THE FISCAL AND ECONOMIC OUTLOOK

ADDRESS TO THE AUSTRALIAN BUSINESS ECONOMISTS

SYDNEY

15 MAY 2007

KEN HENRY

SECRETARY TO THE TREASURY

THE FISCAL AND ECONOMIC OUTLOOK[1]

When I was introduced to the study of macroeconomics at university level in the second half of the 1970s, people were getting used to the notion that economies could simultaneously experience high rates of inflation and substantial unemployment. There was much academic debate about the stability of Phillips curve relationships; whether the short-run and long-run relationships might be quite different; and some emerging work on the role of expectations that became considerably more influential as the decade progressed and well into the 1980s.

In the early 1980s policy advisers on one side of a debate argued that priority should be given to fighting inflation ('fight inflation first') before tackling unemployment, while those on the other side of the debate argued that governments should do whatever it took to get people into jobs, no matter the inflationary consequences.

Had anybody walked in on one of those debates and presented a scatter graph of actual inflation rate and unemployment rate 'pairs' for the Australian economy in the post-1990s recession period, they would have been ridiculed.

That graph, of course, shows a trend decline in the unemployment rate with no inflationary consequences to speak of. Much has been written about the reasons for this extraordinary macroeconomic performance. Some writers have been careful in observing that the performance hasn't been for want of external shocks, including the Asian financial crisis of 1997-98, the industrialised world recession of 2001 and the relatively recent impact of the re-emergence of China and India.

In addressing this group last year I spent quite a lot of time exploring the implications for the Australian economy of the latter.

Chart 1: Terms-of-trade

The most obvious transmission mechanism is the terms-of-trade, which, because of spectacular growth in commodity prices, are now at their highest level in the post-Korean War period. Through its terms-of-trade, Australia has been affected by China's and India's re-emergence more than any other OECD country. Since 2001, Australia's terms-of-trade have risen by nearly 37 per cent. The second largest rise over this period – experienced by Norway – was 33 per cent, followed by Canada with a 13 per cent rise. For most OECD countries, however, the terms-of-trade have

---

[1] As usual, I want to thank a number of Treasury colleagues for their help in preparing this address, especially David Gruen, Steven Kennedy and Jim Thomson.

hardly moved over the past 5 years, with the exception of Japan where they have fallen by 18 per cent.[2]

Our forecasts have the terms-of-trade coming off modestly in 2007-08. And, as usual, we have taken out a bit of insurance in the first two years of the projection period.

While real GDP growth has been running below the levels of the second half of the 1990s, employment growth has been relatively strong. This implies that measured labour productivity growth has not been as strong. Both things – relatively strong employment growth and less strong labour productivity growth – probably have something to do with the strength of commodity prices: apparently, there are very long output lags on investment in new mining capacity, and the higher terms-of-trade that have stimulated that investment may also be generating a structural decline in labour productivity for the reason set out in my address to you last year. Measured labour productivity would also have been affected by the drought. The impact on labour productivity growth of investment lags and the drought should be transitory.

As a consequence of continuing strong employment growth, Australian labour is as fully utilised now as it has been at any time in the past 30 years. The official unemployment rate is at 32-year lows. Some will say that the unemployment rate doesn't provide a complete picture, especially because the ranks of the employed include people who work very few hours. On the other hand, as the participation rate rises, a higher proportion of the working age population is employed even if the unemployment rate is unchanged. If only for these reasons, a broader measure of labour utilisation will be of interest.

Chart 2: FTE employment to working age population ratio

It is possible to construct a measure of 'full-time equivalent' employment, and to express this as a proportion of the working age population; that is, the population aged 15 or more years. In 1979, this ratio stood at 51½ per cent. In the recession of the early 1980s it fell to less than 49 per cent, and in the recession of the early 1990s to about 47½ per cent. In 2006 it was up at 51 per cent. This is only a little less than the immediate pre-recession levels of 1980-81 and 1989-90.[3]

But even this measure is affected by demographic change. Older workers have lower labour force participation rates than younger workers. And on average they work fewer hours. For these reasons, we might want to see how this full-time equivalent employment ratio would behave were we to hold the working age population distribution constant.

This adjusted measure of full-time equivalent employment is about as broad-based a measure of the strength of employment as we can come up with. On this measure,

---

[2] Over the same period, China's terms-of-trade have fallen, but only by about 3 per cent. To ensure comparability, all these terms-of-trade changes are calculated using the latest OECD data.

[3] Levels were even higher in the 1960s and early 1970s, notwithstanding significantly lower rates of female participation.

you would have to say that the labour market is as close to full employment now as it has been at any time in the past 30 years; indeed, it is closer.

Very importantly, we have achieved this employment level without wide-spread wage inflation. No doubt, that has a lot to do with the credibility of monetary and fiscal policy frameworks, and a couple of decades of pro-competitive microeconomic reform – including in product markets, capital markets and, significantly, the labour market.

Achieving a state of something close to full employment is a staggering achievement. But, as I have said on other occasions, it has quite significant implications for policy discipline: not only is there an ongoing need to conduct fiscal and monetary policies in a way that satisfies the requirements of their medium-term frameworks, there is also a need to avoid policy interventions that don't add to supply capacity since these can be expected to be detrimental to productivity and GDP per person growth.

The gap between the unadjusted line and the line that holds demography constant has been growing over time, and will grow further as the population continues to age.

Chart 3: Demographic dynamics

The gap reflects two age-related effects. First, retirees are making up an increasing proportion of the population aged 15-plus; and second, the age composition of the population aged 15 - 64 has also been changing. The first effect has been negative for labour utilisation, and increasingly so. The second effect has had a positive impact on labour utilisation, roughly off-setting the first through to the late 1990s. Since then, however, that positive effect has been weakening, and is set to weaken further, eventually becoming negative.

The proportion of the population of traditional working age, 15 – 64, is at about its peak.

The total dependency ratio – the sum of those aged 0 – 14 and those aged 65-plus divided by the population aged 15 – 64 – is at about its minimum.

This brings me to a consideration of longer run economic developments, fiscal and macroeconomic.

Chart 4: The 3Ps of growth in real GDP per person

IGR2, published last month, presented new figures explaining the sources of GDP per person growth over the past 40 years, and provided projections for the various components of GDP per person growth over the next 40 years. GDP per person growth is set to fall from an annual average of 2.1 per cent to 1.6 per cent. This growth slowdown is due to population ageing. In the past 40 years, due to a collapse in the birth rate, the share of the population aged 15-plus increased strongly. This change in the shape of the population distribution explains 0.4 of the 2.1 percentage point annual growth in GDP per person. In the next 40 years, this population ageing effect will continue, but at a slower pace, adding only 0.1 percentage points a year to GDP per person growth.

Over the past 40 years, average hours worked per person aged 15-plus fell, subtracting 0.1 percentage points from annual GDP per person growth. This was notwithstanding a strong increase in labour force participation rates. The big change here was in average hours per worker, due in part to the stronger growth in part-time employment: average weekly hours per worker fell steadily in the 1960s and 1970s, from 38½ hours in 1964-65 to around 36 hours in the mid-1980s, falling to about 35 hours in the early part of this decade.

Looking ahead, workforce participation is set to fall dramatically, as a consequence of population ageing, cutting 0.3 percentage points a year from GDP per person growth.

Combining the population and participation elements of the '3Ps', we have a broad measure of labour utilisation – specifically, average hours worked per head of population. In the past 40 years, higher labour utilisation contributed 0.3 per cent a year to real GDP per person growth; in the next 40 years lower labour utilisation will subtract 0.2 per cent a year.

Labour productivity growth averaged 1.8 per cent a year in the past 40 years, and we are projecting the same average annual rate of growth over the next 40 years. The projection is the same as that made in IGR1 published in 2002. But the history is quite different. In the 40 years preceding IGR1, labour productivity growth averaged 2.0 per cent a year. Given that history, many readers of IGR1 thought our productivity projection a little conservative. We don't hear the same sort of comment today. My view hasn't changed. I continue to consider 1.8 per cent a year to be an appropriate figure on which to base 40 years projections of productivity growth.

Demographic parameters are no more static than the more familiar macroeconomic parameters that we use in our medium-term economic and fiscal projections. Since IGR1, mortality rates have fallen and life expectancies have increased more rapidly than we expected. The fertility rate has increased. Migrant numbers have increased. And there has been an interesting change in the age distribution of the migrant intake, due principally to the increasing importance of the skilled categories.

Chart 5: Population projections: Percentage change from IGR1 to IGR2

In IGR1, we projected that by 2042 the Australian population would be 25.3 million. We now think it will be 10 per cent larger, at 27.8 million. Less than 1 of the 10 percentage points difference is due to faster-than-projected population growth over the past 5 years. Lower mortality rates explain about 2 percentage points. That is, the faster-than-projected fall in mortality rates over the past 5 years, and the consequent falls in projected future mortality rates, imply an additional half a million Australians in 2042. Higher fertility rates explain 3 percentage points: about three-quarters of a million more Australians in 2042. And higher migrant numbers explain 4 percentage points, or about a million additional Australians in 2042. The change in composition of the migrant intake doesn't have much of an impact on the total population.

These changes in demographic parameters not only affect the total size of the population, they also affect its shape and, principally because labour force participation rates vary by age cohort, also have implications for the level of real GDP per person.

Chart 6: Demographic contributions to real GDP per person

We reckon that the new demographic parameters will reduce GDP per person in 2041-42 by a little less than 1 per cent relative to the IGR1 projections. 20 years from now the reduction will be even larger.

Lower mortality rates, would, on their own, reduce real GDP per person by about 1⅔ per cent. And higher fertility would reduce it by an additional one-half of a per cent – considerably more than that a generation from now. But more than half of this negative impact on real GDP per person will be offset by developments in the immigration area. Higher migrant numbers are expected to increase real GDP per person by about ¾ of a per cent, and the change in the age composition of the migrant intake will add about one-half of a per cent to real per person GDP.

I must admit that, when I first saw this chart, I found it hard to accept that the change in the age structure of the migrant intake in the past 5 years was of such significance.

Labour utilisation and labour productivity explain real GDP per person. In fact, real GDP per person is simply obtained by multiplying these two things: average hours worked per person multiplied by real GDP per hour of work provides GDP per person.

Chart 7: Productivity and labour utilisation

If we construct a chart with labour utilisation on the horizontal axis and labour productivity on the vertical axis, we can chart the course of our economic progress through a map of equal GDP per person contours.

In 1978-79 Australians worked, on average, 15.4 hours a week. If that seems a bit low to you, remember that this includes babies as well as people aged more than 100. Measured in 2006-07 dollars, an average hour of work generated a bit more than $36 of GDP. Hence, in 1978-79 real GDP per person was 15.4 hours a week of labour utilisation multiplied by $36 of labour productivity, or about $560 a week.

By 2005-06, real GDP per person had risen to about $940 a week. Labour utilisation had increased by about 1½ hours a week, or 10 per cent, to 16.9. And labour productivity had grown by some 53 per cent, to be close to $55 an hour.

The path of economic development from the late 1970s to the mid years of this decade wasn't smooth. The recessions of the early 1980s and early 1990s crunched labour utilisation, and stalled productivity. The past decade stands out as relatively benign.

The right-ward drift in history is a product of complex interactions among several demographic and economic parameters: the proportion of the population of working age; age- and gender-specific participation rates; the unemployment rate; and average working hours. We reckon that as much as two-thirds of the trend right-wards movement is due to demography – illustrating the beneficial impact of population ageing in our history.

Before moving on, a quick digression on this concept of labour utilisation might be in order.

Labour utilisation isn't a measure that many of you would use very often. And you might, therefore, not have much of a feel for the significance of the difference between 15.4 hours a week and 16.9 hours a week. It probably doesn't sound like much. But, in fact, it is a very big deal.

A more familiar concept is the employment ratio, representing the proportion of the population aged 15 or more years that has a job. Obviously, the employment ratio is simply the product of the conventional total participation rate and one minus the unemployment rate. 2005-06 labour utilisation of 16.9 hours a week corresponded to an employment ratio of 60.6 per cent. If labour utilisation were to fall from its 2005-06 level of 16.9 hours a week back to its 1978-79 level of 15.4 hours a week solely because of a reduction in the employment ratio – that is, with no change in the shape of the population, and so on – the employment ratio would have to fall from 60.6 per cent to 55.2 per cent. As a rough rule-of-thumb, each one-tenth of an hour per week of labour utilisation equates to a change in the employment ratio of about one-third of one percentage point. [4]

We can transpose the horizontal axis in 'employment ratio equivalent' terms by finding the employment ratio that would produce the same level of labour utilisation – and therefore the same level of GDP – if everything else (demography and average employee working hours) were unchanged. Thus, labour utilisation of 17 hours a week corresponds to an employment ratio of 60.9 per cent, and labour utilisation of 14.5 hours a week corresponds to an employment ratio of 51.9 per cent.

Labour utilisation is sensitive to labour market and tax-transfer policy settings. The Treasury's participation modelling project has the capability to assess the impact on labour supply (or potential labour utilisation) of tax-transfer policy changes in particular. We have run last week's tax cuts through our version of the Melbourne Institute Tax and Transfer Simulator (MITTS). Most of the positive impact on labour supply comes from the increase in the 30 per cent threshold from $25,000 to $30,000 – including the labour supply response of many secondary earners. The increases in LITO and the 40 per cent and 45 per cent thresholds are also positive for labour supply, though smaller. And the increase in the dependent spouse rebate is estimated to produce a very small negative impact on labour supply.

Overall, we calculate that the Budget tax cuts might increase labour supply by about 0.1 hours per week. If this additional supply is fully employed, the increase in labour utilisation will lift the employment ratio by about a third of a percentage point.

But there is another way of looking at this. The increase in labour supply expands potential GDP in the same way as a cut in the NAIRU does. As a rough approximation, an increase in labour supply of 0.1 hours a week would have about the same impact on potential GDP as a cut in the NAIRU of half of a percentage point – quite a significant amount.

The package of tax cuts announced in last year's budget, and in particular the increase in the 30 per cent threshold from $21,600 to $25,000, would have added about

---

[4] See Appendix for more detail.

another 0.1 hours a week to labour supply. Thus, over two budgets, the increase in this threshold could have an impact on potential GDP broadly equivalent to a one percentage point cut in the NAIRU.

Changes to child care arrangements announced in last week's budget can also be expected to make a positive contribution to labour supply. Coming on top of the earlier welfare-to-work measures and the superannuation changes announced in last year's budget, we have had quite a significant policy-induced boost to the economy's supply capacity.

Chart 8: Productivity and labour utilisation

Looking ahead, IGR1 projected that by 2041-42, labour utilisation would have fallen back to 15.7 hours a week. This is equivalent, in terms of its impact on GDP, to the employment ratio decreasing to 56.3 per cent – all else unchanged.

The IGR2 projections show that progress has been made in the past 5 years, with 2041-42 labour utilisation of more than 15.9 hours a week. That picture understates the amount of progress, however. As noted earlier, changes in key demographic parameters have had a negative impact on projected GDP per person. The negative impact comes from lower projected labour utilisation. Were it not for strong growth in labour force participation rates, especially among older workers, over the past 5 years, more than off-setting the changes that have occurred in demographic parameters, IGR2 would have projected 2041-42 labour utilisation of less than 15.6 hours a week, and by 2046-47 IGR2 would have had labour utilisation back below 1978-79 levels.

This illustrates the importance of policy measures to enhance labour force participation. We can calculate that the increase in labour force participation rates between IGR1 and IGR2 was, in terms of its impact on labour utilisation and GDP, as significant as an increase in the employment ratio of about 1.3 percentage points in 2041-42; and it produces a gain in 2041-42 real GDP per person of $38 a week in today's prices.

Those developments are important for their own sake, but they are also important in making an assessment of future fiscal sustainability.

In the IGR, fiscal sustainability is assessed by reference to the movement over time in the budget balance, expressed as a proportion of GDP. Since we hold the revenue-to-GDP ratio constant in making our IGR projections, we are really looking at the behaviour over time of the ratio of government spending to GDP. People who take an interest in 'size of government' questions tend also to look at the behaviour of this ratio over time.

The ratio of government spending to GDP can be expressed as the quotient of two other ratios: government spending per person, and GDP per person. The second of these is a familiar concept. And the first is also an intuitively appealing measure. We know from earlier charts that real GDP per person is affected by population ageing; and so too is real government spending per person – in some quite complex ways. As the population ages, real government spending per person increases because of the age pension and other age-related payments, and falls as payments in respect of

children grow less quickly. But government spending per person is also affected by discretionary budgetary decision-making.

Chart 9: Government spending and GDP

We can construct a chart with real GDP per person on the horizontal axis and real government spending per person on the vertical axis. Provided we use the same deflator for both axes – and it makes sense to use the GDP deflator – we can draw lines (rays from the origin) representing constant ratios of government spending to GDP. The measures of real government spending per person, and the ratio of government spending to GDP, used here exclude interest payments.

In 1972-73 and 1973-74, the ratio of government spending to GDP was a little above 17 per cent. It's been nowhere near that level ever since – the closest being in 1989-90. Its highest level in this 35 year period was in 1984-85, when it hit 25 per cent. The chart illustrates the impact of the recessions of the early 1980s and early 1990s, with real GDP per person falling and real spending per person increasing strongly, partly because of the automatic stabilisers and partly because of discretionary fiscal expansion.

The very large increase in real spending per person associated with the Whitlam Government stands out quite sharply in the chart, as does the 'fiscal consolidation' of the second half of the 1980s.

The figures in the chart are raw Budget numbers excluding interest payments. They have not been adjusted in any way for anything to do with the introduction of *The New Tax System*. [5] Nor have they been adjusted for the downward pressure on the government spending to GDP ratio that is normally observed when the GDP deflator runs well above trend. As to that, over the period 2000-01 to 2005-06 the terms-of-trade boosted nominal GDP by about 5 per cent.

Chart 10: Government spending and GDP

According to IGR2, over the next 40 years, the ratio of government spending to GDP will increase by 4.7 percentage points, from 20.8 per cent to 25.5 per cent, taking the ratio to levels not seen in the past 40 years.[6]

Over the next 40 years we have projected that real GDP per person will grow by about 86 per cent, while real government spending per person will grow by about 128 per cent. Obviously, if the growth in real government spending per person could be restricted to 86 per cent, there would be no increase in the government spending to GDP ratio. There would be no fiscal challenge from population ageing.

---

[5] There is also a structural break in the series between 1998-99 and 1999-2000 due to methodological and data source changes associated with the move to an accrual accounting framework.

[6] In IGR1, the ratio was projected to increase by 4.1 percentage points, from 23.3 per cent to 27.4 per cent: between IGR1 and IGR2, the starting point improved by 2.5 per cent of GDP, about 0.7 percentage points of which relates to a relatively minor methodological change. However, if we take the 35 years common to both sets of projections – 2006-07 to 2041-42 – we find that IGR1 projected a deterioration of 5.8 per cent of GDP, while IGR2 projects a deterioration of 4.1 per cent.

But that might not be the most sensible message to take from the chart. An alternative message would be that we could prevent the government spending to GDP ratio from rising if we were able to secure faster GDP per person growth.[7]

Chart 11: Growth of real GDP and government spending per person

As noted earlier, we are projecting labour productivity growth of 1.8 per cent a year, the same as the average of the past 40 years. Over a 40 year period, that sort of productivity growth compounds up to an increase in real GDP per person of a bit more than 100 per cent. In the past 40 years, we also got a positive contribution from labour utilisation, so that real GDP per person increased by some 128 per cent. But, over the next 40 years, with labour utilisation falling, real GDP per person is projected to increase by only 86 per cent. What this chart shows – quite clearly – is that if it were not for the slowing in real GDP per person growth due to population ageing, the ratio of government spending to GDP would not be projected to rise over the next 40 years: coincidentally, government spending per person would be growing at the same rate as GDP per person.

Concluding remarks

None of the textbooks of the second half of the 1970s described an economy characterised by low and apparently well-anchored consumer price inflation, low rates of unemployment, relatively low and stable interest rates, habitual budget surpluses and negative public sector net debt. An economy in which public policy debate would be pre-occupied with the nation's supply constraints; where macroeconomists would be as obsessed as their microeconomic colleagues with the efficiency of resource allocation and concepts of opportunity cost.

Yet this is the economy that the most ambitious of the policy advisers of the last 30 years wanted to see.

Much remains to be done, of course – a point underscored by the looming demographic challenge, and by the persistence of social and economic disengagement in significant areas of disadvantage, most notably among indigenous Australians. But I have a sense that the parameters of the debate – in the domains of population, participation and productivity – are, these days, well understood.

---

[7] Faster real GDP per person growth would be expected to raise government spending per person somewhat, via its effect on real wages and thereby government benefits linked to wages. Notwithstanding this link, we should expect stronger real GDP per person growth to reduce the ratio of government spending to GDP, relative to what it would otherwise have been.

Chart 1

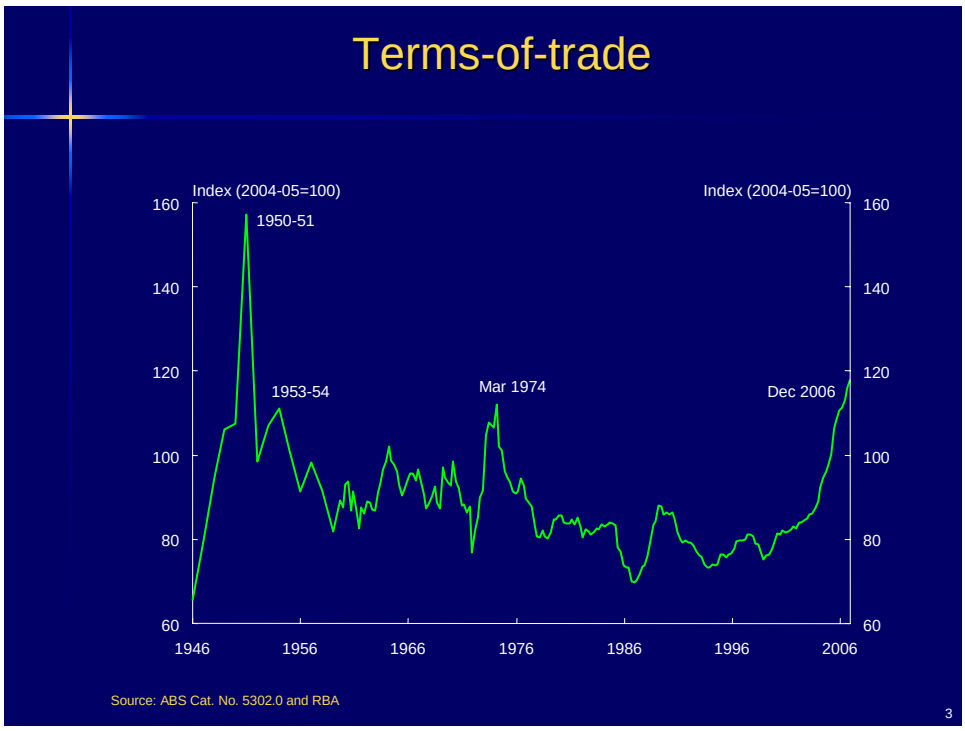

Chart 2

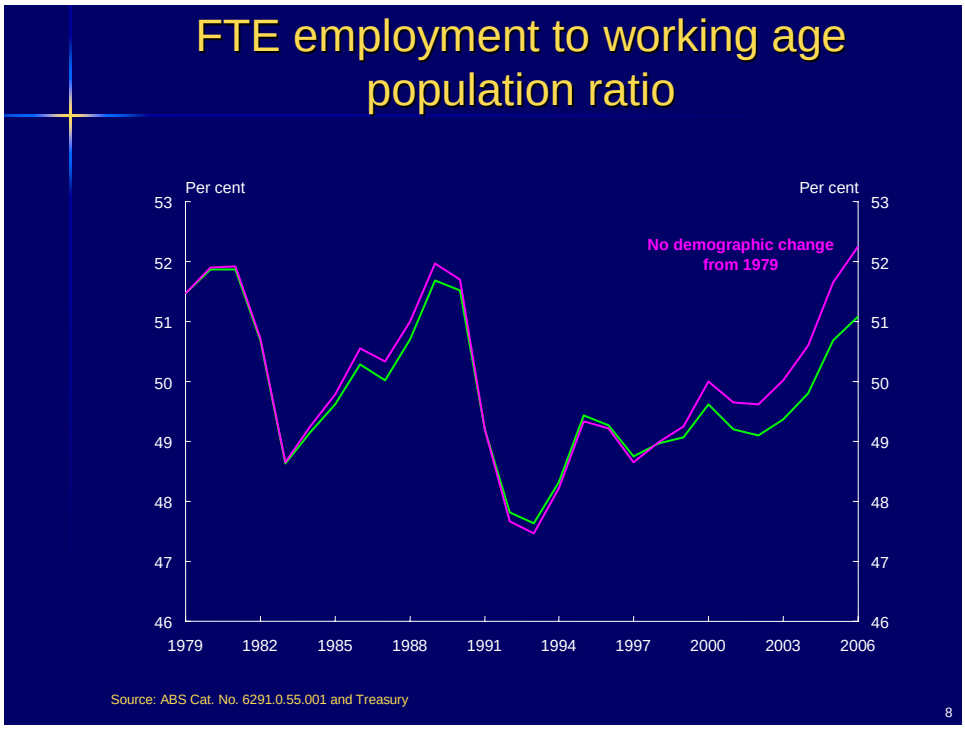

Chart 3

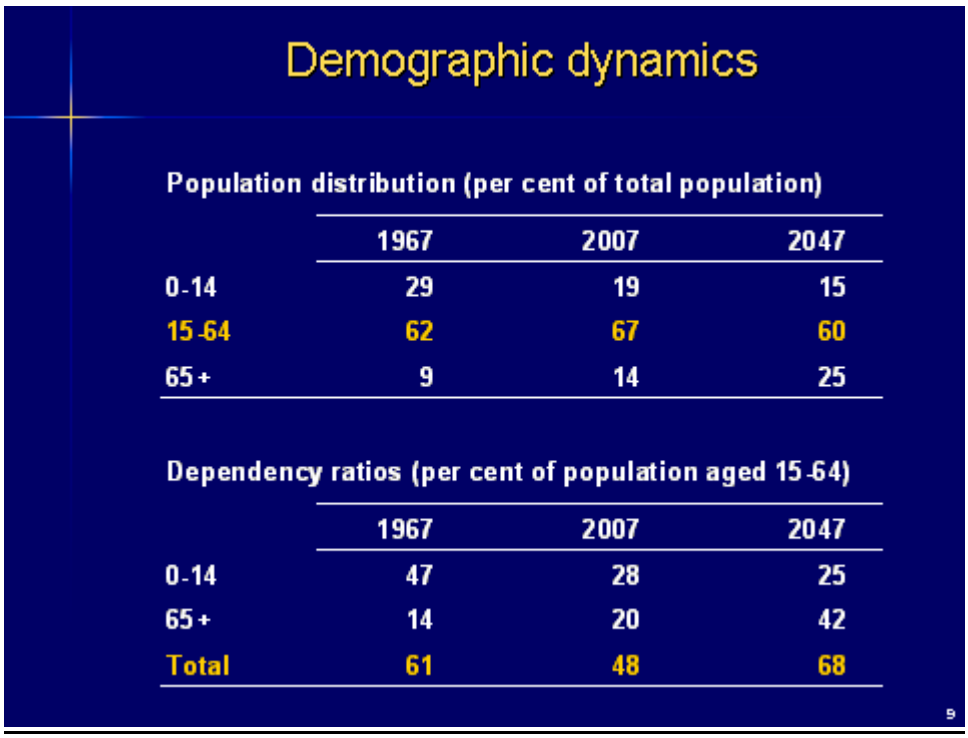

Chart 4

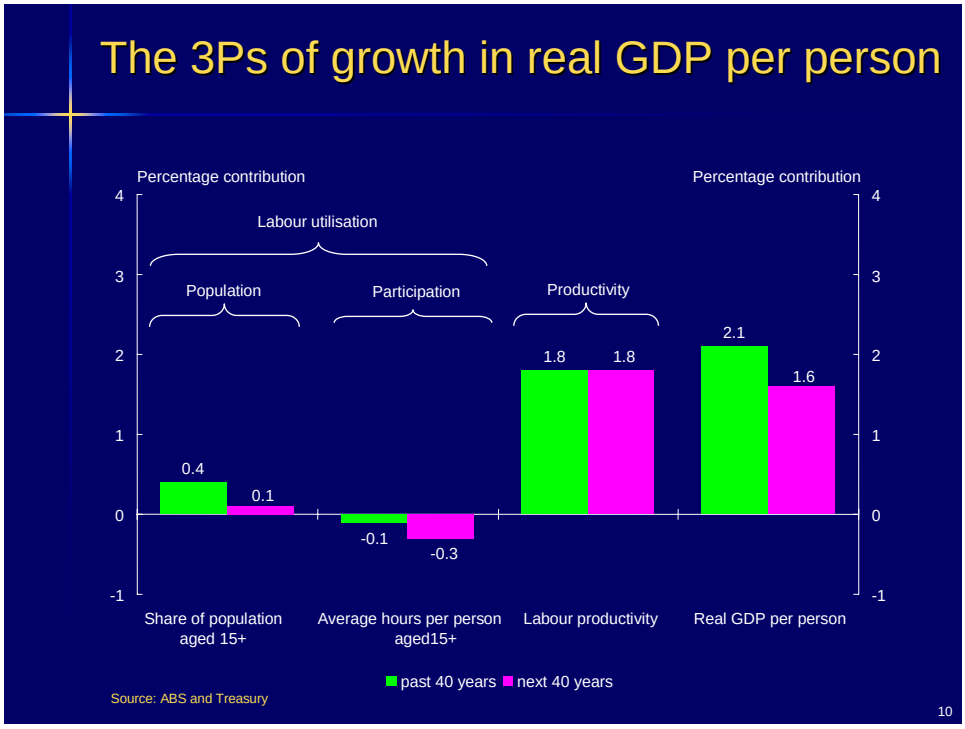

Chart 5

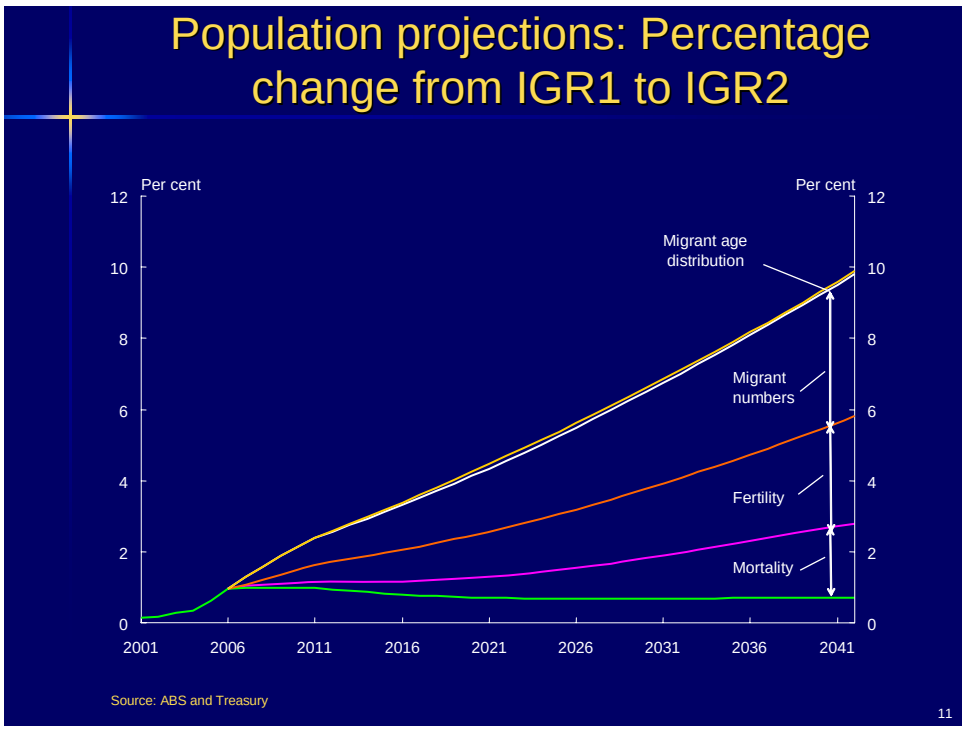

Chart 6

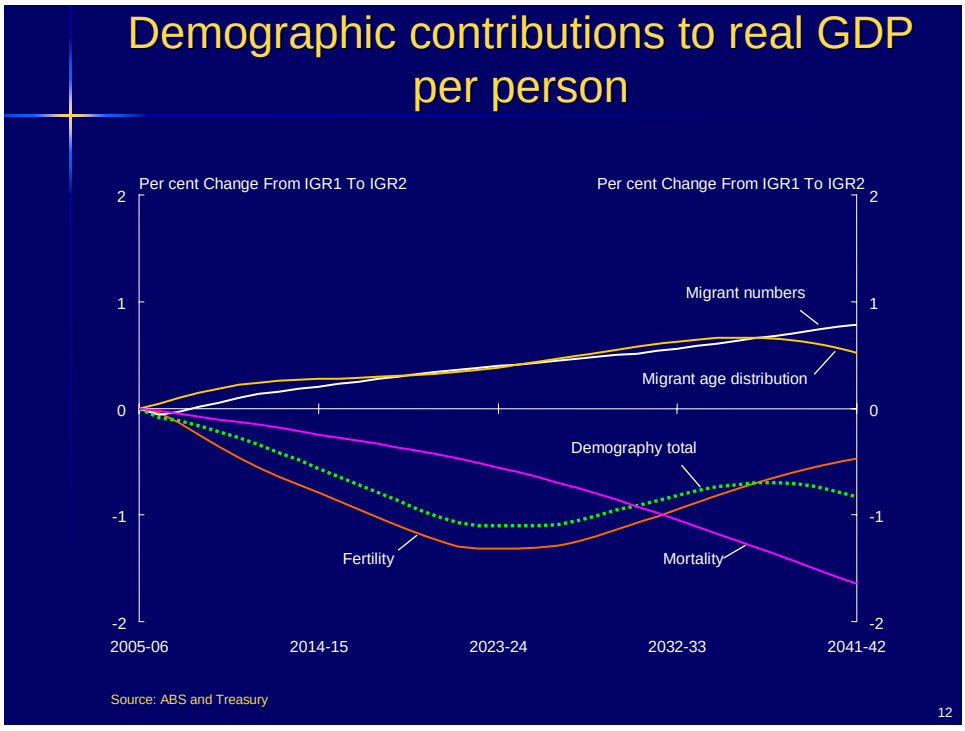

Chart 7

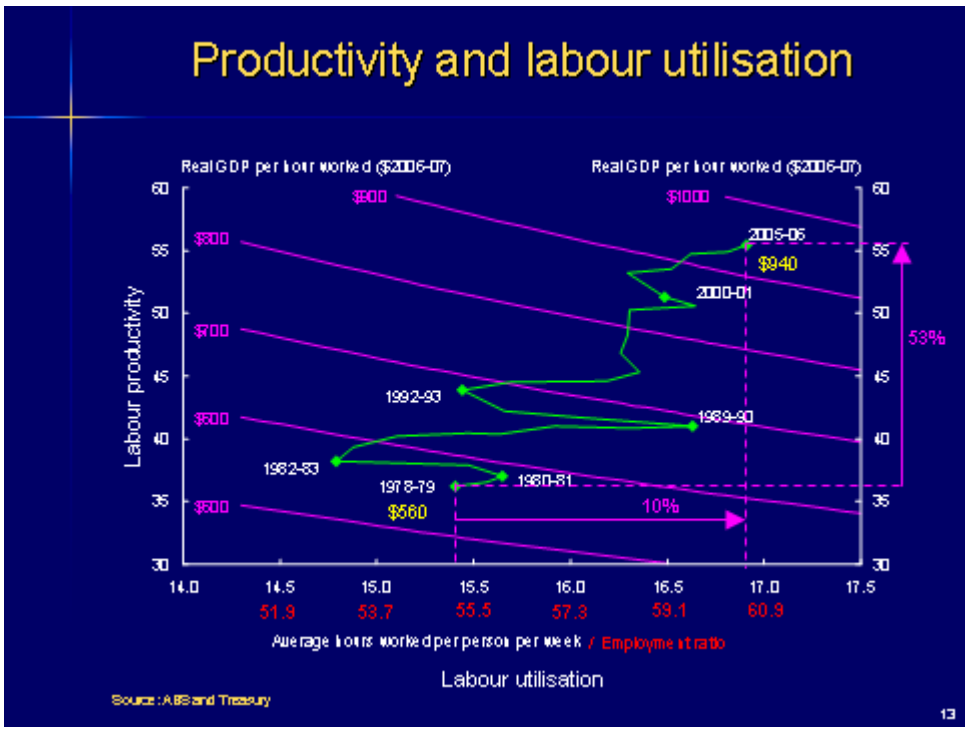

Chart 8

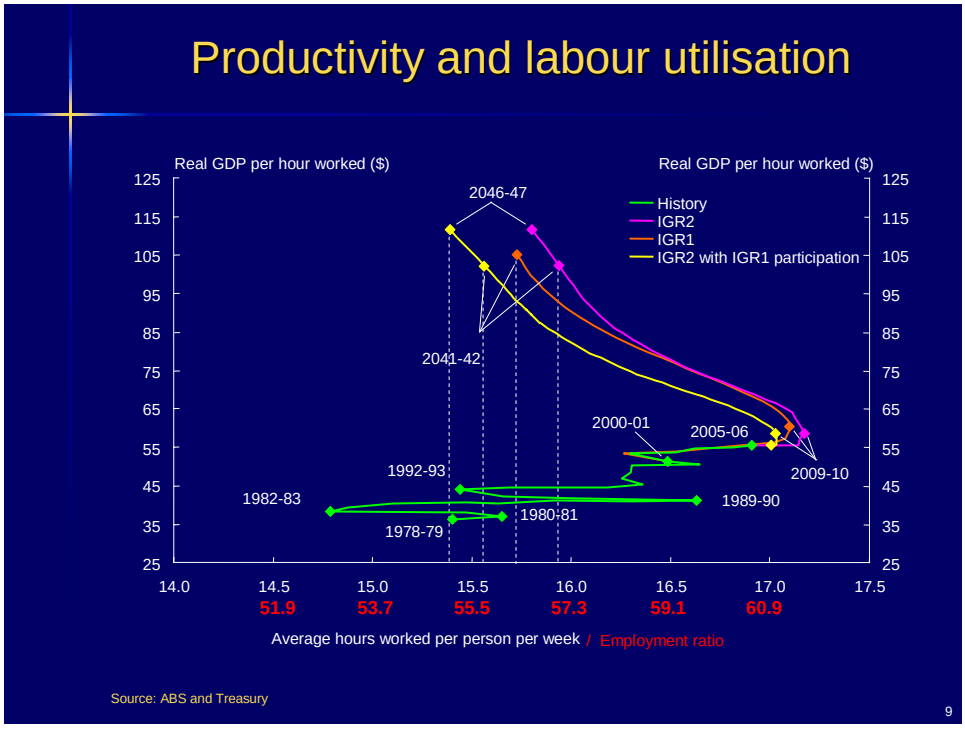

Chart 9

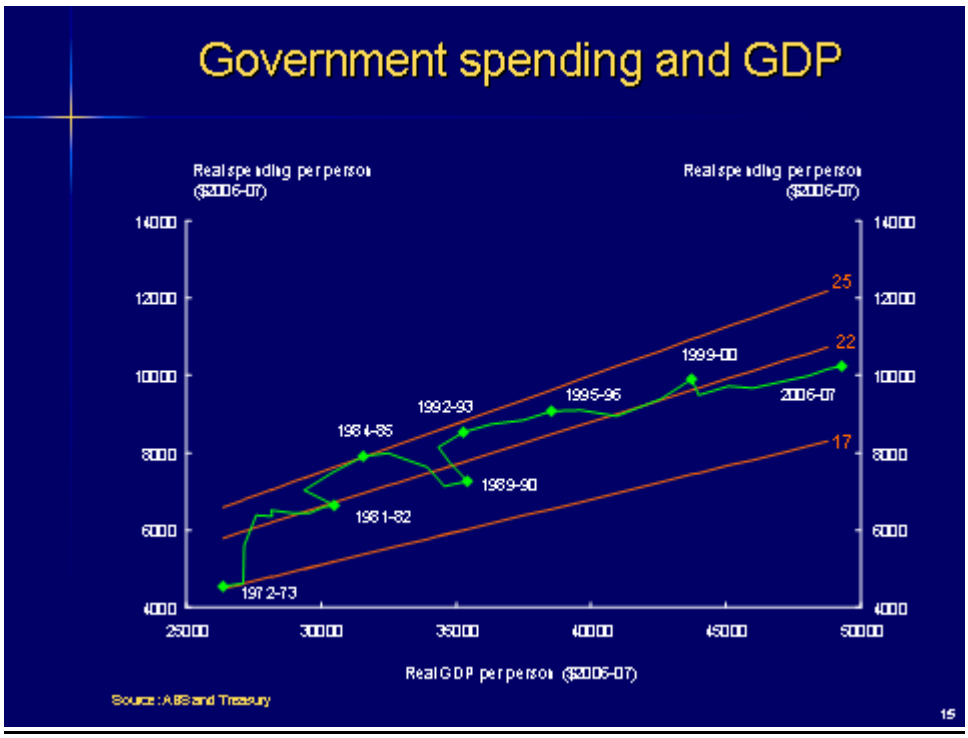

Chart 10

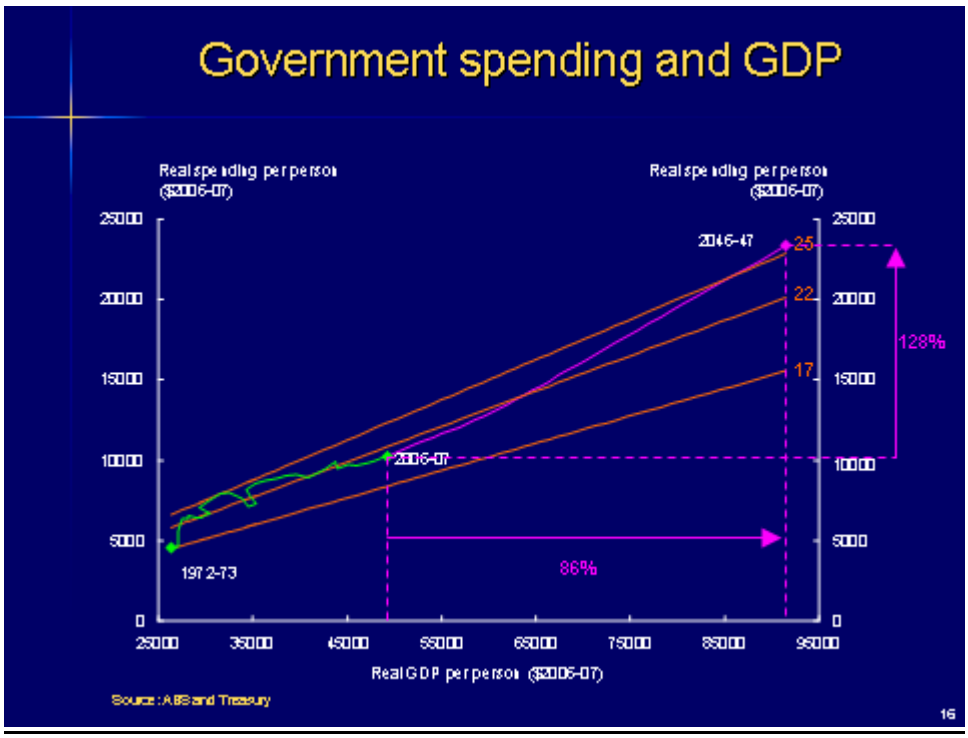

Chart 11

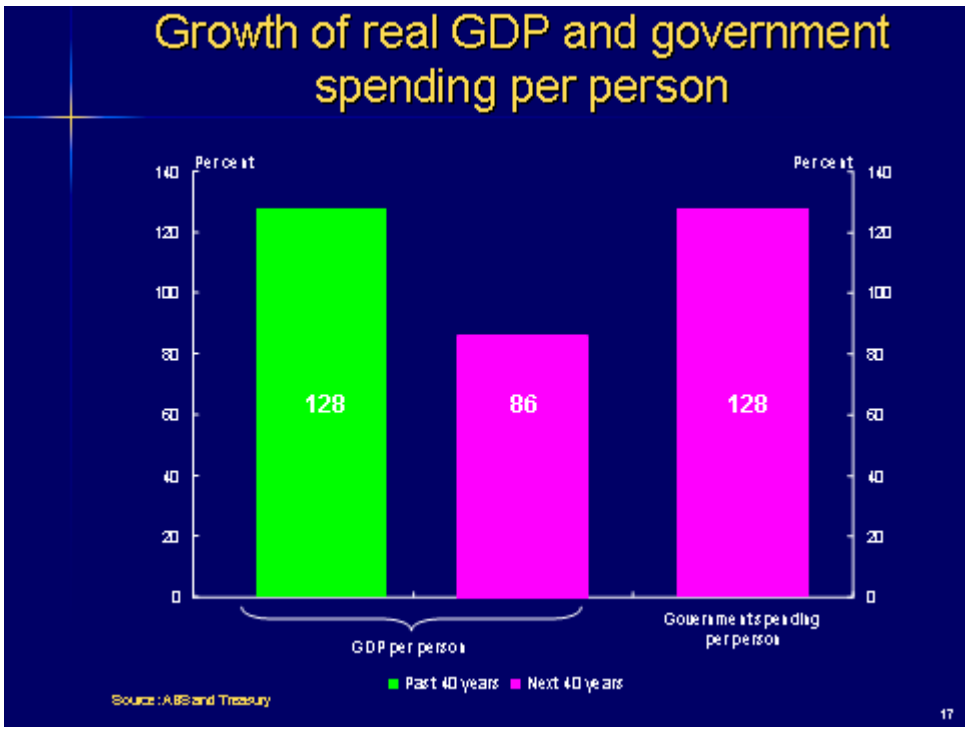
Growth of real GDP and government spending per person
Percent
140
120
100
80
60
40
20
0
128
86
128
GDP per person
Government spending per person
Past 40 years
Next 40 years
Source: ABS and Treasury
17

Appendix

Labour utilisation is a measure of hours worked (in a week) per head of the total population. This may be written algebraically as follows:

(1) $$L = \alpha\rho(1-u)h,$$

where:

α	The proportion of the total population of working age (15-plus).
ρ	The average workforce participation rate.
u	The unemployment rate.
h	Average working hours (i.e., average hours worked in a week by those who have a job).

The ratio of employment to the working-age population, *e*, is the product of the participation rate and the proportion of the labour force that is employed: $e = \rho(1-u)$. Substituting this into (1) and solving for the employment ratio, we have:

(2) $$e = \frac{L}{\alpha h}.$$

In 2005-06, variables of interest had the following values:[8]

α	0.807
e	0.606
h	34.6

Thus, in 2005-06, labour utilisation was:

$$L = 0.807 \times 0.606 \times 34.6 = 16.9.$$

Using (2), we can obtain an expression for the employment ratio that would generate any level of labour utilisation given unchanging values for the other variables:

$$e = \frac{L}{27.92}.$$

---

[8] These numbers are taken from the model-consistent IGR data base, and will differ slightly from other published numbers.

An increase in labour utilisation of 0.1 hours per week, obtained by changes in α or h, could equivalently be obtained by an increase in the employment ratio with those other variables constant. The equivalent change in the employment ratio is:

$$\Delta e = \frac{\Delta L}{27.92} = \frac{0.1}{27.92} = 0.0036 = 0.36\% . \tag{3}$$