

The long-term level “misalignment” of the exchange rate: Some perspectives on causes and consequences

Michael Reddell¹

Speaking notes for presentation at Reserve Bank/Treasury exchange rate forum

26 March 2013

The papers presented this morning focused mainly on the volatility in the exchange rate. My paper takes a different tack. I argue that even if one were totally relaxed about those shorter-term fluctuations in the exchange rate - and I'm not - we should be very concerned about where the exchange rate has been, on average, over the last few decades.

I organise the paper around what I call “three failures”, and use them as the basis for developing a plausible story that explains as many as possible of the uncontested things we know about New Zealand's macroeconomy in recent decades. Much of it is exploratory in nature: joining the dots, but where the dots aren't already numbered. And if a 40 page paper already seems quite long, I am quite well aware that most of those pages could easily warrant 40 more pages to cover all the issues in adequate depth.

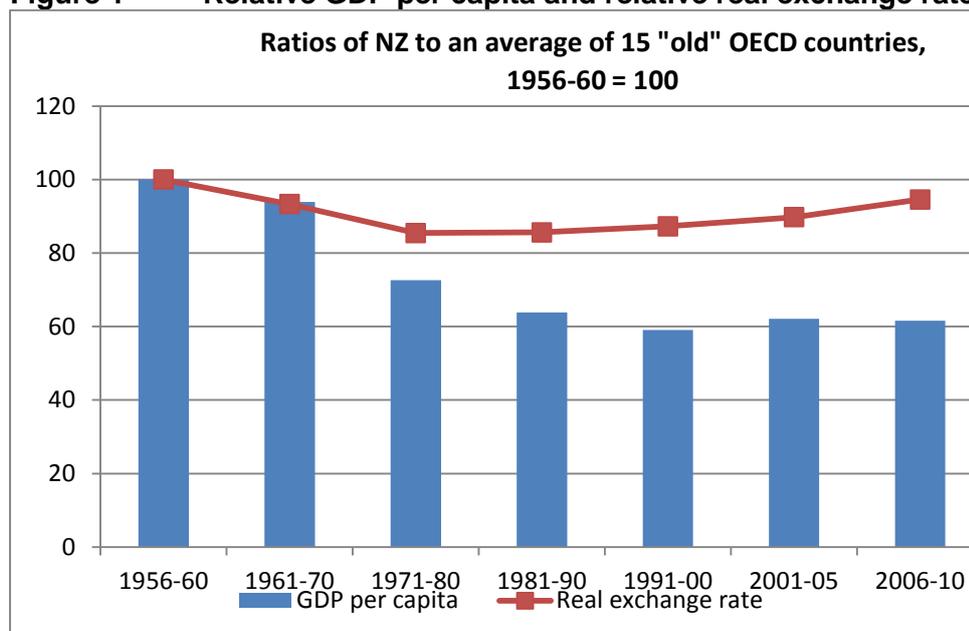
But back to the ‘three failures’.

The one that really matters is the failure to put into reverse the decades-long decline in New Zealand's relative economic performance. Far-reaching reforms were introduced in the mid-late 1980s and the early 1990s, which looked to have positioned us well for catch-up, but it just hasn't happened. We have more or less held our own, at least in per capita income terms, since around 1990 - but even that has been achieved only by working many more hours per capita (eg raising the NZS eligibility age). The productivity performance - what we get per unit of what we put in - has been (bluntly) lousy, and among the worst in the OECD.

My first chart focuses on that real economic decline. I go back a long way, and I use decade or half-decade averages because my focus is on longer-term trends, not cycles. This measure of GDP (nominal GDP per capita, converted at purchasing power parity) captures the direct effect of the recovery in the terms of trade. And as I compare us with a fixed group of “old OECD” countries, the poor relative performance isn't simply reflecting either, say, the dramatic rise of countries like Korea, or the entry to the OECD later in the period of a bunch of Eastern European former communist countries or territories. We've done badly and should be doing better, but we aren't, and there hasn't yet been a compelling story about why.

¹ Reserve Bank of New Zealand. Views expressed are those of the author and should not be attributed to the Reserve Bank,

Figure 1 Relative GDP per capita and relative real exchange rates



Data note: Nominal GDP per capita, converted at PPP exchange rates, and BIS narrow effective exchange rate indices. Prior to 1963, estimated real exchange rates calculated using OECD nominal exchange rate and CPI inflation data.

Source: Penn World Tables, BIS, OECD, author's calculations

My second “failure” is the real exchange rate. There is a variety of different measures of the real exchange rate - this one, largely drawn from the BIS, has the advantage of going back a long way, but alternative measures don't look much different, or more encouraging (Geoff Bertram's, in his *New Oxford History of New Zealand* chapter, looks even worse).

Over almost 60 years, the average level of the real exchange rate hasn't moved very much. But the relative performance of our economy - the relative productivity of its people and firms - has deteriorated markedly. In the early decades, it looked as though our real exchange rate was, on average, also falling, moving in the right direction to match our deteriorating economic performance. Only later did the gap blow out. And if you think this gap looks stark, have a look at the version that uses productivity – the value of output per hour worked - at the back of my paper.

But why does this matter?

It is hard to reverse the effects of decades of underperformance. Doing so required a sustained and substantial period of outperformance - growing materially faster than our peers. Reforms to product and factor markets, and opening up the economy to greater domestic and international competition are a key pre-conditions for making that change. But in a small country in particular, we can't expect to generate the sort of strong sustained growth in output and productivity required just by selling to each other. International experience suggests that realising the gains from reform typically happens through - and is led by - increased sales to the rest of the world, and the investment and productivity that successfully exploiting those much greater opportunities brings.

That, in turn, typically seems to have involved a sustained period with a relatively low real exchange rate - it certainly did, for example, in Singapore, Japan, and Korea. In New Zealand, the export share of GDP (even taking account of the terms of trade) has barely changed in 30 years. In the long-run, we sell to others so that we can consume more ourselves. But starting from a point where export value-added is quite low for a country our size, we need a sustained period of rapid export growth (and probably a rising export share of GDP) to help generate sustainable rapid income and productivity growth. And the sort of sustained rapid export (or tradables) growth that catching up requires is most unlikely to occur while the exchange rate fluctuates around levels that are, on average, so far out of line with our current productivity performance.

In the long-run, if we do catch up, such a high exchange rate might well be appropriate - we see rising real exchange rates in many East European countries as they converge - but we can't put the cart before the house. Catching up with the other advanced countries requires getting a lot of things right, including removing any obstacles that have stopped the real exchange rate being better aligned with long-term fundamentals.

Long-term average exchange rates- don't just appear, or exist in a vacuum. They arise out of what else is going on. So correct diagnosis of the nature of the roadblocks matters.

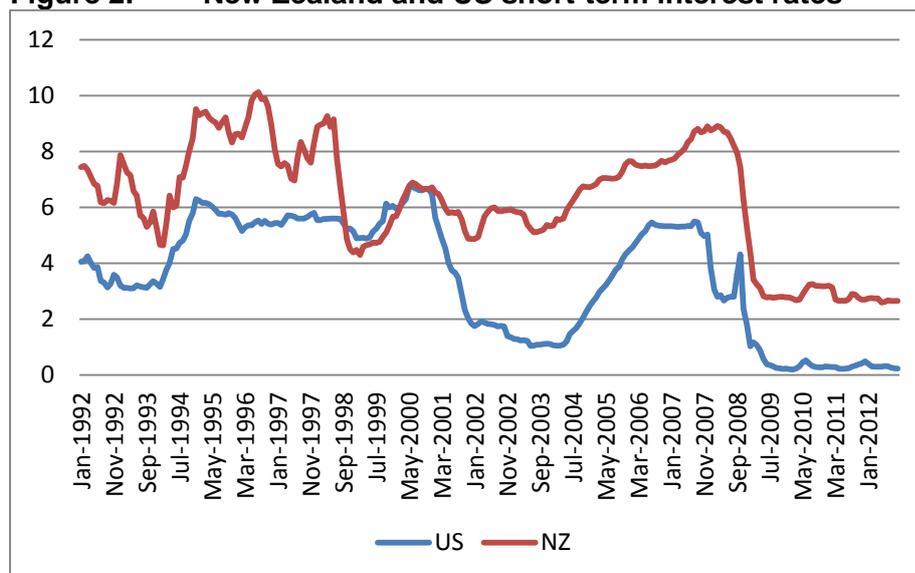
Which brings me to my third "failure" ; the failure of New Zealand interest rates to show any sign of settling down to the levels seen in other advanced countries - again, on average.

Historically, our interest rates weren't obviously higher than those in the rest of the world. Our real interest rates shot up markedly relative to those in other countries in the mid-late 1980s as we reoriented monetary policy to get on top of inflation. But the same thing had happened in, for example, the US and the UK a few years earlier. It was quite widely expected that once people were confident inflation was in check, and especially taking into account the huge improvement in the structural fiscal position, that New Zealand dollar real interest rates would drop back to around those in other advanced countries. We got inflation to normal international levels 20 years ago, but interest rates didn't converge. And there is still no sign of it sustainably happening.

In the paper, I devote quite a lot of space to discussing the slope of the interest rate yield curves. It is there mostly to illustrate one important point: it was not only central bankers and other official observers who thought that on average our interest rates would get closer to those in other countries. It was markets too. We have been wrong, and the private sector, represented in financial market prices, has been wrong too. But their views directly matter for the exchange rate.

Interest rate surprises tend to be reflected in the exchange rate. Why? Because money is pretty mobile and investors (New Zealand and foreign) have every incentive to eliminate apparent "free lunches". If New Zealand interest rates rise, the exchange rate also tends to rise pretty quickly to the point where the (implicitly) expected future depreciation in the exchange rate largely offsets the higher interest rates. The longer local interest rates are expected to stay up relative to those abroad, the more the exchange rate has to rise if this rough equalisation of expected returns is to hold. One can put all that in very technical language and dress it up with lots of algebra, but all it means is simply "apparent free lunches typically don't last long".

Figure 2: New Zealand and US short-term interest rates



Source: OECD

Even relative to the rest of the world, New Zealand's interest rates have fluctuated a lot. Relative to the United States, this simple chart from the paper illustrates it. But the focus here is more on the average gap over 20-25 years, which hasn't narrowed, even though it had repeatedly been expected to do so. Surprises of that sort - realising that, once again, the gap isn't closing - tend to take the exchange rate higher. Since, even in the last decade, there has been no sign of interest rates actually converging - even though they had been expected to do so - it shouldn't be much surprise that the average real exchange rate has been even higher than it was in the previous decade. Indeed, what might be most concerning is if markets resignedly concluded that the real risk-adjusted gap was going to be there indefinitely - on that view, there could be material further increases in the real exchange rate.

Interest rates also don't occur in a vacuum. Central banks set very short-term interest rates directly, but they do so in response to everything else going on in the economy, to deliver an inflation target. And long-term interest rates are heavily influenced by expected short-term interest rates. As is, I think, now well-recognised, there is nothing unusual about New Zealand's inflation target itself.

That suggests we need to look primarily at what it is in the demand (and supply) pressures in New Zealand itself that mean that so much higher average interest rates - a higher "neutral rate" if you like - have repeatedly been required to deliver much the same inflation outcomes as in the rest of the advanced world. What is it that officials and markets have missed?

In thinking about the issue, we often talk about savings/investment imbalances. That is just another way of saying that at any particular interest rate, the total amount people want to spend (consumption and investment spending) creates a demand for resources (especially labour) that exceeds the resources that are readily available. If so, an increase in the interest rate (and in the real exchange rate) is a rationing device, to help reconcile all our conflicting desires.

So the fact that our real interest rates have needed to be, on average, persistently higher than those abroad - “needed”, remember, to achieve much the same sort of inflation target as other countries are using – suggests that the gap between what we want to spend and what is available has been larger here, most of the time, than the gaps in most other countries. One way that the effects of that imbalance have shown up is in the persistent quite large current account deficits (and the high level of external debt despite a poor economic record). The other way the imbalances have shown up is in surprisingly high interest rates and, in response, a high real exchange rate. High interest rates deter some spending (especially investment spending that would otherwise have taken place). And the high exchange rate diverts some resources, which would have been used to produce exports, to meet domestic aspirations. The high exchange rate also tends to skew the business investment that does take place away from export-oriented industries to those meeting domestic demand.

Some have argued that our interest rates are high, not because of on-going incipient excess demand, but because of investor fears associated with the large stock of foreign debt New Zealand has run up over previous decades. I deal with this possibility more fully in an Appendix to my paper. For now, just note that the process by which such a risk premium would get into New Zealand interest rates would also be one that tended to **lower** the exchange rate, not raise it. The claim here is that persistently high average interest rates (and the real exchange rate consequences) stem from the choices, including the results of policy choices, that New Zealanders themselves have made.

The rest of the paper is about trying to get to identify the nature of those choices. Inevitably it is more exploratory.

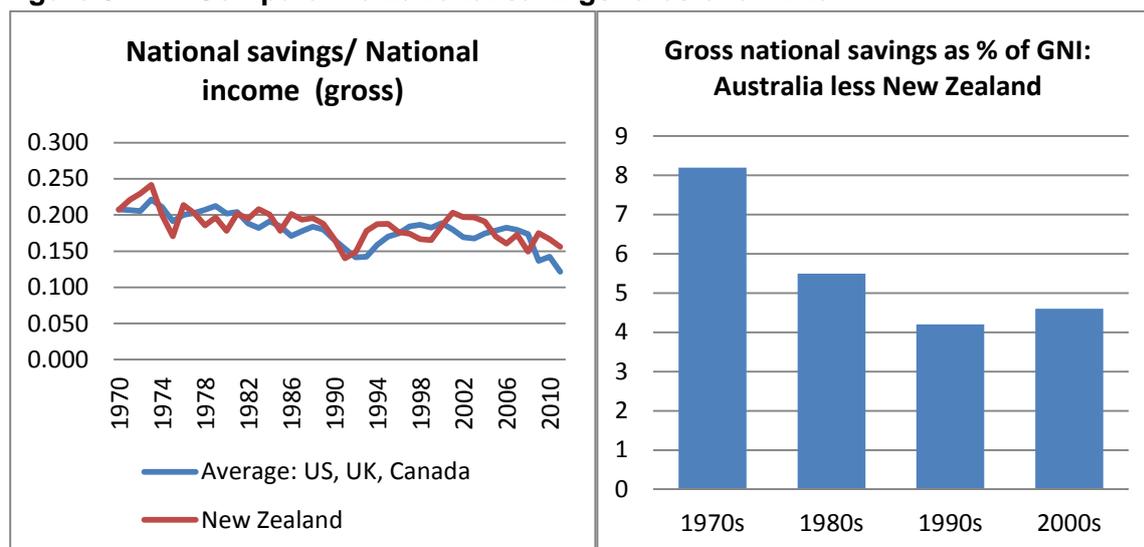
Of course, not all such apparent imbalances are problematic. A country with rapidly improving income prospects and high levels of business investment to take advantage of those prospects might also have seen imbalances: desired spending at a “world” interest rate running ahead of readily available resources.. That is, Australia’s story, needing to make room for the highly profitable part of the business sector that is expanding rapidly. It isn’t a picture that describes New Zealand.

In trying to understand what makes New Zealand different, the savings rate has been a popular place to look. So let’s do so.

For macroeconomic purposes, we need to focus mainly on national saving – the sum of the saving of New Zealand firms, households and governments. The national savings rate is relatively low in New Zealand – but has been for a long time. Good comparable cross-country data only go back to 1970: in the first half of the 70s we had the sixth lowest savings rate in the OECD, as we did in the second half of the last decade.

And although our savings rate is quite low relative to the OECD group as a whole, I was quite struck when I stumbled on this chart, showing how similar our savings rate has been over decades to a simple average of 3 other Anglo countries (Canada, US, UK). Of course, as everyone knows, Australia is different. But Australia’s national savings rate has been consistently higher than New Zealand (or the other Anglo countries) for decades - if anything, the gap has been narrower in recent decades than earlier, notwithstanding the compulsory private savings system introduced 20 years or so ago.

Figure 3: Comparative national savings rates over time



Source: OECD, Statistics NZ

In the paper, I work through a variety of stories that might explain our modest savings rate. Here I am going to address only three of the more prominent ones:

- The welfare and retirement income system
- The tax treatment of savings
- House prices

Welfare systems clearly tend to discourage private savings. But all OECD countries have welfare systems, and many of them - including those of high savings Northern European countries - are considerably more generous than New Zealand's. There is little reason to think that welfare spending discourages saving more here than in other countries, and some possibility that it is actually less discouraging (our unemployment and disability benefits are all flat rate, so provide limited income replacement for middle income people). Similarly, while our New Zealand Superannuation system provides effectively and universally against elderly poverty, New Zealand spends among the lowest shares of GDP on public pensions of any advanced economy. Our flat rate and non means-tested system leaves plenty of more need, and limited disincentive, for middle and higher income people to save privately.

The tax treatment of savings in New Zealand is unusual. We are the only OECD country to use the TTE system - funds going into retirement savings vehicles have already been taxed, and earnings on those funds are taxed, while final withdrawals - like withdrawals from a bank account - are tax-free. Deviations from TTE around the PIE system and Kiwisaver are fairly limited. Most other advanced economies provide the opportunity for substantial savings in vehicles where only a single T applies. I have argued elsewhere that there is a good case for a materially less onerous tax treatment of capital income generally. However, the literature cautions against expecting reforms to the tax treatment of savings to change in household savings in particular very much. Which instrument savings are held in is quite sensitive to changes in relative after-tax returns, but total savings is much less responsive. As perhaps just one illustration, the tax treatment of savings differs widely across New Zealand, Canada, the UK and the US - and has varied in each country quite markedly over time - and yet the national savings outcomes have been quite similar.

Many people believe that high and/or rising house prices have worsened New Zealand's national savings outcomes. I think there is little or no evidence to support that claim. Rising real house prices can pose financial stability risks, and they can and do lead to severe, probably adverse, distributional implications – unexpectedly and arbitrarily transferring wealth, for example, from the relatively young to the relatively old. But rising house prices don't make the country as a whole better off, or boost what collectively we can afford to consume in future. Whether it is wealth or borrowing constraints, one person's gain from rising house prices is another's loss. One of my favourite charts is a simple plot of consumption spending as a share of national income. Despite the biggest house price boom in modern New Zealand history in the last decade, there was simply nothing unusual in consumption behaviour. And although I'm not particular fond of using the household savings rate (as a share of household disposable income), it is estimated to have reached its low point - around -9% - right at the beginning of the house price boom (the year to March 2003), and then gradually rose for the rest of the decade.

If there was something obvious about policy frameworks that were clearly biasing savings rates down in New Zealand more than in other countries, there would be a good case for providing advice to take steps to rectify the distortion. But there is little sign that policy is unduly distortionary of savings behaviour.

What does that leave? Business saving tends to be a neglected part of any discussions about savings behaviour in New Zealand. Business saving is that portion of the income generated by the business that the owners choose to reinvest. Expected future returns are, of course, a key element in that calculation. Abstracting from the ups and downs of business cycles, in an economy with reasonable market disciplines, business savings are likely to reflect perceptions of future opportunities - if they are poor, business savings will be low too. So if business investment prospects in New Zealand were better - if, for example, the real exchange rate was cycling around a materially lower level - the average business savings rate would also, over time, probably be materially higher. But business savings probably isn't some autonomous drag on New Zealand's economic performance.

Finally, it is possible that New Zealanders - and their US, Canadian, and UK peers – just have a different set of cultural attitudes to savings and consumption than citizens in many Northern European or Asian countries do. Economists aren't usually comfortably invoking cultural differences as explanations - such differences often have their routes in something more "fundamental". But it is no more arbitrary or unreasonable to invoke cultural difference, having examined all other plausible candidates, than to simply rule out, by assumption, the possibility of such differences. Whatever the reason, savings rates do differ, and there is no compelling explanation, based in policy framework differences, that points directly to a case for government action designed to alter private savings. [Private preferences of citizens, expressed consistently over decades, should be respected, and factored into other government planning as required, not be lightly over-ridden or nudged in other directions under the thrall of models structured on the assumption that policymakers and economists can identify the best lifetime interests of adults better than those adults can themselves.]

None of this is to suggest that savings don't matter. If all New Zealanders woke tomorrow keen to save 20 per cent of their incomes rather than 15 per cent, outcomes would be materially different - and by most standards judged more favourable. But if firms and households are making considered choices over decades, and New Zealand has still ended

up with the anomalously high real interest and exchange rates, perhaps we need to check the other side of savings/investment balance picture.

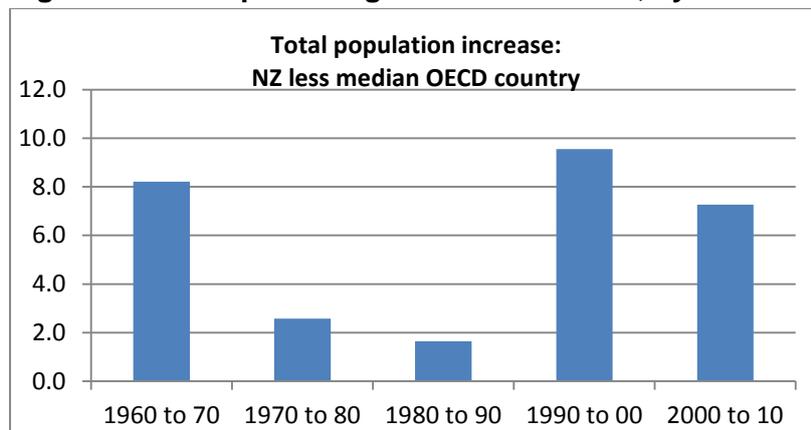
In some New Zealand discussion of macro imbalances there has been a tendency to downplay the investment side of things - an investment to GDP ratio around OECD averages might suggest nothing interesting. But, of course, what we observe is actual investment, at the real interest and exchange rates that have prevailed. But it is the gap between “desired investment” - what would occur at (say) a world interest rate, and the exchange rate that would result from such interest rates - and the investment we actually observe that should be our focus. The difference between those two is what has had to be “crowded out”.

The OECD did the New Zealand debate a service in their 2011 survey by including a nice chart breaking down investment to GDP into its institutional sectoral components. That set out starkly how low business investment had been in New Zealand as a share of GDP over the last 20 years or so. Business investment is typically how many of the productivity gains that drive future prosperity are achieved – supporting new products, new markets, and (not least) enabling smarter better combinations of labour and capital. But business investment is also driven primarily by expected economic returns and, hence, is likely to have been much more sensitive to market prices such as the cost of capital, and to the exchange rate. By contrast, people need to live somewhere and government investment isn’t notably tightly tied to expected economic returns

An oft-heard line is that New Zealand has overinvested in housing. Given our population growth, that claim simply doesn’t stack up. We have devoted a relatively large share of GDP to house-building, but less than our population growth would have led one to expect. The sharp rise in real house prices certainly suggests too few houses rather than too many.

But that “given our population growth” is a critical caveat. In fact, over the last couple of decades, New Zealand has had one of the more rapid rates of population growth. Perhaps partly consistent with that, government investment as a share of GDP has also been among the highest in the OECD - roads, schools, hospitals etc.

Figure 4: Population growth differentials, by decade

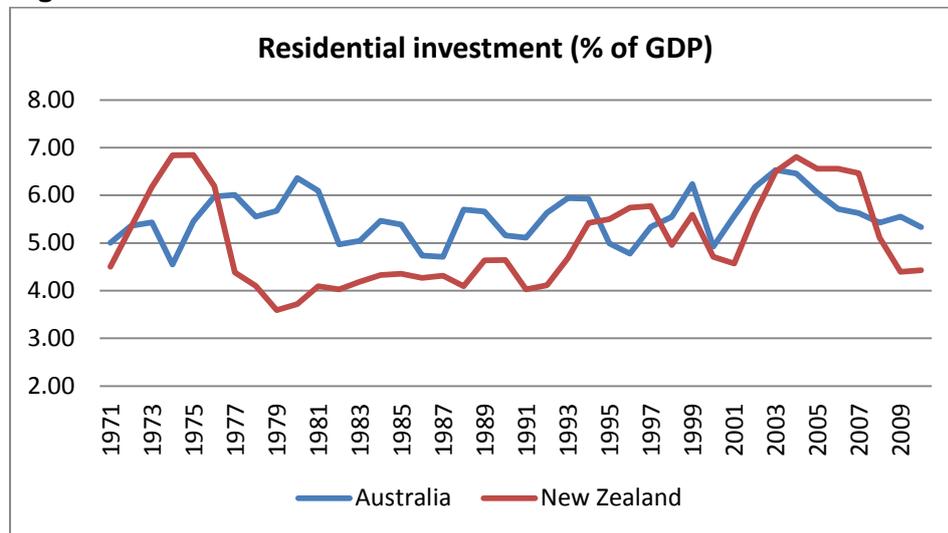


Source: OECD (Sample includes all current OECD countries)

New Zealand’s population has typically grown at a materially faster rate than the populations of most other OECD countries. That changed markedly for 15 years or so from the mid

1970s. As New Zealand's relative incomes and productivity performance fell away sharply we saw, for the first time, large sustained net outflows of New Zealanders - mostly pursuing better opportunities across the Tasman. Net inflows of non-New Zealanders continued, but at a pretty low level, and New Zealand experienced its lowest sustained rates of population growth in modern history.

Figure 5: Residential investment in New Zealand and Australia



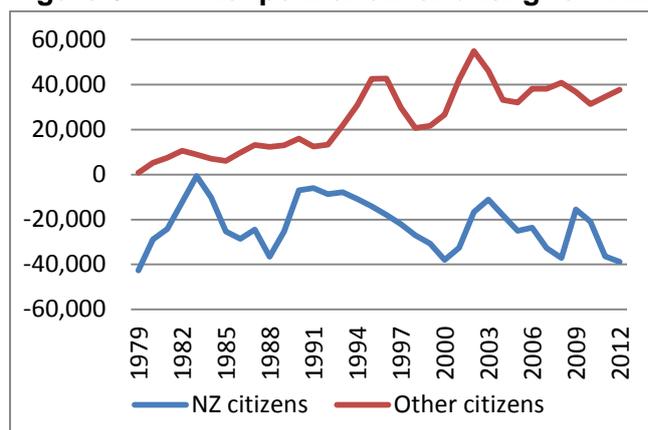
Source: OECD

This chart shows what difference it made to house-building activity in the 70s and 80s; comparing New Zealand's experience to Australia (the two countries had similar underlying demographics). Resources not spent on (quite labour intensive) house-building were free to be used on other things. Or, given our modest savings rate, other investment didn't have to be squeezed out to make room for the house-building previously required to accommodate a fast-growing population. Unfortunately, in the specific circumstances of that period we used them first on Think Big and, in the first blush of liberalised markets, on a massive commercial property boom.

Of course, the sharp dip in our relative population growth didn't last. What changed, so that our population growth was once again among the faster in the OECD? In short, migration policy did.

Substantial reforms in the late 1980s and the early 1990s led to, and were intended to lead to, a substantial increase in the net inflow of non New Zealanders.

Figure 6: Net permanent and long-term migration, by citizenship



Source: Statistics NZ

Indeed, over the last 20 years net migration of non New Zealand citizens has accounted for around 80 per cent of total population growth (the remainder is the underlying natural increase, offset by the very large net outflow of New Zealanders).

Each 1 per cent addition to the population typically requires another 3-4 percentage points of GDP in investment. If the savings rate does not lift when the population growth does - and you can see from earlier charts that it didn't in New Zealand - then that additional demand has to be met through some combination of a larger current deficit, or by squeezing out some other spending.

The pressure that large migration flows put on available savings, through the need for high investment - across all types of capital - just to maintain pre-existing capital/output ratios, was a recurrent feature in the New Zealand economic analysis of migration in the 1950s and 1960s. By contrast, little attention appears to have been given to this factor in the policy thinking on the liberalisation in the late 80s and early 90s, perhaps because of the liberalised approach to private financial capital. But if money is very mobile, physical capital isn't.

Repeated waves of new migrants over the past 25 years, drawn in as a matter of government policy, in an economy with only a modest savings rates, have, it appears, directly contributed to holding up average real interest rates and the average real exchange rate. The most easily squeezed component of spending is business investment, especially that in the tradables sector. "Desired" business investment - at a world interest rate and a real exchange rate consistent with that - may well have been high. The earlier reforms and the needs of a rapidly growing workforce make that quite plausible. But actual business investment looks to have borne the brunt of the relative prices (interest and exchange rates) we've actually experienced.

In other words, the policy-led reacceleration in population growth may have made catching up with the rest of the advanced world a whole lot harder. The New Zealanders leaving for better opportunities in Australia were probably helping both countries - just as 19th century Irish migration to the US and UK did. In our context, the net outflow of New Zealanders puts downward pressure on real interest and exchange rates, creating the opportunity to "crowd in" more business investment, at least once the extensive reforms had created a more propitious climate for private sector led growth. When government policy stepped in to, in

effect, actively replace those who were leaving, rather than waiting until materially improved economic conditions altered the relative attractiveness of staying in (or returning to) New Zealand, that worsened the prospects for those who remained. Total GDP grew quite rapidly, but at the expense of GDP per capita (and per hour worked). In terms I used earlier, the cart was put before the horse.

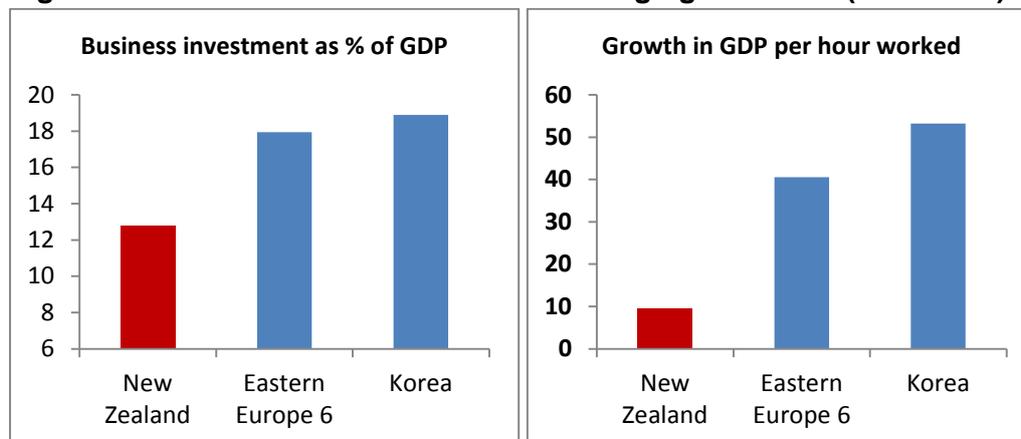
The tension would not have been anywhere near as great in a high savings economy. But that simply isn't the New Zealand of the last 20 (or 40) years. And the constraint of a modest savings rate wasn't a new development that became apparent only after the migration policy changes had been made.

Migration probably isn't typically a "bad thing" in economic terms. It certainly benefits the migrants (or they wouldn't come) but the evidence that it benefits citizens of the recipient country is much harder to find. Circumstances probably matter a lot. In very high savings Singapore and Switzerland it may well have helped - utilising more effectively, in a better-balanced way, the abundant national savings.

All else equal, a country with fast population growth should be expected to devote a larger share of GDP to investment (not just houses, but all types of capital). In the paper I included some simple scatter plots suggesting - somewhat counter-intuitively - that over the last decade or so those OECD countries with faster population growth had actually devoted a smaller share of GDP to non-housing investment, and had achieved lower growth in total factor productivity.

That led me to reflect more specifically on the experience of eastern European OECD countries. They liberalised and reformed their economies (from much worse starting points than ours) at much the same time we did, effectively facing the same challenge we did - trying to converge with the richer OECD countries. But consider the suggestive experience of the last decade (so not even the first few post-liberalisation years). They - like Korea, another convergence country - saw a high share of GDP devoted to business investment. - responding to favourable market signals and opportunities. Unlike New Zealand. They, like Korea, achieved rapid real productivity growth. Again, unlike New Zealand.

Figure 7: New Zealand and some converging countries (1999-2009)



Source: OECD

And they - like Korea, typically with quite low birth-rates – experienced little or no population growth, often experiencing quite material net migration outflows to western Europe. This isn't a mono-causal story. These economies also had higher savings rates than New Zealand did. But in a sense that is my point: their stars were aligned. Relatively high savings rates, much-improved micro policy frameworks, in combination with very low population rates, helped make room for rapid productivity growth and convergence in per capita GDP.

We, by contrast, appear to have inadvertently put a major element of policy directly at odds with the longstanding savings choices of our own people. Doing so, stifled the adjustment our citizens were already bringing about, exacerbated the macroeconomic imbalances, and seriously impeded the prospects for realising the full gains of our own extensive microeconomic reform programme. Reconciling those imbalances and conflicting pressures (through persistently high real interest rates and the high average real exchange rate) appears to have crowded out material amounts of business investment, probably reduced the rate of business saving, and skewed what business investment has occurred towards the non-tradables sector. And it has left our net external debt higher than it otherwise would have been. .

When a family has fallen on hard times, and has to devote lots of energy to stabilizing the situation and restoring the family finances, having another child at the same time isn't a terribly prudent economic choice. Having another child then will almost invariably worsen the family's economic position (whatever other joy it brings). It is a folksy comparison and breaks down at some points, but New Zealand is in some respects that family: choosing to have lots more kids, as it were, with all the attendant pressures, just when we were otherwise getting into a position to capitalise on the reforms put in place in the late 1980s and early 1990s and restore the family finances and income prospects.