel, 1990), and Australia (Lenehan, 1996). In the latter two papers GDP growth is also shown to be correlated with the age structure.¹ These results suggest that age structure may be useful for forecasting inflation and potential GDP in the medium and long run.

This paper studies post-war Swedish time series data to assess the potential for such forecasts. The results corroborate the findings in other data sets. With some exception dependent age groups—like children and pensioners—have a negative relation to GDP growth and a positive relation to inflation, while working age groups show the reverse correlation. If these correlations remain stable, the next decade in Sweden will be characterized by high growth and low inflation. Around 2010 the large birth cohorts from the 1940s start to retire and the situation will reverse. There is, however, some evidence that the inflation targeting regime in Sweden in the 1990s has at least dampened a deflation impulse from the age structure. Even so, the underlying deflationary pressure from the age structure is still there and may have implications for the design of monetary policy. Qualitative statements on the direction of demographic effects on inflation and growth trends can be made with a high degree of reliability. Precise quantitative forecasts are, of course, more uncertain. Collinearity plagues coefficient estimates. Per se that should not pose any great forecasting problems provided the correlations within the age distribution were stationary. Unfortunately this is not the case. However, as will be demonstrated below forecasts 3-5 years ahead have forecast errors similar to those of conventional models at shorter horizons. Moreover, age structure information seems more or less statistically equivalent to other information commonly used in inflation and potential GDP modelling.

Why would inflation and growth be correlated with the age structure? On a general level—for the moment disregarding policy interventions—a simple life cycle saving and human capital accumulation story (Modigliani and Brumberg, 1954; and Becker, 1962) could be told. When a large proportion of the population is saving for consumption during retirement, demand tends to be lower than supply, generating deflationary pressures. A large proportion

¹Higher growth with higher population shares in working ages is also found in the OECD by Lindh and Malmberg (1999b), for Asian countries by Bloom and Williamson (1998), and in world-wide samples (e.g. Bloom and Sachs 1998; and Bloom et al. 1999). In addition such relations are found in Swedish (Malmberg, 1994) and Scandinavian time series (Andersson, 1998), and in regional data in the United States (Persson, 1998).

of children or retirees implies a tendency for excess demand and inflationary pressures. An older and more experienced work force, and a faster rate of capital accumulation in a net saving population should also stimulate economic growth. The age distribution impact on inflation can thus be phrased both as a demand pull and cost push story. Cohort-size effects can affect inflation both directly through excess demand effects and indirectly through changes in factor supplies which change the productive capacity of the economy. The correlations seen in data are mostly consistent with any of these mechanisms.

But there may well be other mechanisms at work behind the correlations in the empirical record. For example, the demand for low-productivity-growth services (at least as measured in GDP) from children and pensioners may contribute to lower measured growth. Large population shares in these age brackets will also affect government budgets. Tax revenue falls with fewer active income earners, and public expenditures rise with increasing education and health needs in the population. Structural adjustments in the production of goods and services will also take place when the content of the average consumption basket shifts as different segments of the age distribution dominate demand and supply relations. No attempt is made here to differentiate between these and other possible mechanisms. The age coefficients estimated in this paper are reduced form parameters that cannot be given any rigorous structural interpretation.

To what extent we will observe age structure pressures in macroeconomic data will, of course, depend on institutional factors and economic policy. A central bank could forecast age pressures and adjust interest rates and money supply to erase any correlations from the actual inflation record. Of course, we would then not conclude that age structure is neutral with respect to inflation. On the contrary, such policies would unequivocally prove its importance. It is, however, a fact that central banks have so far not engineered policy to fully counteract all kinds of disturbances, including age structure changes. The strong correlations between age structure and inflation found here provide corroborative evidence for this.

It is an empirical question whether age structure contains information that is useful for forecasts of medium-term inflation and potential GDP growth. Age structure offers a forecasting tool with several advantages. Although demographic projections are uncertain like all forecasts, the age distribution as a whole is only marginally affected by these uncertainties in the medium run.

But there is no such thing as a free lunch. While the variation in the age distribution is quite considerable it is comparatively slow and changes in different parts of the distribution are often highly correlated for natural

reasons. The pattern of effects is therefore hard to identify. Consequently reliable quantitative forecasts are not easy to make either. The evidence presented here indicates, however, that out-of-sample forecasts perform rather well on a medium-run horizon, and the GDP gaps implied by age predicted growth are remarkably similar to those found by completely different methods.

2 Introductory remarks and caveats

The outline of this paper is as follows. The empirical study starts off in Section 3 with descriptive regressions of macroeconomic variables on age structure, and a discussion of potential explanations for the patterns found in the Swedish data. Then a specification search is undertaken to find good forecasting models for inflation and growth: including control variables, taking care of outlier observations, checking statistical properties of the residuals, checking for possible endogeneity problems, parameter stability and so on. In Section 4 the results from Section 3 are put to use by estimating a system of inflation and growth equations and using this for out-of-sample forecasting and testing the performance. Finally the connections to conventional Phillips-Okun models of inflation is explored by computing the GDP gaps implied by age predicted growth and using this to model inflation.

Below some important points regarding the implementation and interpretation of the empirical results are briefly discussed.

2.1 Three points to keep in mind

Inflation forecasts conditional on monetary policy are, of course, more policy relevant than forecasts made independent of policy. Unfortunately it is too hard in this context to condition the estimates on Swedish monetary policy. Over the post-war period monetary policy in Sweden has changed objectives several times, and for most of the period it has been tightly interwoven with other economic policy objectives as well. Only from 1992 has there been a fixed inflation target of 2 percent \pm 1 percentage point. Forecasts conditional on monetary policy reactions would need to include policy instruments and their effects in the information set. To do so for the entire post-war period would require both more data and more precise knowledge of the economic mechanisms than is available. The analysis here therefore proceeds on the simplifying assumption that a stable pattern can be found in spite of variations in monetary policy. That might be a partially unwarranted assumption with respect to the 1990s as will be discussed further down.

Two other general points are also useful to bear in mind. First, the slow variation in age structure means it cannot be expected to reflect short-term variation. Consequently, forecasts do not perform well in the short run. However, most methods to forecast inflation and output gaps are in fact making some decomposition into what may be called a trend part—which might be more or less varying over time—and some short-run part. The trend part, whether stochastic or deterministic, may be motivated and estimated in several different ways, varying from simple Hodrick-Prescott filters to more elaborate structural models, see e.g. Apel and Jansson (1999a). Age structure based forecasts could be viewed as an alternative way to predict the trend rather than the short-term variation. Below it is demonstrated that with respect to the output gap this works rather well.

Second, age structure is not a single number but the distribution $A_t(a)$ of age groups in the population. Interpreted as a continuous distribution over the age span this means that macroeconomic variables are hypothesized to be functions of functions, functionals, and not ordinary functions. For nice, well-behaved distributions that can be characterized by just a few parameters this would not be much of a problem. Age distributions are, however, not very well-behaved in general. The demographic transition—experienced by all industrialized countries—from a high-mortality and high-fertility equilibrium to a low-mortality and low-fertility equilibrium, has been accompanied by recurring baby booms and baby busts. The reasons for this are not entirely clear, but it is an empirical fact with far-reaching consequences.

Figure 1 shows the variation in the Swedish post-war age distribution for some aggregated age groups and the projected variation up to 2030. While the share of children in the Swedish population has decreased from a maximum close to 24 percent in the 1950s to around 18 percent today the 75+ group has increased its share from a mere 3 percent in 1946 to around 9 percent today. This aging trend is projected to continue more or less in the future. Young retirees increased from slightly more than 6 percent in 1940 up to nearly 10 percent in the end of the 1980s, when Sweden had the largest share of people aged 65 and above in the world. During the 1990s this group has become populated by the sparse birth cohorts from the 1930s and has therefore decreased its share by nearly a fifth. The middle aged group has fluctuated between a maximum in the 1960s of nearly 19 percent to a minimum around 1990 slightly more than 15 percent, even lower than in 1946. The large birth cohorts from the 1940s started flowing into this age group after 1990 so its share has been increasing sharply and is projected to reach a new maximum of 20 percent in about 5 years from now.

The booms and busts in fertility have resulted in an age distribution that is far from stable, with a strongly multi-modal form which shifts as time

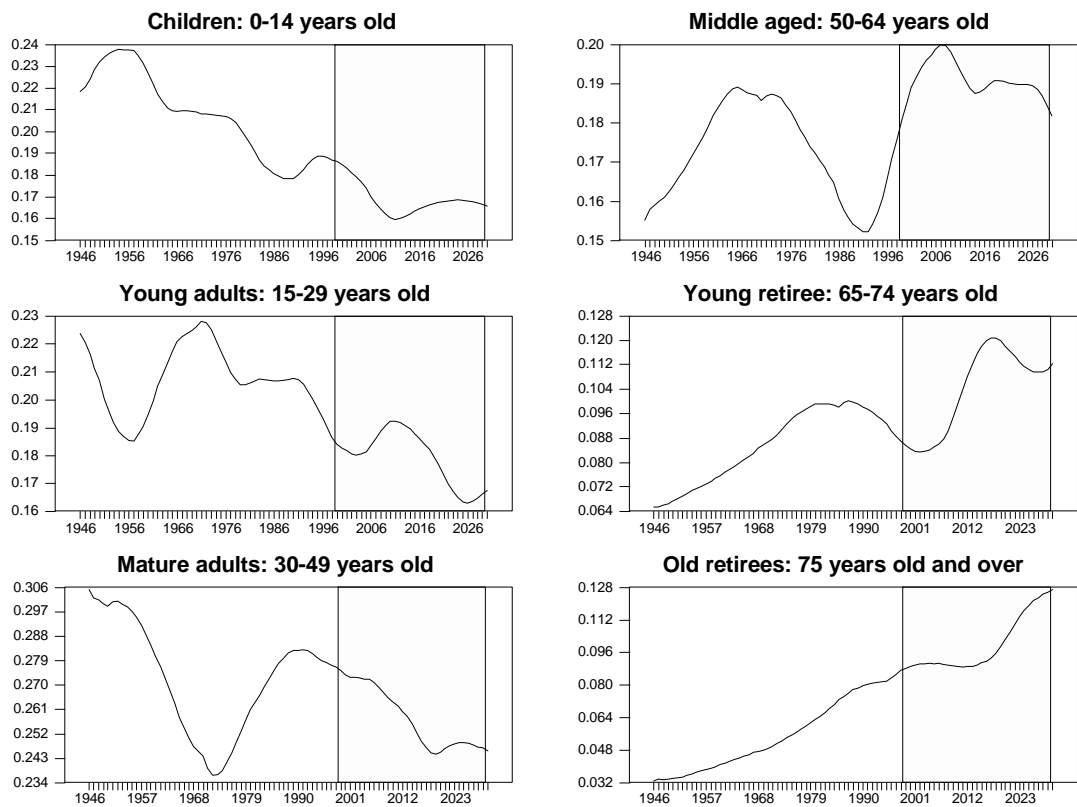


Figure 1: Swedish population age shares during the post-war period 1946-1998. The shaded area contains projections 1999-2030.

goes by. In Figure 1 we can follow the dents made by the very different fertility rates in the 1930s and the 1940s. These dents have two important consequences. On one hand, this variation is the basic reason why we can distinguish any age effects at all in the empirical record. A stable population where age shares remained constant would preclude identification of any age effects. Nor would age effects be of any importance for the aggregate economy, since relative cohort sizes would not change. On the other hand, the instability of the distribution means that the effects are hard to reduce to changes in a few aggregate statistics, even if it may be possible sometimes. The next subsection discusses this in some more detail.

2.2 Age distribution data and methodological issues

The age distribution data used here are annual estimates of the number of people in five-year cohorts taken from official Swedish statistics. Forecasts of the age distribution of the population are from the Statistics Sweden 1997 projection (latest update in 1999).

Often, age distribution effects are represented by a single aggregate measure: for example, old age dependency rate (population share of people aged 65 and over), youth dependency rate (share of people aged 0-14) or total dependency rate (old age plus youth dependency rate). A weakness with this approach is that only a small part of the total variation in the age distribution is used in the estimation of age effects. In particular, important variations in economic behavior and resources that occur during the course of normal working life may be ignored.

Ideally, one would like to take account of the variation in every age group since economically important changes in behavior and resources for different age groups occur in a gradual fashion. Due to the collinearity across age groups and perhaps also because important phases in life—like retirement, age of marriage and so on—tend to shift a bit over time, it is not possible to estimate separate coefficients for each one-year age group. The general form of the average life cycle patterns is, however, fairly stable, so aggregation should be helpful in smoothing out minor variations.

One econometric approach to circumvent this, pioneered by Fair and Dominguez (1991), is to use a polynomial restriction. The coefficients of the age groups are constrained to a polynomial defined on the mean age of each age group. Then only two or three parameters need to be estimated. The regression variables used are age weighted sums of population age shares with different powers of age as weights, i.e. different moments of the age distribution. But in many cases this restriction is rejected by the data, since it is unable to accommodate the comparatively abrupt changes in economic

effects that take place around retirement, see e.g. Lindh and Malmberg (1999a).

A third way to represent the age distribution—which will be used here—is to include population shares for a relatively small set of aggregated age groups that captures the most important phases of an individual’s economic life cycle. In comparison to the single-measure approach this allows a fuller representation of the age structure, and at the same time is less restrictive on the form of the age effects than the polynomial approach. A subdivision into six age groups has been used here: children 0-14 years old, young adults 15-29 years old, mature adults 30-49 years old, middle aged 50-64 years old, young retirees 65-74 years old, and old retirees above 75 years of age.

This general division can be motivated on theoretical grounds even if the exact delimitation is a bit arbitrary. Children, ...rst, do not take economic decisions themselves. Young adults often live single or are still living with parents. In OECD countries they are also to a high extent still studying and have quite distinct consumption habits. Mature adults are raising families, buying homes and starting in earnest to accumulate wealth. The middle aged are generally past their family years, have high incomes and are more immediately concerned with their retirement prospects. Young retirees are still rather active but have started to dissave, at least in terms of their pension claims. The old retirees have considerably more health problems, higher mortality and are more concerned with bequests.

Let z_t be a variable to be explained by the age distribution. If we make the simplifying assumption that we only consider linear effects the aggregated approach amounts to slicing up the distribution into larger chunks and assuming that each chunk has a constant effect. The dependent variable in the regression equation is thus assumed to be (ignoring other independent variables)

$$z_t = \int_{A_t(0)}^{A_t(125)} b(a) dA_t(a) + \sum_{i=0}^{14} b_{i,k} A_t(i)$$

where we approximate the $b(a)$ coefficients over a range of ages by a single number. The polynomial approach in contrast assumes that

$$b(a) = \sum_{n=0}^k b_n a^n$$

and estimates the parameters b_n using moments of the distribution.

The aggregated age groups are cross-correlated, hence the estimated coefficients will not be independent, but may be blown up or shift signs due to

collinearity. The correlation pattern between the age groups will shift with sample periods inducing some further instability into coefficients. For example young and old retirees are trended in the same direction up to around 1990, when the correlation shifts sign. Shifts in the internal distribution of each age group will also induce variations in the appropriate coefficient. However, for predictive purposes the group of variables may still perform well if the instability in the coefficients mainly is due to multi-collinearity.

Since the sum of age shares is unity either the intercept or one of the age groups must be dropped from a regression equation to avoid perfect collinearity. To make interpretation more intuitive I have chosen to drop the intercept, but it should be noted that the actual sign of each coefficient is still dependent on the intercept. Shifts in the intercept will, therefore, induce shifts in the coefficients even if the relative pattern is quite stable. It is the relative pattern of age coefficients that is important rather than the absolute level of each coefficient.

In view of these qualifications the age structure regressions below turn out to be remarkably stable. Still, there is some fluctuation in coefficient values and collinearity is a problem. But that is only to be expected and should not discourage the use of one of the very few variables that can be independently and fairly reliably projected 3-5 years ahead.

3 Exploration of age effects on the economy

Besides age data, annual macroeconomic data on real GDP, gross investment, current account balances, consumer price index, and unemployment have been linked from a variety of sources (see the data appendix) to yield as long series as possible of post-war data 1946-1998. In order not to lose data because of lags and the computation of growth rates, the data actually used stretches back into the war years. Although data for the war years may be of inferior quality—due both to the peculiarities of war economies with rationing, restricted trade, and due to the large war related public expenditures—it is likely to be better than no data at all. It is important to have long series in order to achieve robust estimates of age effects. This is not primarily a matter of degrees of freedom—if it were I could have used quarterly data—but to increase the variation of the age groups by extending the series over as many birth cohorts as possible. Even with a series stretching back to 1946 no cohort is observed during a full life time, so only a part of the potential variation is actually observed.

So why not go even further back? Because it is essential for a forecasting model that the information incorporated into the model is relevant for the

forecasting period. The further back we go, the less likely this is. A balance must be struck between the length of the series and how well this time period reflects currently relevant information. To this should be added the concern that data quality tends to deteriorate when we go further back in time. Preliminary investigations of macroeconomic series 1860-1998 showed that an appropriate cut-off seems to be after World War II. Both World Wars represent major structural breaks in the series, and there was evidence that the correlations to age structure were different before World War II. For forecasting purposes the information gain from including older data is therefore likely to be very small.²

The patterns of age effects on CPI inflation, GDP growth, the gross investment share of GDP and the national saving rate are explored in this section. GDP growth and inflation is computed as the annual log difference of real GDP and CPI, respectively. The gross investment share is the ratio between gross investment and GDP in current prices. The national saving rate is computed by adding the current account balance divided by GDP to the investment share. This is not entirely correct. Net factor income from abroad should also have been added. This is probably of minor importance and there were difficulties in obtaining comparable data for the whole sample.

The age data in the sources refer to the population 31 December of the given year. In order to ensure that the age shares are predetermined relative to the macroeconomic variables, the age data used here have been lagged one year, so they refer to the age distribution in the beginning of the year.

3.1 The Post-War Period

As a preliminary investigation inflation, growth, the investment share and the saving rate were regressed on the six age shares recursively backwards from 1998. The estimated coefficients start to stabilize as the sample gets back to around 1970, i.e. around 30 years of data in the regressions. However, around 1951 the coefficient estimates start to fluctuate again. In 1951 Sweden experienced a strong inflation shock due to the Korea war. Adding a time dummy for 1951, inflation and growth coefficients become stable back to 1946.

The coefficients in the saving rate and investment share equations still became unstable—in spite of the 1951 dummy—as the years before 1951 were included. Data are linked over 1950 and there is reason to believe that

²For the study of Sweden's economic history, these results, however, indicate a new and fruitful approach to explain phenomena like the huge current account deficits before World War I, and Sweden's rapid convergence rate after growth took off in the end of the 19th century.

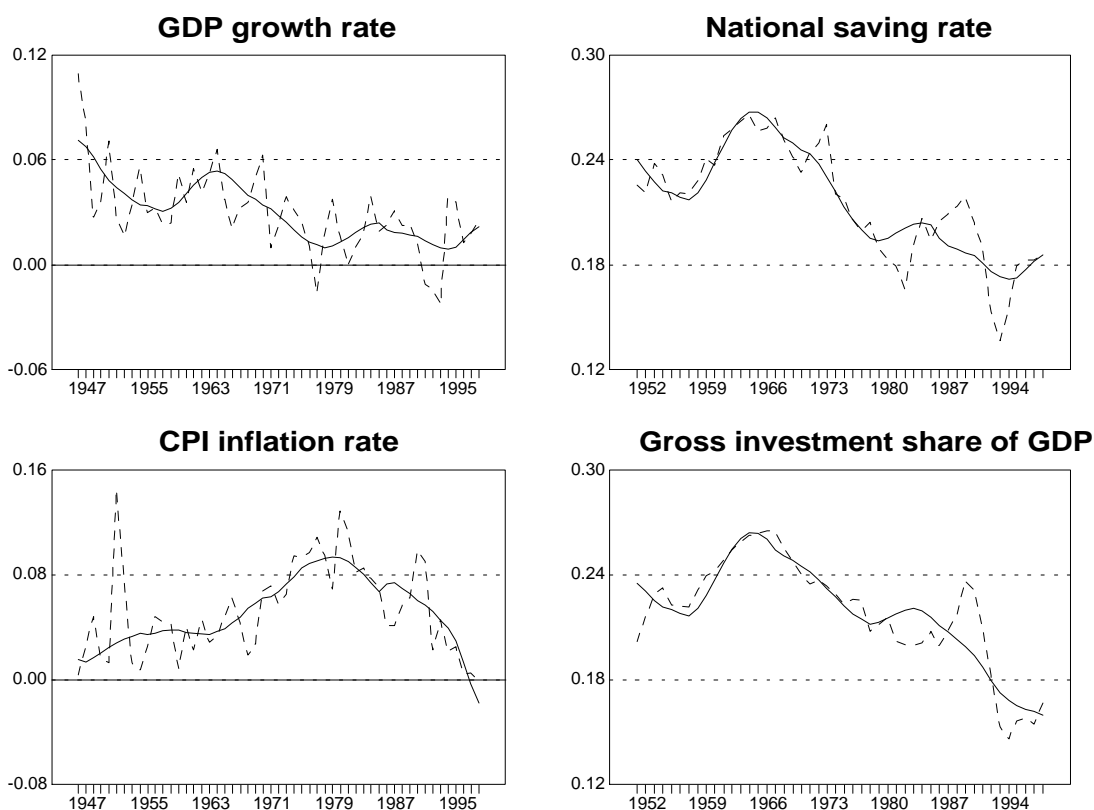


Figure 2: Age structure models of macroeconomic time series compared to actual data. See Table 1 for the estimated equations.

the quality of investment and the current account data may deteriorate, be differently defined or be affected by a number of other special circumstances related to the war, which are less important for the growth rates in CPI and GDP. The samples for the savings and investment regressions were, therefore, truncated to 1951-1998.

In Figure 2 the age predicted model is compared to actual data. The medium-run trends are picked up reasonably well in the growth and inflation equations but less well in the latter part of the time series for gross investment and the saving rate. In particular the age structure does not seem to be correlated with the overheating of the Swedish economy in the late 1980s. But age structure catches the trend shifts in the saving and investment rates that takes place in the 1970s rather well.

In Table 1 the regression equations corresponding to Figure 2 are reported. Inflation is positively associated with children and especially young

retirees while the active groups are deflationary. The strong negative effect from the oldest is the only surprising feature, since the saving rate regression indicates that they do have a negative effect on saving. On the other hand factor supplies should increase with the active age groups. Individuals passing from active ages to retirement thus affect the factor supply negatively, decreasing the supply of both labor and capital. The latter should theoretically not be very important in an open economy.³ An individual passing from young retiree to old retiree does not affect labor supply, however. Moreover, mortality rates are much higher above 75, so large shares of the population in that age group means an increased flow of free slots in the housing stock, thus depressing residential investment (1999d). There is also microeconomic evidence that they consume less than expected from life-cycle theory.

In the growth equation the pattern is rather the opposite of the inflation pattern. Dependents put a drag on growth and the active groups, in particular the 50-64 group, pull up growth. Saving and investment have a similar hump shape although young retirees have a positive effect on investment in contrast to their strong negative effect on saving. These patterns correspond very well to previous findings on OECD material. With respect to life cycle theory the strong positive savings effect from young adults is not expected. Credit rationing and household formation could be two potential explanations. However, it may also be an artifact due to collinearity.

Saving rates and investment shares are trended and there is significant first-order correlation in the residuals, so although the patterns are not too unreasonable from a theoretical point of view, a more thorough investigation is needed to draw any conclusions. Note, however, that standard errors have been adjusted by the procedure suggested by Newey and West (1987). Unless the serial correlation indicates misspecification the test statistics in Table 1 are still valid. The space here does not permit pursuing this further, and we will drop these variables except for use as controls in the following.⁴ The difference in effects of young retirees on saving (negative) and investment (positive) is, however, consistent with the different inflation effects of young and old retirees. Dragging down saving and boosting investment implies a strong positive effect on aggregate demand from young retirees. While old retirees also have a negative effect on saving, their effect on investment is

³However, the positive investment effect may be caused by the fact that this age group is still converting real assets to financial assets. There is strong evidence of this in Swedish microeconomic data, see Andersson (1999). The conversion of real to financial assets stops around age 75.

⁴Discussion and empirical results concerning the age dependence of these variables, and by implication the current account, can be found in Higgins (1998), Higgins and Williamson (1997), Herbertson and Zoega (1999), and Lindh and Malmberg (1999a).

Dep var	Inflation	GDP Growth	Saving rate	Investment rate
Age 0-14	0.318 (0.94)	-1.118 (3.61)	-1.483 (3.18)	-1.725 (4.01)
Age 15-29	-0.210 (1.18)	0.384 (2.37)	1.084 (3.59)	0.797 (2.14)
Age 30-49	-0.283 (1.40)	0.623 (4.06)	0.854 (3.12)	1.015 (4.16)
Age 50-64	-0.924 (3.42)	0.915 (3.16)	1.639 (4.42)	1.423 (3.45)
Age 65-74	4.476 (8.87)	-1.146 (3.86)	-1.441 (2.36)	0.624 (1.049)
Age 75 above	-2.148 (5.22)	-0.850 (2.57)	-1.580 (2.68)	-2.884 (5.09)
Dummy 1951	0.118 (27.6)	-0.016 (4.46)		
\hat{A}^2 age var (p-value)	669 (0.000)	749 (0.000)	32733 (0.000)	38131 (0.000)
adj R ²	0.667	0.417	0.791	0.791
# obs	53	53	48	48
D-W	1.70	1.64	0.83	0.50

Table 1: Post-war sample 1946-1998 for inflation and growth, 1951-1998 for the investment share and the saving rate. OLS regression estimates with a pure age share model. Absolute t-values are in parentheses. Standard errors are Newey-West corrected for heteroskedasticity and autocorrelation.

even more negative, so the net effect is that the aging of the population tends to create a current account deficit as long as both groups are growing.

3.1.1 The complex economics behind age structure effects

Although fairly simple life cycle models seem sufficient to explain the general form of the age effects, it is harder to reconcile the large quantitative effects with microeconomic evidence on life cycle behavior. In particular demographic savings effects have been discounted in the literature on that ground (see e.g. Bosworth et al., 1991). Apart from the econometric uncertainties recounted in Section 2, there are also more fundamental reasons why the quantitative effects are hard to explain directly from e.g. household saving behavior. Changes in the age structure affects almost all economic equilibria, as can be illustrated in stylized form in overlapping generations (OLG) models. The large changes in the population's age composition also have political and social consequences that feed back into the economic relations in ways that may both reinforce or cancel out in the aggregate.

Space does not permit any deeper analysis of this interaction here, but at this point it may be useful to make a slight narrative digression. Although the main objective of this paper is to explore the forecasting potential of age structure, an intuitive picture of how age structure interact with the economy at the macro level is useful to avoid an oversimplified understanding of the data patterns. A way to acquire that intuition is to consider a somewhat speculative story about the economics behind the correlations in Table 1 and the macroeconomic age effects in Figure 2.

In Figure 2 we can start to discern an economic story about what was driving the inflationary trend in the 1970s according to these age models. National saving is decreasing much faster than investment creating a current account deficit. This is largely a consequence of the difference between the two oldest age groups in their effects on investment. The reason why this dominates the macroeconomic evolution is that both these groups are increasing at approximately the same rate at this time. The rapid decrease in the middle aged group, however, does not matter much to the current account since it has approximately the same effect on saving and investment. This is also true for the other age groups.

Hence we have the combination of a growth slowdown and a deterioration of the current account. In the absence of a tight monetary policy during this period, it is not surprising that this results in rising inflation rates, especially since government consumption at the same time expands rapidly. This increase in government consumption—which is an important part of the downturn in the saving rate—is in turn at least partly driven by increasing

demand for health services from the old and a rapid increase in the share of mature adults.

Mature adults at that time exhibit increasing labor participation rates and are more highly educated than both the preceding and the succeeding cohorts (Ohlsson, 1986). An increasing demand for child care services and for more voluminous housing is not unexpected as these groups increase. These age groups were also increasing their relative weight in the electorate and thus had considerable political clout, making it hard for politicians to resist demands for heavy subsidies for housing and child care. A similar story holds for retirees and their demand for health care and pension benefits, although these demands started to increase earlier in the 1960s. See Lindert (1994) about political economy effects from ageing.

The dominating paradigm in the economics profession was at that time Keynesian, making it easy to motivate expansionary policies when growth went down and unemployment went up. Growth falls with the retirement of the middle aged baby boomers born just before World War I. Unemployment rises with the entry of baby boomers born in the 1940s into the labor market, since young adults with higher mobility have higher unemployment rates (Shimer, 1998).

In the 1950s and even more in the 1960s high growth and low unemployment led to expectations of a very high growth path as a steady state. It was then understandable that many economists had difficulties in adapting their advice to the entirely different situation in the 1970s. Politicians had even greater difficulties to adapt to the new stagnation situation since they were trying to appease the political turmoil and unrest that had accompanied the entry of the post-war baby booms in the industrialized part of the world into adulthood in the end of the 1960s.

This very brief and far from complete age structure story is not only consistent with the estimated age effects and the actual development of the Swedish economy. It is also consistent with conventional economic explanations emphasizing lax monetary and fiscal policies. It only takes a step further back in the causal chain and gives a coherent explanation of why the policies were lax, and why it took so long to come to grips with the stagnation phenomenon. It is, of course, a complex story that is not easily formalized.⁵

Yet there are several other consequences of demographic change which are also likely to play a part, either by reinforcing some of the mentioned mechanisms or by damping them. Let us consider one example. Divorce

⁵But simple OLG models can—even if analytically intractable—reproduce parts of the story. See e.g. Blomquist and Wijkander (1994) where simulations of a three-generation model with a disturbance in the generation size—calibrated on the Swedish baby boom in the 1940s—is capable of reproducing several features of the post-war development.

rates increased rapidly in the beginning of the 1970s, partly because of legislation that made it easier to dissolve the marriage but also because divorces are generally more frequent among young couples, who were relatively more common at that time. This, of course, increased the demand for child care as well as female labor participation. An increasing trend of co-habitation among young couples—as a substitute for marriage—also contributed to an increasing rate of separating couples. Since divorces and separations split households and there are scale effects in consumption—take housing consumption as an obvious instance—the rise in divorce rates contributed further to demands for increased subsidies and the general tendency towards excess demand.

This is not the place to go deeper into the complex demographic interactions with the economy. For the interested reader Easterlin (1981) is a good starting point. Age structure affects the economy through a multitude of economic, social and political mechanisms and is itself also affected by the economy. The final outcome may be possible to describe by simple models like life cycle saving, but the aggregate effects are not a simple summation of individual behavior. Nor is age structure an alternative explanation of economic events, but a complementary perspective going back to ultimate rather than proximate causes.

3.1.2 What have we learned?

Let us now return from this narrative digression to the main purpose of this paper and sum up what we have learned from the data exploration in this section. In spite of potential complexity and large structural changes in the Swedish economy over the post-war period the age distribution shows fairly stable correlations over the whole period to growth and inflation and somewhat less stable to investment and saving. The coefficient patterns are by and large consistent with expected life cycle patterns, although there are anomalies. This may be due to complex interactions with the economy, but the patterns make sense also for more straight-forward explanations.

The magnitude of the effects from individual age groups may be exaggerated by multicollinearity, even if the effects from the group as a whole are reliable. In fact, similar patterns estimated on OECD panel data where it is possible to control for time effects do have smaller coefficients, but still catch the same trends.⁶ The signs of the coefficients may depend on a non-identifiable intercept, and therefore should not be taken too literally. For example, the positive investment effect from young retirees may actually be

⁶See Lindh, 1999; Lindh and Malmberg 1998, 1999a, 1999b and 1999c.

a comparatively weak negative effect.

3.2 Specification issues

The descriptive regressions in the previous subsection are suggestive, but this could easily be a lucky—or not so lucky—coincidence. Before relying on the correlation pattern in the previous section for forecasting some further checking of the econometric specification is in order. In this subsection we leave investment and saving aside and concentrate on growth and inflation as the central variables which we want to forecast using age structure.

3.2.1 Inflation

As remarked above backward recursion of the regressions made the Korea-inflation 1951 stand out. But using a dummy for this year and checking recursive regressions forward from 1946 instead, 1991 and 1992 stand out. With dummies for these years the age pattern of inflation in Table 1 becomes forward stable. In the first column of Table 2 the pure age model with dummies for 1951, 1991 and 1992 is reported. The Jarque-Bera statistic does not reject the null hypothesis of normality, but the Ljung-Box Q-statistic indicates significant high order autocorrelation. In column 2 a regression with two lags of the dependent variable is reported. This does not affect the age coefficients much but high order serial correlation in the residuals is now rejected by the test.

To check for omitted variables a set of controls—investment shares, current account shares, population growth, unemployment rates and GDP growth rates—were included with current values and up to two lags. A great number of other controls could be thought of, but this set of variables proxies for several important influences on money demand and potential GDP. Investment shares measure the demand for investment relative to output. Current account shares is a partial measure of the net flow of foreign currency. Population growth measures both an increase in domestic demand, *ceteris paribus*, and migration that may affect both potential GDP and foreign transactions outside of the current account (by remittances for example). Unemployment rates are standard ingredients in Phillips-Okun type models of inflation. GDP growth rates affect the gap between potential and actual GDP.

Several of these controls had significant t-statistics for one or more of the lags. By successive elimination of variables with t-statistics insignificant at the 1 percent level only current growth and lagged population growth remained. This specification is reported in column 3. Both coefficients are negative. Short-run increases in GDP growth in the current period should

tend to narrow the GDP gap and induce higher inflation. Theoretically we would, therefore, expect a positive effect on inflation. But growth may be simultaneously determined with inflation. Instrumented regressions (not reported) are, however, nearly identical to the reported LS result. The Ljung-Box Q-statistic is not significant but normality can be rejected at the 5 percent level. In column 4 growth was therefore replaced by its lag. This was sufficient to restore normality and also yielded the expected positive coefficient.

To the extent that population growth measures demand pressure, the sign of this variable is also unexpected. Obviously population growth is correlated to age structure, however, and from column 3 we see that the only significant change in the age coefficients is the reversed sign of young adults, which also happens to be the main migration group. Part of this interaction disappears in column 4 when current growth is replaced by lagged growth. Population growth also becomes insignificant. A reasonable conjecture is therefore that the population growth variable in this regression mainly picks up the influence of migration on potential GDP, and hence the coefficient becomes negative.

The age coefficients vary a bit over these specifications as would be expected when we include variables that are correlated with the age variables. The basic form of the coefficient patterns is more persistent. Working age groups and old retirees have a negative relation to inflation and children and young retirees have a positive relation to inflation. The positive coefficient on young adults in the third column is the only real exception to this and is likely to be an interaction effect between population growth, current growth and the age group most sensitive to migration.

3.2.2 GDP growth

The growth equation shows a great deal of instability when the group of retirees is split into young and old since these groups are highly correlated and have similar effects on growth. Therefore they were aggregated for the specification tests. The estimated coefficients also varied a bit as the recession in the 1990s was included in the sample. After some experimentation a dummy for the whole period 1991-1994 turned out to work well in stabilizing coefficient estimates. With the 1951 dummy and the recession dummy the age coefficients are then practically the same for estimations ending in the mid 1970s up to the present. The pure age regression with dummies—reported in the first column of Table 3—indicates higher-order serial correlation by the Ljung-Box Q-statistic.

Including lags of growth in the regression there is a highly significant second lag that improves the fit quite considerably. The magnitude of the age

Dep. var: Inflation	Pure age	Dep var lag	Omitt var I	Omitt var II
Age 0-14	0.364 (1.16)	0.285 (1.11)	0.317 (1.45)	0.782 (3.10)
Age 15-29	-0.208 (1.17)	-0.204 (1.16)	0.496 (3.01)	-0.006 (0.04)
Age 30-49	-0.308 (1.61)	-0.257 (1.60)	-0.177 (1.34)	-0.495 (3.33)
Age 50-64	-0.942 (3.74)	-0.815 (3.67)	-1.064 (6.10)	-1.215 (5.89)
Age 65-74	4.440 (8.83)	3.976 (5.34)	3.043 (4.18)	3.240 (4.68)
Age 75w	-2.095 (5.45)	-1.833 (3.98)	-1.971 (4.28)	-1.055 (2.48)
Inflation(-1)		0.340 (4.99)	0.211 (2.75)	0.413 (7.00)
Inflation(-2)		-0.237 (2.71)	-0.249 (5.61)	-0.163 (1.84)
1951 dummy	0.118 (27.4)	0.122 (42.3)	0.104 (25.7)	0.111 (34.1)
1991 dummy	0.033 (7.47)	0.019 (4.88)	0.018 (4.18)	0.021 (3.52)
1992 dummy	-0.030 (7.53)	-0.033 (8.96)	-0.038 (9.23)	-0.028 (4.67)
Growth			-0.553 (6.81)	
Growth(-1)				0.351 (4.51)
Pop growth(-1)			-3.233 (3.50)	-1.455 (1.41)
Ljung-Box Q(13) p-value	0.016	0.147	0.090	0.156
Jarque-Bera p-value	0.924	0.819	0.048	0.639
adj R ²	0.689	0.756	0.813	0.782
# obs	53	53	53	53
D-W	1.70	1.70	1.87	2.08

Table 2: Post-war sample 1946-1998. OLS regression estimates with specification tests. Absolute t-values are in parentheses. Standard errors are Newey-West corrected for heteroskedasticity and autocorrelation.

coefficients also increases. Tests on the residuals now reject serial correlation and cannot reject that they are normally distributed. Some experimentation showed that a finer division of the young adults group actually made the second lag coefficient insignificant. Young adults have changed their labor force behavior over the period considerably. In the 1950s it was still common to quit school at 16 and start working while this is extremely rare today, when school enrollment is almost total up to the age of 19. However, finer division of age groups leads to more severe collinearity problems. The lag specification is very stable, and does not cause problems with forecasting. Although more, perhaps, could be explained by age structure alone, it did not seem necessary to change the age group division.

There might still be omitted variable bias, so the other variables in the data set were again included with lags to check for that. Investment shares are standard variables in Solow-type growth regressions. Population growth is also often included under the somewhat dubious assumption that it represents work force growth. However, the results in the inflation equation above hints that the migration component of population growth might actually work as a proxy for long-run growth potential. Current account shares reflect the difference between domestic investment and saving, and to the extent that international capital markets provide only imperfect substitution for domestic saving, deficits can reflect capital shortages in production. Unemployment rates reflect closeness to full capacity and might thus pick up adjustment inertia in production of different kinds. Inflation has been found to depress long-run growth in cross-country regressions (Barro, 1997). In short-run models the reverse effect is expected.

After successive elimination of insignificant variables on the 1 percent level the specification in column 3 remained. The most noticeable change in the age pattern is the comparatively stronger positive effect of young adults. Current unemployment and inflation are likely to be correlated with the errors in the growth equation. Instrumentation did not affect the age pattern much, but generally decreased the magnitude of the coefficients a little.

Another change is that the dummies become insignificant. The unemployment variables with almost equal coefficients of opposite signs catch the recession effect as witnessed by the insignificant recession dummy. Swedish unemployment rates have a very obvious structural break that coincides with the recession, and are in comparison more or less constant at a low level before the 1990s and at a high level in the 1990s. It is not surprising that the recession dummy and the unemployment variables proxy each other. Removing unemployment and its lag restores significance to the recession dummy and the age coefficients for the working age groups becomes statistically in-

distinguishable. Unemployment is considerably higher among the young than among the older. The negative accelerator effect in column 3—the differenced unemployment rate has a negative effect on growth—is therefore picking up changes in the composition of the work force.

The 1951 dummy is, however, still insignificant, since the inflation series has a large spike in 1951 that picks up the same effect. Removing the current value of inflation restores significance to the 1951 dummy but the second lag of inflation remains significant and negative, providing some support for the negative effect of inflation on long-run growth. But the residuals now show evidence of autocorrelation. Removing the 1951 dummy in column 5 this disappears.

The coefficient patterns in Table 3 raise some questions. Should dummies be preferred to a genuine explanatory variable like unemployment? Should a highly significant dummy be removed if it induces autocorrelation? I believe the answer to both questions is yes. In the first case unemployment effects on growth mostly cancel out except in the recession period. Even if the accelerator effect from the change in unemployment is theoretically plausible, it is obviously well proxied by the age structure, except in the structural shift in unemployment rates that takes place in the recession years. For a forecasting model it therefore seems to be little to gain by retaining a hard-to-forecast variable like unemployment. In the second case, a highly significant dummy signals an outlier, but the information in that outlier might still be important for a correct identification of the coefficients. If including a dummy for 1951 causes specification problems rather than solving them it seems safer to abstain from it.

4 Medium-Term Forecast Performance

In the specification search above growth and inflation were treated as mutually independent variables. That is not a reasonable assumption. Since they also are clearly collinear with the age variables one should expect a fair amount of interaction with the estimated coefficients. Moreover, the errors in the inflation equation are likely to be correlated with the errors in the growth equation. Forecasting using these variables on the right hand side also requires that the equations are estimated as a system. Thus the equations were simultaneously estimated by Seemingly Unrelated Regression, SUR.

The starting point was the last columns in Table 2 and 3, respectively, but it turned out that lagged growth and the dummies for 1991 and 1992 became insignificant in the inflation equation, so these variables were dropped. To increase reliability the estimation was iterated on the covariance matrix,

Dep var: GDP growth	Pure age	Dep lag	Omitt I	Omitt II	Omitt III
Age 0-14	-1.038 (4.66)	-1.382 (6.93)	-0.616 (3.62)	-1.009 (5.25)	-1.109 (5.22)
Age 15-29	0.398 (2.86)	0.619 (5.17)	0.672 (8.59)	0.673 (9.93)	0.652 (8.72)
Age 30-49	0.621 (5.50)	0.816 (7.04)	0.373 (4.28)	0.617 (6.05)	0.644 (5.77)
Age 50-64	0.690 (3.43)	0.945 (4.82)	0.315 (1.60)	0.637 (3.67)	0.749 (3.81)
Age 65w	-0.877 (6.61)	-1.262 (11.3)	-0.761 (9.34)	-1.027 (8.86)	-1.052 (8.29)
Growth(-2)		-0.558 (8.27)	-0.456 (7.73)	-0.654 (12.3)	-0.636 (13.5)
1951 dummy	-0.018 (4.86)	-0.022 (10.4)	-0.009 (1.43)	-0.027 (12.9)	
1991-1994 dummy	-0.019 (2.87)	-0.026 (4.56)	-0.004 (0.95)	-0.028 (7.82)	-0.027 (7.25)
Inflation			-0.163 (3.45)		
Inflation(-2)			-0.076 (3.58)	-0.209 (3.45)	-0.193 (2.73)
Unemployment			-0.015 (5.99)		
Unemployment(-1)			0.014 (5.11)		
Ljung-Box Q(13) p-value	0.011	0.524	0.206	0.062	0.347
Jarque-Bera p-value	0.678	0.749	0.769	0.531	0.559
adj R ²	0.453	0.658	0.834	0.721	0.700
# obs	53	53	53	53	53
D-W	1.73	2.17	2.25	2.58	2.51

Table 3: Post-war sample 1946-1998. OLS regression estimates with specification tests. Absolute t-values are in parentheses. Standard errors are Newey-West corrected for heteroskedasticity and autocorrelation.

a procedure which approximates a Maximum Likelihood Estimator if the disturbances are normally distributed.

In Table 4 estimates from three different samples are reported. These estimates have been performed for every year 1990 up to 1998, although only a representative subset is reported here. The individual parameter estimates are surprisingly stable.

There is a variation in the estimated age coefficients, mainly in the direction of decreasing magnitude of the age coefficients. In particular, the difference between the two oldest groups diminishes. This is largely due to a change in the collinearity between these groups as we extend the sample over the 1990s. The cross-correlation between these age shares drops relatively much during the 1990s. The recession dummy in the growth equation of course contributes to the stability. But even if that period dummy is dropped, the stability in both equations is still good. This is rather remarkable over a period that is one of the most turbulent in Swedish post war experience.

4.1 Forecasting the 1990s out-of-sample

The reliability of the demographic projections as compared to projections of other explanatory variables is the main reason to conjecture that forecast performance can be considerably improved by using age correlations. Exogenous variables in the forecast equations that must be forecasted themselves cause a considerably higher degree of uncertainty in the forecasts than is generally accounted for in the reported prediction intervals (Diebold, 1998). A recent study concludes that using forecasts conditional on variables that themselves must be forecasted easily doubles the uncertainty of the forecast (Tashman et al., 1998).

The uncertainty in demographic projections concern mortality, fertility and migration, i.e. the inflows and outflows to the population. Mortality projections for the oldest tend to err on the pessimistic side. Future increases in life expectancy are often underestimated, so projections tend to understate the aging of the population. But the errors in absolute numbers over a five-year horizon are comparatively small. Fertility and migration are more troublesome, the errors are larger and are more dependent on economic conditions. However, over a five-year horizon fertility forecast errors will exclusively show up in the group of children and does not affect the older population. Migration is not that concentrated to one age group but still migrants will primarily be in the group of young adults, and errors in the prediction of migration will mostly show up in that group. On the longer horizons for which demographic projections are undertaken, 50 or 100 years

SUR estimates	Inflation			GDP growth		
	1946-90	1946-94	1946-98	1946-90	1946-94	1946-98
Age 0-14	0.639 (1.15)	0.636 (1.29)	0.779 (1.69)	-1.175 (3.88)	-1.119 (3.54)	-1.088 (3.65)
Age 15-29	0.128 (0.39)	0.107 (0.32)	0.212 (0.72)	0.673 (4.40)	0.633 (3.95)	0.626 (4.69)
Age 30-49	-0.351 (1.20)	-0.355 (1.40)	-0.442 (1.89)	0.674 (4.37)	0.644 (3.99)	0.633 (4.17)
Age 50-64	-1.538 (3.26)	-1.501 (3.07)	-1.368 (3.09)	0.761 (2.47)	0.772 (2.39)	0.732 (2.50)
Age 65-74	4.832 (2.89)	4.603 (3.54)	3.524 (4.49)	-1.111 (6.45)	-1.055 (5.95)	-1.021 (6.30)
Age 75w	-2.653 (1.81)	-2.367 (2.51)	-1.540 (2.99)			
Inflation(-1)	0.271 (2.92)	0.288 (3.15)	0.301 (3.50)			
Inflation(-2)	-0.306 (3.01)	-0.287 (2.91)	-0.272 (2.93)	-0.114 (1.58)	-0.181 (2.60)	-0.182 (3.06)
Growth(-2)				-0.546 (5.82)	-0.588 (6.28)	-0.570 (6.83)
Pop. growth (-1)	-2.657 (1.77)	-2.115 (1.40)	-2.805 (2.27)			
Dummy 1951	0.097 (6.52)	0.098 (5.97)	0.098 (6.52)			
Dummy 1991-94					-0.024 (3.19)	-0.027 (4.12)
Correlation resid.	-0.513	-0.514	-0.518			
adj R ²	0.709	0.687	0.728	0.664	0.692	0.697
# obs	45	49	53	45	49	53

Table 4: Variation of the sample using the ...nal forecasting speci...cation. System estimation by iterated SUR. Absolute t-values are in parentheses. Standard errors are Newey-West corrected for heteroskedasticity and auto-correlation.

into the future, the cumulated errors will spread over the whole population and become quite considerable. The first 5 or 10 years of a projection are fairly reliable though, and in comparison to standard projections of economic variables at that horizon the forecast error is much smaller.

From the previous regressions we can see that the slow-moving age variables are quite good at picking out medium-run persistence or trend variations in the data. But suspicious econometricians may well ask if the age variables are picking up a time trend or a persistence phenomenon with other causes that just happen to coincide nicely with the variation in age structure over this period. The relative stability over samples and specifications cannot be said to confirm those suspicions, but a more stringent test is prediction out-of-sample.

Stability of the coefficients and a fairly good fit are not sufficient reasons to expect good forecasting performance. To get as close as possible to an out-of-sample forecast, the model was first estimated using actual data up to 1991. Then the forecast was made using predicted values for the lags of growth and inflation and replacing the actual demographic data by projections made by United Nations (1994).⁷ The main difference between the projections and actual demographic data is that the projections are based on assumptions of a continuation of the fertility rates in Sweden around 1990. Since fertility rates have fallen drastically from around replacement level of 2.1 children per woman 1990 to currently around 1.5 children per woman, the projection has a too high share of children in the 1990s. Statistics Sweden made similar assumptions in the beginning of the 1990s, since the drop in fertility rates after 1990 was unexpected. So was the recession that probably caused much of this drop in the total fertility rate by the accompanying cuts in the welfare systems.

In Figure 3 the dynamic forecasts using data up to 1991 and SUR estimation can be compared to the actual data. The forecasting model was estimated from the specifications in Table 4, and forecasts were made using the UN data to simulate demographic projections in 1991. Forecasted values were used for growth and inflation lags. The sample interval here was chosen to include 1991 in order to get a coefficient estimate for the recession dummy, so it could be included in the growth forecast. Recall that the 1991 and 1992 dummy were dropped in the inflation equation for the system estimates.

The prediction intervals in Figure 3 have been computed by a RATS procedure that ignores the sampling error (in coefficient estimates) and only

⁷The UN Population Division makes their own estimates and projections using methods designed to get comparability across countries. The differences to Statistics Sweden estimates and projections are, however, relatively small, since Swedish raw data are very reliable.

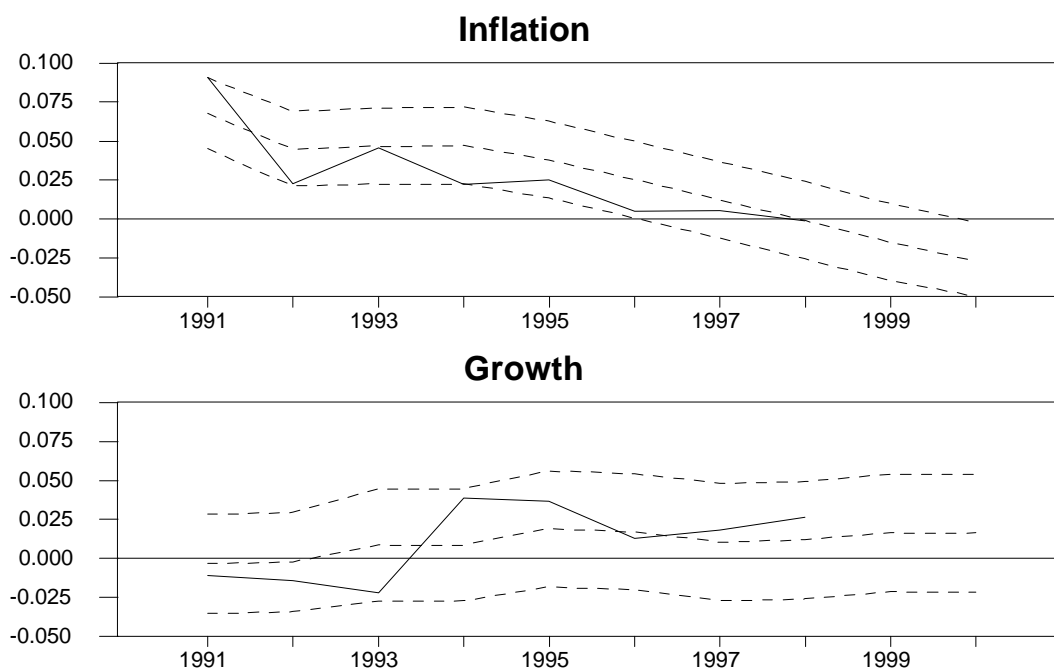


Figure 3: Out-of-sample forecasts 1991-2000 from model estimated on data 1946-1991 by SUR. 95 % prediction intervals. The unbroken curves are actual inflation and actual GDP growth.

takes into account the innovations. The reason is that the lag coefficient uncertainty enters the sampling errors in a very complex way. These intervals are therefore lower bounds on the true 95 percent prediction intervals.

The actual observations are nevertheless within these prediction intervals. The inflation forecast tend to be systematically higher than the actual values—although this will not be true for 1999. The growth forecasts tend to undershoot actual values in the end of the 1990s. Both these features have to do with the overshooting projections for the children 0-14 years old. Using actual age data to forecast, the impact of the faulty fertility assumptions in the demographic projection can be assessed. The inflation forecast then tends to be more deflationary and the growth forecast tends to be more optimistic.

To test the stability of the forecasts the same forecast setup was repeated by adding a year of data and then forecasting on the UN data. This is not quite a realistic out-of-sample experiment, since we would like to update the demographic projections as the estimation sample is extended. In Figure 4 these recursive forecasts are shown without prediction intervals in order

to be distinguishable. For the same reason only forecasts separated by two years are shown. The inflation forecasts for 1999 at different horizons tend to undershoot. This goes for other years in the end of the decade also but otherwise the forecasts are fairly accurate up to 1996. The growth forecasts, on the other hand do not predict the recession but become increasingly accurate over time, even if the point estimates undershoot actual growth somewhat towards the end.

Using actual age data instead (not shown) the inflation forecasts show undershooting also earlier in the decade while the growth forecast improves considerably with all point forecasts for 1999 concentrated in a rather narrow band between 3 and 4 percent growth except the 1990 forecast of 2.6 percent. Since updated demographic projections would have adjusted the fertility assumptions one can conjecture that the true out-of-sample performance of the age forecasts should lie somewhere in between, i.e. fairly good for both equations but with a tendency for underprediction at long horizons for the inflation equation.

In Table 5 the out-of-sample forecast statistics are reported for the forecasts in Figures 3 and 4. There are 8 successive forecast series, and the number of steps in each forecast series decreases as we get closer to the sample end. Therefore there are only four observations of 5-step forecasts as compared to 8 of the 1-step forecasts. The mean error is simply the average of the observed forecast errors. MAE is the mean absolute error, which should be greater in absolute value than the mean error unless there is a systematic under- or overprediction. The RMS is the square root of the mean squared error, which penalizes large errors relatively more than the MAE and therefore should be higher unless all errors are of equal size. The Theil U statistic compares the mean forecast error to the mean error made by a "naive" forecast.⁸ The latter is the forecast that the last observed value will hold in the future, too, i.e. the best forecast if the forecasted variable follows a random walk without drift. If Theil's U statistic is unity we are doing no better than the naive forecast, if it is below unity we are doing better.

Given that there is a downward trend in inflation over the period it might seem that a random walk incorporating that drift would be a better comparison. However, it is obvious from Figure 2 that an estimate in 1990 of the

⁸The exact definition of Theil's U for an i-step ahead forecast is

$$U = \frac{\sqrt{\frac{1}{n} \sum_{t=1}^{n-i} \frac{z_{t+i} - \hat{z}_{t+i}}{z_t}^2}}{\sqrt{\frac{1}{n} \sum_{t=1}^{n-i} \frac{z_{t+i} - z_t}{z_t}^2}}$$

where \hat{z}_{t+i} is the forecasted value and z_{t+i} is the actual outcome.

linear trend in inflation would be highly dependent on the estimation period. Using the whole post-war sample would yield a positive drift, starting in 1980 the trend would be negative (but less so than in the 1990s) while weighting recent observations in the end of the 1980s more heavily would tend to shift the trend towards positive drift again. A comparison to a random walk with drift is thus very sensitive to how the drift is estimated. If a positive drift was estimated in 1990 the age forecasts would do even better in comparison. If a negative drift was estimated they would do relatively worse.

Both the growth equation and the inflation equation are doing better than the naive forecast and in fact better and better the longer the forecast horizon is. The tendency to underpredict towards the end of the decade in the growth equation is discernible in the closeness of MAE and the mean error for the 4-step and 5-step forecasts, which only are observed in the latter half of the decade. The MAE and RMS measures are rather close. But, of course a mean absolute error of 1-1.5 percentage points in the growth rate and a little less in the inflation rate (except for the larger errors in the 1-step forecasts) may still be considered rather large.⁹

Errors of 1 percentage point or more are frequently made in standard forecasts of inflation and growth, and for the latter variable this even holds for preliminary estimates after the year in question has passed. What distinguishes the age forecasts is rather that the forecast errors do not show any tendency to rise with the time horizon (up to 5 years ahead). It is true that the short-run forecasts perform badly and this may pose problems for attempts to use age forecasts in monetary policy decisions aimed at stabilization.

It could be argued that a good trend forecast in fact should be filtering out the short-term variation just like a Hodrick-Prescott filter does. Then the RMS and the MAE are not the relevant measures of forecast accuracy since they measure the size of such short-term errors. The mean error, however, lets forecast errors around the trend line cancel out. It is therefore a relevant measure for how good the trend forecast is.

The same forecast statistics as in Table 5 for forecasts using actual age

⁹As a comparison the root mean squared forecast errors over the period 1985 to 1993 in the United States for a selection of two-year ahead forecasts of GDP inflation based on leading indicators in Staiger et al. (1997) are comparable in size to those reported here for inflation for 5-year ahead forecasts. Over the period 1975-1984 Staiger et al. report RMS of more than 2 percentage points. It might be argued that inflation volatility was higher for US inflation over those periods than for Sweden 1990-1998 and thus the comparison is not fair. However, computing standard deviations for yearly inflation rates from monthly CPI indices for these periods, the standard deviation is 2.97 percent for the US 1975-1984, 1.10 percent for the US 1985-1993 and for Sweden 1990-1998 the standard deviation is 3.55 percent.

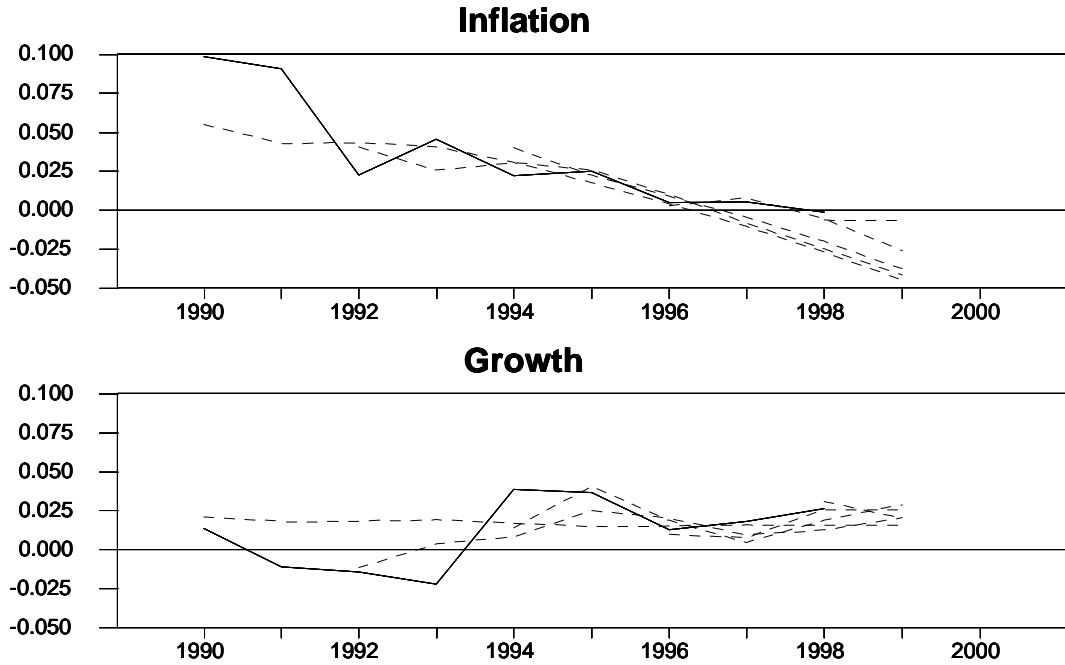


Figure 4: Out-of-sample forecasts from estimated age models using data up to 1990, 1992, 1994, 1996, and 1998, respectively. The unbroken curves are actual data.

	Steps ahead	Obs	Mean error	MAE	RMS	Theil's U
Growth	1	8	-0.26	1.53	1.94	0.77
	2	7	-0.27	1.58	2.06	0.59
	3	6	0.13	1.69	2.19	0.59
	4	5	1.00	1.19	1.41	0.43
	5	4	0.96	1.22	1.34	0.40
Inflation	1	8	-0.13	1.85	2.18	0.78
	2	7	-0.42	0.76	0.84	0.24
	3	6	-0.16	1.15	1.28	0.31
	4	5	-0.18	0.94	1.15	0.23
	5	4	0.12	1.20	1.37	0.22

Table 5: Forecast statistics for out-of-sample forecasts of growth and inflation over the period 1990-1998. Errors are measured in percentage units.

data do not change the picture radically. The inflation equation then performs worse and the growth equation performs better, but the trend that Theil's U decreases and that there is no tendency for RMS and MAE to increase with the forecast horizon still holds. The inflation equation, however, shows signs of systematic underprediction. Table 5 therefore probably overstates the performance of the inflation forecasts during the 1990s and understates the performance of the growth forecasts. If that conjecture is true the conclusion is that monetary policy in Sweden during the 1990s have been partly effective in counteracting the deflationary pressure from the age structure. Not wholly effective since the target band has been missed from below by the annual CPI inflation rates for the last three years.

As was pointed out in the introductory remarks there has been a regime shift in Swedish monetary policy since 1992 with an explicit target band for the CPI inflation. This regime shift may have changed the correlation between age structure and inflation, not by any explicit policy to sterilize age effects, but simply by aiming at the target rate, it will have partly counteracted the deflationary age structure influence. Taking explicit account of the age pressures could have been helpful by suggesting a somewhat higher aim, in much the same way as an archer compensates for gravity by aiming above the bull's eye. Thus, age forecasts of the inflation trend could be helpful in policy design even if there are systematic biases in the forecasts.

4.2 Phillips-Okun relations

Much of the literature around the empirical implementation of inflation forecasting and output gaps either builds on some type of Phillips-Okun relation or some filtering approach, sometimes combined into one.¹⁰ Crucial to the more structural approaches is the identification of potential GDP. That is, of course, also the weak point, since potential GDP is not only unobservable, there is no consensus about exactly what it is. Somewhat vaguely there is agreement that potential GDP should measure the long-run steady-state GDP that could be achieved by full utilization of all resources without introducing any excessive inflation pressure into the economy. How to measure this concept in practice is an entirely different question. Measures of capacity utilization are of little help when it comes to prediction, so often practitioners simply use some smoothed measure of average growth to generate their potential GDP measures. More sophisticated approaches (see e.g. Apel and Jansson, 1999b) attempt to identify unobserved components in the data and

¹⁰Concerning practical matters see the overview on potential output by Giorno et al. (1995), for a more policy-oriented approach see Clarida et al. (1999).

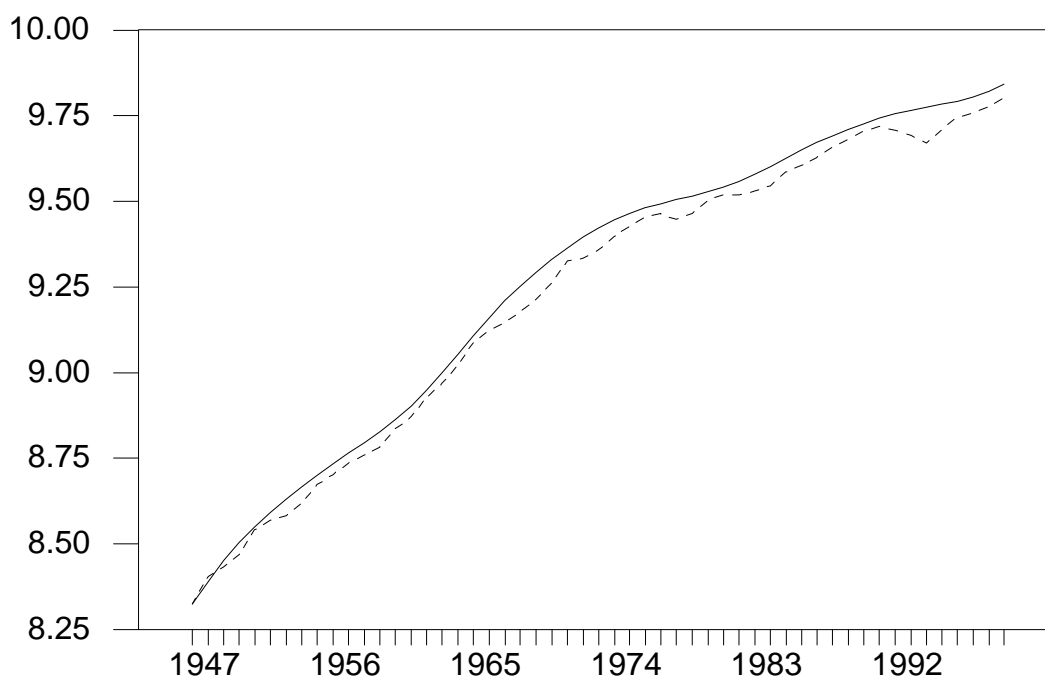


Figure 5: "Age Potential GDP". The logarithm of the GDP measure implied by the age structure effect on growth 1946-1998. The broken curve is the logarithm of real GDP.

acknowledge that there is more to the Phillips relation than just a simple bivariate relation.

In this context we have available another notion of potential GDP. Age structure has a very strong correlation to growth, verified in several empirical studies over the last few years, as well as in the estimates above. It is clear from the estimations above that this is due to a medium-frequency correlation that picks up medium-term trends and not short-run variation. It therefore serves as a proxy for what we might call potential GDP. If we take potential GDP in a more practical sense as the full capacity output, it is not unreasonable to assume that this is dependent on the human resources in the economy, which in turn are closely related to the age structure of the population. Whether this measure is consistent with a non-inflationary output level is another matter which lies outside the scope of this paper.

Figure 5 shows the result of unraveling the growth estimate from a pure age share regression 1946-1998 to construct a potential GDP measure. Starting from the logarithm of GDP in 1946 the age predicted growth rate, $\hat{g}_{1947} = \log \hat{Y}_{1947} - \log Y_{1946}$, can be used to solve for $\log \hat{Y}_{1947}$, which in turn

Correlations	Unobs. comp.	Whitt.-Hend.	Prod. function
Age GDP gap	0.80	0.77	0.71
Unobserved components		0.85	0.95
Whittaker-Henderson			0.84

Table 6: Cross-correlations between different GDP gap measures

can be inserted into \hat{g}_{1948} to solve for $\log \hat{Y}_{1948}$, etc. The resulting series seems to be an uncannily good measure of a capacity bound in the economy. But this is actually a coincidence since the initial value for GDP in 1946 happens to be in a high growth year (also as predicted by age structure).

The log difference $y = \log Y_t - \log Y^{\text{pot}}$ is the conventional definition of the GDP gap. In this case it will be negative for every year (except 1947). This is purely a matter of scale. Figure 6 compares this gap to gap series on a quarterly frequency 1980 to 1998 which have been computed at the central bank of Sweden (Sveriges Riksbank). The UC curve in Figure 6 is from an unobserved components model, the PF curve is a production function based measure and the WH curve is the Whittaker-Henderson gap implied from a Hodrick-Prescott filtered GDP series.¹¹ The age generated gap has been expanded from annual frequency to quarterly, using the same value for each quarter. Apart from the step character of the curve that follows from this expansion it differs from the other measures primarily by being all negative, a scale effect of no consequence.

The variation in the curve is otherwise rather similar to the average annual variation we would get by taking averages of the other series. The correlation between these different measures in Table 6 confirm that impression. The age generated GDP gap is nearly as closely correlated to the other measures as they are to each other. In fact the difference is mainly due to the difference in frequency. There is one difference in Figure 6, though, that is discernible by the naked eye. The other three measures, especially the WH filter, imply a downturn in the potential GDP gap from 1995 to 1996, which is absent in the age gap. In spite of minor differences it seems safe to conclude that the age generated measure of potential GDP catches much the same trend as other measures do. It therefore seems worthwhile to use the gap as regressor in order to see whether we can arrive at a conventional Phillips-Okun relation.

In Table 7 the results are reported for a number of successively extended specifications. Only lags that were statistically significant or close to being significant have been included in the table. The age GDP gap is significant

¹¹The methods are documented in Apel et al. (1996). The series used here are the latest updates available. A fourth measure is presented in Jacobson et al. (1999).

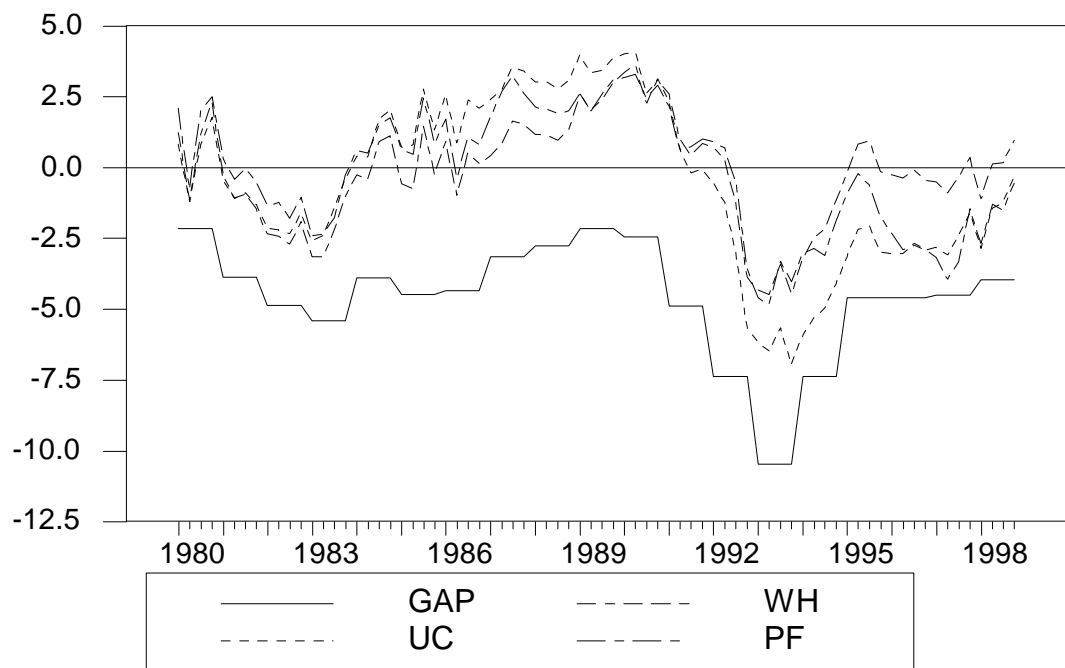


Figure 6: Comparison of the age structure implied GDP gap (GAP) and three other measures: an unobserved component (UC) measure, a production function based (PF) measure, and a Whittaker-Henderson (WH) gap. The latter is the resulting gap from a Hodrick-Prescott ...ltered GDP series.

but does not by itself explain very much of the variation in inflation.¹² The first lag of the GDP gap consistently has a positive coefficient meaning that a large (negative) GDP gap last year will depress inflation this year. The second lag is negative so there is an accelerator mechanism where an increasing gap depresses inflation and vice versa. In fact the third column hints at the possibility that it is exclusively an accelerator with no effect from the level of the GDP gap. But deleting the second lag leaves the first lag still significant with a coefficient that is approximately half the previous estimate.¹³ Excluding the GDP gap in column 4 worsens the fit considerably and also lowers the significance of the lagged inflation and unemployment variables. An accelerator mechanism with some role for the level therefore seems consistent with this evidence.

But there are still age effects that are not captured by the GDP gap variable. In the last column we see that the remaining age pattern is fairly similar to the one obtained in previous regressions. The main difference is that the negative effect from the oldest disappears. This is an indication, albeit weak, that the main factor behind the negative inflation effect of old retirees is the strong negative effect on investment in Table 1, since investment is fairly closely correlated to GDP variations. Given that we have controlled for age effects through the GDP gap one need to examine the interpretation of the GDP gap which conventionally is thought of as a measure of excess demand.

The age generated GDP gap does not fit so nicely into this interpretation. In fact a demand explanation for the remaining age effects is more convincing. That means the net consuming groups fan inflation and the net saving groups dampen inflation. In comparison to the pure age model in Table 1 the dampening effect of the oldest group now becomes insignificant and the positive effect from young retirees a little more moderate. Children on the other hand have a stronger positive effect and the point estimate for young adults becomes positive. Thus, one would conjecture that the GDP gap implied by age effects on growth works as a supply shock through factor supplies in this context while the main age effect on inflation is a demand effect through consumption which is not captured by the GDP gap. Inclusion of the age variables renders the third lag of the unemployment rate insignificant. In fact, none of the current or first three lags of the unemployment rate are significant when the age variables are included.

Tests of the forecasting performance of these equations were disappointing. At the 3-5 year horizon it did considerably worse than the simple

¹²Since it is a generated regressor the standard errors from the OLS regression may be misleading, see Pagan (1984).

¹³Apel and Jansson (1999a) find a similar accelerator mechanism from unemployment on Swedish inflation.

Dep var	Inflation				
Constant	0.071 (5.27)	0.026 (2.94)	0.043 (3.44)	0.035 (3.75)	
GDP gap(-1)	0.783 (3.05)	1.047 (3.39)	1.150 (4.10)		1.197 (3.76)
GDP gap(-2)	-0.379 (1.03)	-0.833 (2.73)	-1.014 (4.15)		-0.475 (1.55)
Inflation(-1)		0.682 (6.03)	0.611 (5.24)	0.533 (3.89)	0.225 (2.06)
Unempl(-3)			-0.652 (4.04)	-0.411 (2.96)	-0.568 (2.05)
Age 0-14					1.391 (4.00)
Age 15-29					0.247 (0.98)
Age 30-49					-0.876 (4.74)
Age 50-64					-1.186 (3.71)
Age 65-74					1.816 (1.81)
Age 75 above					0.582 (0.47)
Ljung-Box Q(13) (p-value)	61.46 (0.000)	4.44 (0.986)	6.09 (0.943)	11.42 (0.575)	7.65 (0.866)
adj R ²	0.081	0.494	0.560	0.363	0.635
# obs	51	51	51	53	51
D-W	0.65	1.97	2.04	1.96	1.94

Table 7: Post-war sample 1946-1998. Using a GDP gap variable computed from an age model of growth. Absolute t-values are in parentheses. Standard errors are Newey-West corrected for heteroskedasticity and autocorrelation.

inflation-growth system reported above. It is still possible that a more careful specification might lead to something useful, but it did not seem worthwhile to pursue this here.

Summing up, the evidence presented here supports using a direct approach to forecast inflation and growth simultaneously by age structure variables. Predicting potential GDP first seems to be an unnecessary detour if the purpose only is to forecast, an argument also made in Jacobson et al. (1999).

However, in an explicit reaction function derived from an objective function for the Central Bank it might be more useful. That is provided that the Central Bank attaches some weight to the GDP gap in its loss function. In that case it is, of course, no detour to estimate and forecast the potential GDP by a growth regression on age structure. However, before such an exercise will be worthwhile a closer study of how the policy tools interact with these measures is called for. Recall that this potential GDP concept as such carries no implication of being neutral with respect to inflation. A further understanding of the policy interactions may also yield means to control for the possible change in the relation between age structure and inflation that is induced by policy.

4.3 Two further sensitivity tests

Obviously the trend forecasts are missing out on the short-term variation, as they should I hasten to add. This raises the question whether the trend forecasts can be combined with short-term predictors to achieve an overall improvement in forecasts? Time series of leading indicators are not available for the long time spans required to get sensible age structure estimates, so to investigate that question a detour was necessary. First, the age predicted inflation from the standard forecast equation above was expanded to quarterly frequency and smoothed by a centered moving average filter to get rid of the step character. Then this age predicted series was deducted from quarterly inflation figures 1984Q1 to 1998Q1. Finally this residual series was regressed on sets of leading indicators with several lags. These leading indicators are used by Sveriges Riksbank in business cycle analysis: electricity production, retail sales, hours worked and manufacturing production. The idea was that the indicators should pick up residual variance that was not predicted by age variables.

This idea failed, however, in an interesting way. The leading indicators did not have any predictive value for the residual inflation. The adjusted R^2 was in fact strongly negative. Regressing the smoothed age predicted inflation series on the indicators yielded an adjusted R^2 of 0.72 while using

actual inflation as the dependent variable the adjusted R^2 still was negative, although less strongly. The set of regressors can surely be trimmed to yield some predictive value for the quarterly inflation series. The point is, however, that this was not needed for the age predicted series. The conclusion is that the age structure seems to have an information content that strongly overlaps the information set of these leading indicators.

As a formal test of whether multicollinearity was driving the results vector autoregressions (VAR) were run using inflation and growth as dependent variables including age structure as deterministic variables. The results were much the same as for the previous SUR-estimates and are therefore not reported here. Age coefficients do change when lags are included, and the basic patterns may shift both vertically—i.e. with respect to the intercept—and horizontally—i.e. with respect to age groups—when the number of lags included and the sample period is varied. However, forecasts were fairly similar to those obtained from the specification in Table 4 when using actual age data. That is to say a tendency to underpredict inflation at longer horizons, but getting growth fairly right in the multi-step forecasts, while short-term performance is considerably worse. I therefore did not proceed to genuine out-of-sample forecasts. The VARs pick up some of the age variation but seem to be more or less equivalent to the SUR-estimates when it comes to forecasting. This might be taken as corroborative evidence that the regime shift in monetary policy indeed has suppressed inflation responses to age structure changes in the 1990s.

5 Conclusions

This paper has studied Swedish post-war macroeconomic time series and finds a comparatively stable post-war correlation between age structure and growth and inflation, respectively. Evidence has been presented that statistical models of inflation and GDP growth explained mainly by population age shares do not perform too badly in out-of-sample forecasting on a horizon of 3-5 years ahead during the 1990s. There are, however, indications of systematic underprediction for the inflation rate. This may be a consequence of the inflation target for monetary policy in Sweden that has been operational for most of the 1990s. However, it may also be a consequence of a general structural break in the Swedish economy in the beginning of the 1990s. For example the disastrous rise in Swedish unemployment from an average annual unemployment rate 1946-1991 of less than 2 percent that never exceeded 3.5 percent to over 8 percent of the labor force in 1993. The growth forecasts are considerably more stable and less biased.

Constructing potential GDP measures from the fitted values of a growth model predicted exclusively by age shares of the population conform well to standard GDP gap measures. This is rather remarkable since these measures have been derived from quarterly data of much shorter series with entirely different methods. Using the age generated GDP gap in a Phillips-Okun equation works but there are still remaining age effects on inflation. This equation also forecasts less well than an inflation-growth system.

To be useful for monetary policy analysis further research is needed to clarify the connection between policy instruments and the economic mechanisms that cause the strong correlation between age structure and macroeconomic variables like inflation and growth. The aim of monetary policy is short-run stabilization but it is commonly believed that policy instruments work with a substantial lag. In order to avoid overshooting reactions when the trends change a close look at longer horizons is well advised. More research on the potential for using demographic projections to that end can have huge pay-offs.

More generally this paper confirms previous results on many other samples, that demography matters for central macroeconomic variables, in this case growth and inflation. Even though this introduces further complexity into macroeconomic modelling, the sheer size of the effects—often around half of the total variation is explained by demographic variables—is a reason to spend considerable research effort on clarifying and understanding the economic mechanisms that mediate this impact.

References

- Andersson, B.: 1998, Scandinavian evidence on growth and age structure, Working Paper 1998:4, Department of Economics, Uppsala University, Uppsala.
- Andersson, B.: 1999, Age dispersion and risk aversion: A cohort analysis of Swedish households' asset portfolios 1982-1992. Work in progress.
- Apel, M., Hansen, J. and Lindberg, H.: 1996, Potential output and output gap, Sveriges Riksbank Quarterly Review (3), 24–35.
- Apel, M. and Jansson, P.: 1999a, System estimates of potential output and the NAIRU, *Empirical Economics* 24(3), 373–388.
- Apel, M. and Jansson, P.: 1999b, A theory-consistent system approach for estimating potential output and the NAIRU, *Economics Letters* 64, 271–275.
- Barro, R. J.: 1997, *Determinants of Economic Growth: A Cross-Country Empirical Study*, MIT press, Cambridge, Massachusetts.
- Becker, G. S.: 1962, Investment in human capital: A theoretical analysis, *American Economic Review* 52, 9–49.
- Blomquist, N. S. and Wijkander, H.: 1994, Fertility waves, aggregate savings and the rate of interest, *Journal of Population Economics* 7, 27–48.
- Bloom, D. E., Canning, D. and Malaney, P. N.: 1999, Demographic change and economic growth in Asia. mimeo, Harvard Institute of International Development, Cambridge, Massachusetts.
- Bloom, D. E. and Sachs, J. D.: 1998, Geography, demography, and economic growth in Africa, *Brookings Papers on Economic Activity* (1998:2), 207–273.
- Bloom, D. E. and Williamson, J. G.: 1998, Demographic transitions and economic miracles in emerging Asia, *World Bank Economic Review* 12(3), 419–455.
- Bosworth, B., Burtless, G. and Sabelhaus, J.: 1991, Fertility waves, aggregate savings and the rate of interest, *Brookings Papers on Economic Activity* 1:1991, 183–256.

- Clarida, R., Gali, J. and Gertler, M.: 1999, The science of monetary policy: A new keynesian perspective, *Journal of Economic Literature* . Forthcoming.
- Diebold, F.: 1998, *Elements of Forecasting*, South-Western, Cincinnati.
- Easterlin, R. A.: 1981, *Birth and Fortune*, Grant McIntyre, London.
- Fair, R. C. and Dominguez, K.: 1991, Effects of the changing U.S. age distribution on macro-economic equations, *American Economic Review* 81, 1276–1294.
- Giorno, C., Richardson, P., Roseveare, D. and Van Den Noord, P.: 1995, Potential output, output gaps and structural budget balances, *OECD Economic Studies* (24), 167–209.
- Herbertsson, T. T. and Zoega, G.: 1999, Trade surpluses and life-cycle saving behaviour, *Economics Letters* 65(2), 227–237. Forthcoming.
- Higgins, M.: 1998, Demography, national savings, and international capital flows, *International Economic Review* 39(2), 343–369.
- Higgins, M. and Williamson, J. G.: 1997, Asian demography and foreign capital dependence, *Population and Development Review* 23(2), 261–293.
- Jacobson, T., Jansson, P., Vredin, A. and Warne, A.: 1999, A VAR model for monetary policy analysis in a small open economy, Working Paper 77, Sveriges Riksbank, Stockholm.
- Johansson, Ö.: 1967, *The Gross Domestic Product of Sweden and its Composition 1861-1955*, Almqvist & Wicksell, Stockholm.
- Jones, M. T. and Obstfeld, M.: 1997, Saving, investment, and gold: A reassessment of historical current account data, Working Paper 6103, National Bureau of Economic Research, Cambridge, Massachusetts.
- Krantz, O.: 1997, Swedish historical national accounts 1800-1990 – aggregated output series. manuscript.
- Krantz, O. and Nilsson, C.-A.: 1975, *Swedish National Product 1861-1970: New Aspects on Methods and Measurement*, CWK Gleerup, Lund.
- Lenehan, A.: 1996, The macroeconomic effects of the postwar baby boom: Evidence from Australia, *Journal of Macroeconomics* 18, 155–169.

- Lindert, P. H.: 1994, The rise of social spending, 1880-1930, *Explorations in Economic History* 31(1), 1–37.
- Lindh, T.: 1999, Age structure and economic policy—the case of saving and growth, *Population Research and Policy Review* 18(3), 261–277.
- Lindh, T. and Malmberg, B.: 1998, Age structure and inflation—a wicksellian interpretation of the OECD data, *Journal of Economic Behavior and Organization* 36(1), 19–37.
- Lindh, T. and Malmberg, B.: 1999a, Age structure and the current account—A changing relation?, Working Paper 1999:21, Department of Economics, Uppsala University, Uppsala.
- Lindh, T. and Malmberg, B.: 1999b, Age structure effects and growth in the OECD, 1950-90, *Journal of Population Economics* 12(3), 431–449.
- Lindh, T. and Malmberg, B.: 1999c, Can age structure forecast inflation trends?, *Journal of Economics and Business* p. Forthcoming.
- Lindh, T. and Malmberg, B.: 1999d, Demography and housing demand—what can we learn from residential construction data? Paper presented at Workshop on Age Effects on the Economy, Stockholm 2-3 June.
- Malmberg, B.: 1994, Age structure effects on economic growth: Swedish evidence, *Scandinavian Economic History Review* 42(3), 279–295.
- McMillan, H. M. and Baesel, J. B.: 1990, The macroeconomic impact of the baby boom generation, *Journal of Macroeconomics* 12(2), 167–195.
- Modigliani, F. and Brumberg, F.: 1954, Utility analysis and the consumption function: An interpretation of cross-section data, in K. Kurihara (ed.), *Post-Keynesian Economics*, New Brunswick.
- Newey, W. K. and West, K. D.: 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55(3), 703–708.
- Ohlsson, L.: 1969, *Utrikeshandeln och den ekonomiska tillväxten i Sverige 1871-1966*, Industriens Utredningsinstitut, Stockholm.
- Ohlsson, R.: 1986, *Högre utbildning och demografisk förändring*, Ekonomisk-historiska föreningen, Lund.

- Pagan, A.: 1984, Econometric issues in the analysis of regressions with generated regressors, *International Economic Review* 25(1), 221–247.
- Persson, J.: 1998, Demographics, human capital, convergence and growth—an analysis of the US states 1920-1990, *Essays on Economic Growth*, Institute of International Economic Studies, Stockholm University, Stockholm. Doctoral thesis.
- Shimer, R.: 1998, Why is the U.S. unemployment rate so much lower?, *NBER Macroeconomics Annual* pp. 13–61.
- Staiger, D., Stock, J. H. and Watson, M. W.: 1997, The NAIRU, unemployment and monetary policy, *Journal of Economic Perspectives* 11(1), 33–49.
- Tashman, L. J., Bakken, T. and Buzas, J.: 1998, Effect of regressor forecast error on the variance of regression forecasts. mimeo, University of Vermont.
- United Nations: 1994, *Global Estimates and Projections of Population by Sex and Age: The 1994 Revision*, New York.

Data appendix

This appendix describes sources and definitions of the data. The original data set comprised annual time series 1860-1998 of economic and demographic data. Actually used in this paper is the period 1944-1998 where the first two years were added to provide lags for estimations over the post-war period 1946-1998. The detailed accounts below do not give any details for the parts of the data that has not been used here.

The leading indicators 1984Q1-1998Q4 and the gap variable data 1980Q1-1998Q4 were compiled by Tomas Lindström from data bases at Sveriges Riksbank.

GDP

GDP at factor costs in current prices 1860-1980 and a deflator are from Krantz (1997). It should be noted that these estimates are still not official. An account of methods and more details on the construction of the estimates can be found in Krantz and Nilsson (1975). For the period 1941-1950 the sum of sector GDPs were used since the aggregate for that period is not accurate according to the author. These series were spliced to the latest update from Statistics Sweden (October 1999) by ratio linking in 1980. Real GDP is the deflated series. Due to changes in the National Account system the two last years 1997-1998 refer to GDP at market prices, and have been ratio-linked to the factor cost series.

CPI

The consumer price index is spliced in 1980 by ratio linking using the latest update of annual CPI from Statistics Sweden (October 1999) and historical cost-of-living indices from Statistics Sweden that go back well beyond 1860.

Population

The basic age structure data are five-year cohort population numbers on annual basis. 1911 up to 1967 have been compiled from official statistics by Bo Malmberg. 1968 to 1998 are from the latest updates of Statistics Sweden. Projections 1999-2030 are an updated version of the 1997 forecasts of Statistics Sweden. Note that all these population data refer to 31 December in the current year, but in the estimations they have been lagged one year. Projections of the Swedish population in the 1990s are from United Nations (1994).

Current Account

1861-1945 is a series constructed by Jones and Obstfeld (1997) which builds on earlier Swedish sources. For the period 1950-1998 the latest update from the Central Bank of Sweden was obtained. The missing years 1946-1949 were filled in by data from Ohlsson (1969) table B:1 column 6 without linking. The overlap to the Central Bank data have somewhat larger differences but they are diminishing as we go back in time, so it seemed reasonable not to use any ratio link in 1950 either. Ohlsson (1969) compiles earlier data sources on the balance of payments up to 1966.

Gross investment

Current value estimates 1861-1945 are taken from J&O which is essentially K&N only adding change in live stocks from Johansson (1967) in order to arrive at a measure more suited to compute savings from. The lacuna 1946-49 was closed by splicing to K&N domestic investment by ratio linking in 1945. The latest updates from Statistics Sweden (April 1999) were then ratio-linked in 1950 to this series. Note that due to definitional changes these data are spliced by ratio-linking in 1980, too.

Unemployment

Unemployment rates 1911-1997 were graciously provided by Anders Forslund, who has linked modern data from Labour Force Surveys (AKU) to statistics derived from trade union data on recipients of unemployment benefits. The last years up to 1998 have been updated by the latest Labour Force Survey statistics.