# Modelling Regional Labour Market Adjustment in New Zealand

Wai Kin Choy, David C. Maré
and Peter Mawson

NEW ZEALAND TREASURY
WORKING PAPER 02/01

MARCH 2002

THE TREASURY
Kaitohutohu Kaupapa Rawa

**NZ TREASURY WORKING PAPER SERIES 02/01**

# Modelling Regional Labour Market Adjustment in New Zealand

**AUTHORS**

Wai Kin Choy

Reserve Bank of New Zealand,
PO Box 2498, Wellington, New Zealand

Email choywk@rbnz.govt.nz

Telephone 64-4-471 3642

Fax 64-4-471 3964

David C. Maré

Motu Economic and Public Policy Research Trust,
19 Milne Terrace, Island Bay, Wellington, New Zealand

Email dave.mare@motu.org.nz

Telephone 64-4-939-4252

Fax 64-4-383-4270

Peter Mawson

The Treasury,
PO Box 3724, Wellington, New Zealand

Email peter.mawson@treasury.govt.nz

Telephone 64-4-471-5288

Fax 64-4-499-0992

**ACKNOWLEDGEMENTS**

The authors would like to thank: Cyril Mako and Jeremy Foster of Statistics New Zealand for their assistance in the provision of the data used in this study; and Guy Debelle for making available the programs and other material used in Debelle and Vickery's (1999) study. The authors acknowledge funding for this project from the Departmental Contestable Research Pool (DCRP). The authors would also like to thank Catherine Atkins, Maryanne Aynsley, Bob Buckle, Guy Debelle, Melody Guy, Tim Hampton, Dean Hyslop, Veronica Jacobsen, Philip Morrison, James Newell, Peter Thomson, and other participants at seminars at The Treasury for valuable comments and suggestions. At the time this paper was written, Wai Kin Choy was working at The New Zealand Treasury.

**NZ TREASURY**

New Zealand Treasury, PO Box 3724, Wellington, New Zealand

Email information@treasury.govt.nz

Telephone 64-4-472 2733

Website www.treasury.govt.nz

**DISCLAIMER**

# Abstract

This paper adopts a vector autoregressive (VAR) approach to analyse the labour market adjustment mechanisms for 12 New Zealand regions over the period 1985 to 2001. It examines the effects of a region-specific shock to employment on itself, the unemployment rate, the participation rate, and the wage rate. The role of migration as a channel of regional labour market adjustment is also inferred. We find that adjustment occurs predominantly through inter-regional migration although the unemployment and participation rates also play a role. Wages, on the other hand, account for very little adjustment. The importance of inter-regional migration in New Zealand matches the results found in Sweden, but stands in contrast to the picture in many European countries. Migration appears to be a more dominant adjustment channel compared to the US and Australian cases. However, this has to be placed into context – New Zealand regions are much smaller in terms of population size.

**JEL CLASSIFICATION** R23 - Regional Migration; Regional Labour Markets; Population; J61 - Geographic Labour Mobility; Immigrant Workers

**KEYWORDS** Regional labour market adjustment; Internal migration

# Contents

# List of Tables



# List of Figures

# Modelling Regional Labour Market Adjustment in New Zealand

## 1 Introduction

When a region experiences a shock to employment demand[1], what are the implications for the region's labour market? Are there adjustment processes that act to mitigate the initial impact of the shock or is the region permanently affected? If adjustment processes are present, what are they? This paper addresses these questions in the New Zealand context by using a vector autoregressive (VAR) approach to analyse the labour market responses of 12 regions to a shock in employment demand.

There are several reasons for studying how New Zealand's regional labour markets respond when hit by employment shocks. First, theory suggests that regional labour markets should adjust in response to shocks through at least one of the following channels: changes in unemployment; changes in labour force participation; changes in worker remuneration; and migration. This is the theory, but does adjustment actually take place in practice for New Zealand regions, and if so, through which channels? This study empirically addresses these questions.

Secondly, the issue of how regional labour markets adjust to shocks is of international significance. Several overseas studies have previously applied the methodology used in this study to other countries or international regions. The primary adjustment channels have been found to differ internationally. For example, overseas studies have found that worker migration plays a substantial role in the local labour market adjustment process in the US and Australia. However, such labour mobility plays a much smaller role in the adjustment of many European labour markets to region-specific shocks. Instead, labour force participation changes bear most of the adjustment. How does adjustment for New Zealand regions compare?

Finally, the particular channels through which adjustment occurs can have implications for policy (see Maré and Choy, 2001). For example, in the regional development policy area,

---

[1] In the spirit of other studies of this type, this study abstracts from region-specific shocks to labour supply. In other words, we are only concerned with changes in labour supply induced by inter-regional migration. We therefore abstract from the fact that a disproportionately large number of immigrants settle first in Auckland, and that there may be exogenous shifts in preference towards living in certain places (eg, retirement migration). This abstraction is reasonable given that short term changes in employment measures are primarily due to changes to labour demand rather than labour supply.

there is a debate surrounding the appropriateness of people-based policies versus place-based strategies. One of the key issues in this debate is that attempts to improve the prospects of people in particular regions via place-based strategies may be confounded if the in-migration response is large. In other words, assistance could end up benefiting new entrants to the region rather than the initial target population or community.

There are a few reasons, noted by Decressin and Fatas (1995), why studies focus on sub national areas or regions.[2] Region-specific shocks may trigger different adjustment mechanisms compared to national shocks. For example, one might expect more inter-regional migration in response to region-specific shocks than international migration in response to national shocks. Secondly, if there is more specialisation in the production of goods and services at the regional compared to the national level, analysing the national labour market dynamics will give a picture that may be too aggregated. In this case, looking at regional dynamics is likely to provide interesting and relevant results. Furthermore, regional differences in labour supply characteristics could mean that regions may differ in their response to labour market shocks.

A negative labour demand shock implies a reduction in employment in the affected region. How do firms and individuals respond to this? One possibility is that unemployment increases; individuals who had a job no longer have one. But what do these individuals do? It may be that they decide that the region no longer provides them with the opportunities they desire and consequently leave the region in search of better job opportunities. It is likely that not all the people who decide to leave were those who directly lost their own jobs as a result of the shock. Family members of those who lost their jobs may also leave. People who were unemployed prior to the shock may leave as they see the likelihood of finding employment has decreased due to the shock. Even people with jobs may leave if they fear that the region is in decline and they may be next in line for job losses.

Another possible adjustment is that labour force participation may drop. Some of those people who lose their jobs may decide to take early retirement or decide that the best thing for their future prospects is to undertake further training. Some who have lost their jobs and those who were already unemployed may become discouraged and no longer actively seek work.

Yet another possible form of adjustment is through changes to the price of labour. The decrease in labour demand may have a negative impact on wages.[3] A decrease in wages may help mitigate the effect of the initial reduction in demand for labour by attracting new firms into the region, hence creating some new jobs to offset the initial shock, or discouraging existing firms from cutting jobs.

---

[2] Generalisations about regions and regional adjustment may in fact not apply to the whole region. This raises the question as to how regional labour market should be delimited or what level of disaggregation we should choose. Definitions about what constitutes a regional or local labour market vary considerably in the literature (see for example Box 1 in OECD (2000)). This study focuses on the regional council level. The question of alternative definitions of regional or local labour markets is being considered by other researchers. For example, Newell and Papps (2001) classify New Zealand into local labour market areas using travel-to-work data.

[3] In this paper, we use wages as a proxy for the price of labour. However, it is worth noting that wages are not the only component in this price. Other components could include such things as redundancy entitlements and superannuation benefits.

What this discussion suggests is that when analysing the impact of an employment shock, it will be important to capture the adjustment channels mentioned above. Consequently, the potential labour market adjustment channels included in our analysis are changes to the unemployment rate, employment level, the labour force participation rate, and wages. The role of inter-regional migration is also inferred from the analysis.

However, other adjustments such as intra-regional migration or commuting may also be important (see Maré and Choy, 2001). Furthermore, we provide a picture of regional labour market adjustment for the typical New Zealand region. There could be some idiosyncratic features of some regional labour markets, where case studies may be more appropriate. It is also worth noting that this study examines regional labour market adjustment in relation to the labour force as a whole, rather than specific labour force groups.

As a preliminary step, the next section takes a univariate look at what has been happening to employment, unemployment, labour force participation and wages in New Zealand and its regions over the study period. The remainder of the paper is organised as follows. Section 3 reviews the literature on regional labour market adjustment. Section 4 outlines the Vector Autoregressive (VAR) modelling approach that we adopt. In Section 5, we discuss the data to be used in the analysis, and present some preliminary analyses. The main empirical results are presented in Section 6 as well as a discussion on how sensitive our results are to a number of necessary modeling assumptions. In Section 7, we summarise the key findings from our results and discuss how our New Zealand results compare with those from other international studies. Section 7 also examines what the results mean for policy, and potential areas for future research.

# 2 New Zealand labour markets

This study of New Zealand regional labour market adjustment is based on data for the period 1985 to 2001. The results of any study need to be placed in the context of the environment from which they are obtained. This section provides a brief overview of what happened to key labour market variables over the study period at both the national and regional level.

## 2.1 The New Zealand labour market (1985 – 2001)

The study period largely coincided with a period of significant economic reform for the New Zealand economy. Maloney and Savage (1996, p. 187) note that:

> [t]he period from 1984 to 1990 was one of very extensive product market reform. These reforms included deregulation of factor markets, such as finance, energy and transport; removal of import protection and export incentives in the tradeables sector and a comprehensive overhaul of business laws and the tax structure. The continued effect of these reforms was to push many New Zealand industries towards a much more competitive price setting environment than they had experienced previously.

They also note that the little evidence on how these product market reforms affected the labour market was consistent with it reducing rents and aligning wage setting behaviour more closely to market pressures. Up until 1991, the reforms had indirect rather than direct impacts on the labour market. The Employment Contracts Act 1991 (ECA) changed this. According to Maloney and Savage, the ECA was an explicit shift away from collectivism and centralisation towards a focus on individual employment contracts in an attempt to get a decentralised, competitive wage setting system based on individual firm level bargaining. They go on to argue that the most important aspect of the ECA was that it abolished the national award system, although this was not explicitly mentioned in the legislation.

With the move away from the national award system, one might have expected that the ECA could have influenced the nature of regional labour market adjustment in New Zealand. In particular, wage adjustments might have had greater opportunity to assist regional labour market adjustment. Unfortunately the wage data that we have available do not provide sufficient information on wages prior to the introduction of the ECA to adequately test this, although we do observe the extent of wage adjustment post 1989.

Figure 1 plots national total and full-time employment throughout the study period. This period can be characterised by a fall in employment following the share market crash in 1987 to its trough in 1991. Between 1992 and 1996, there was a strong surge in employment, with subdued but positive trend growth subsequent to this.

***Figure 1 - National employment (1985:4 to 2001:2)***


Source: Household Labour Force Survey

The economic downturn post 1987 is also reflected in the unemployment figures with a dramatic increase in unemployment occurring between 1987 and early 1992 – in fact unemployment almost tripled in this period (see Figure 2).

The large increase in employment between 1992 and 1996 is largely reflected in the unemployment figures, with unemployment falling steadily until about mid 1995. While employment largely levelled off in the period 1996 to 1999, unemployment began another rise, albeit not to the same levels it had reached in the early 1990s. In recent years, unemployment has fallen back again, although there are still more people unemployed than was the case in the mid 1980s.[4]

---

[4] Broadly similar patterns are observed if one focuses on unemployment rates rather than the number of unemployed. Morrison (2001) provides a more comprehensive discussion of New Zealand's employment and unemployment experiences throughout the 1990s.

*Figure 2 - National unemployment (1985:4 to 2001:2)*


Source: Household Labour Force Survey

While not shown graphically, it is worth noting that the national labour force participation rate[5] declined between 1987 and mid 1989. In the second quarter of 1987, New Zealand's participation rate was close to 75%. This fell to as low as 72% in 1989. The labour force participation rate then increased between 1993 and 1996 before levelling off at about 75%.

## 2.2 Regional labour market profiles (1985 – 2001)

The national labour market profiles provided above largely reflect the fortunes of the national or macro economy. However, there are regional differences in terms of labour market performance. For example, Morrison (1999a) examines the labour market differences across 14 New Zealand regions in 1991 using census data. The labour market indicators that the study focused on were the labour force participation rate, the unemployment rate, the full-time work rate and full-time wage income. The results suggest that not only do differences exist but that there is also substantial correlation across the four indicators of labour market performance by region. However, Morrison's study does not look at regional labour market performance over time (ie, regional labour market adjustment). Ultimately, this study investigates how regional labour markets adjust in response to region-specific rather than macro shocks. We therefore now examine the extent to which the labour market fortunes of different regions have varied in comparison with the national picture and other regions.

It is worth highlighting that it is not just labour market outcomes that vary between regions. The sizes of the regions used in the study are not uniform and exhibit considerable variation in terms of population size. To give a sense of the variation across regions, Figure 3 plots the 1996 working age population for the 12 New Zealand regions used in this study.

---

5 The participation rate is measured as the proportion of the working age population (15-64 years) that is part of the labour force. In other words, the proportion of the working age population that is either employed, or unemployed but seeking work.

*Figure 3 - Size of regions in 1996: Working age population (15 – 64 years)*[6]

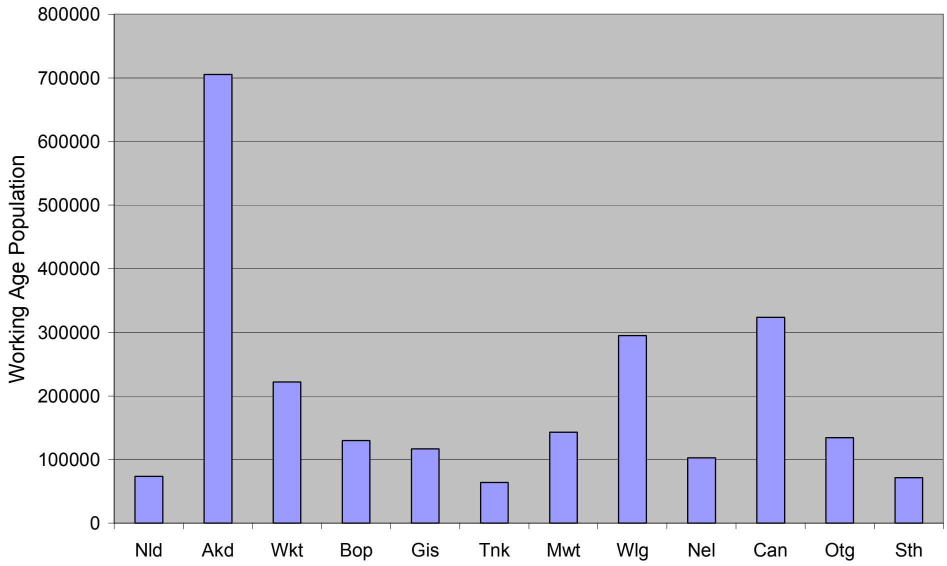


Source: Household Labour Force Survey

### 2.2.1 Employment

Figure 4 shows the employment performance of the 12 regions used in the study. The panels show both the employment level for each region as well as national employment (the scales have been adjusted to reflect differences in population size). Auckland has the largest employment level, followed by Canterbury and Wellington. The average number of people employed in these three regions over the study period was 470,000, 215,000 and 202,000 respectively. At the other end of the scale, the average employment levels for Southland and Taranaki were both just over 44,000. Even with the different scales, there is enough variation in the employment patterns that have been plotted to clearly indicate that regions differed in their employment experiences over the study period.

The differing employment fortunes of the different regions are perhaps even more apparent when employment shares (the proportion of national employment that the region's employment level represents) are plotted (see Figure 5). There is a trend decline in Manawatu's employment share (from 1985), as well as in Gisborne, Otago, and Southland (the last three from around 1992). On the other hand, there is a trend increase for Canterbury. It is also worth noting that Auckland's employment share exhibits the largest deviations in absolute terms (ie, it has the highest standard deviation).[7]

---

[6] The names of the regions in Figure 3 are spelled out in full as follows: Northland, Auckland, Waikato, Bay of Plenty, Gisborne, Taranaki, Manawatu, Wellington, Nelson, Canterbury, Otago and Southland.

[7] The movement appears to track quite closely with international net migration trends (see Glass and Choy, 2001). In fact, census data show that overseas inflows (particularly those not from Australia and the United Kingdom) are heavily concentrated in Auckland, while there is lack of census data on overseas outflows (see Mare and Choy, 2001).

### *Figure 4 - Employment levels by region*

Left Vertical Axis displays: total number of people employed in region (000s)
Right Vertical Axis displays: total number of people employed nationally (000s)
Horizontal Axis displays: year

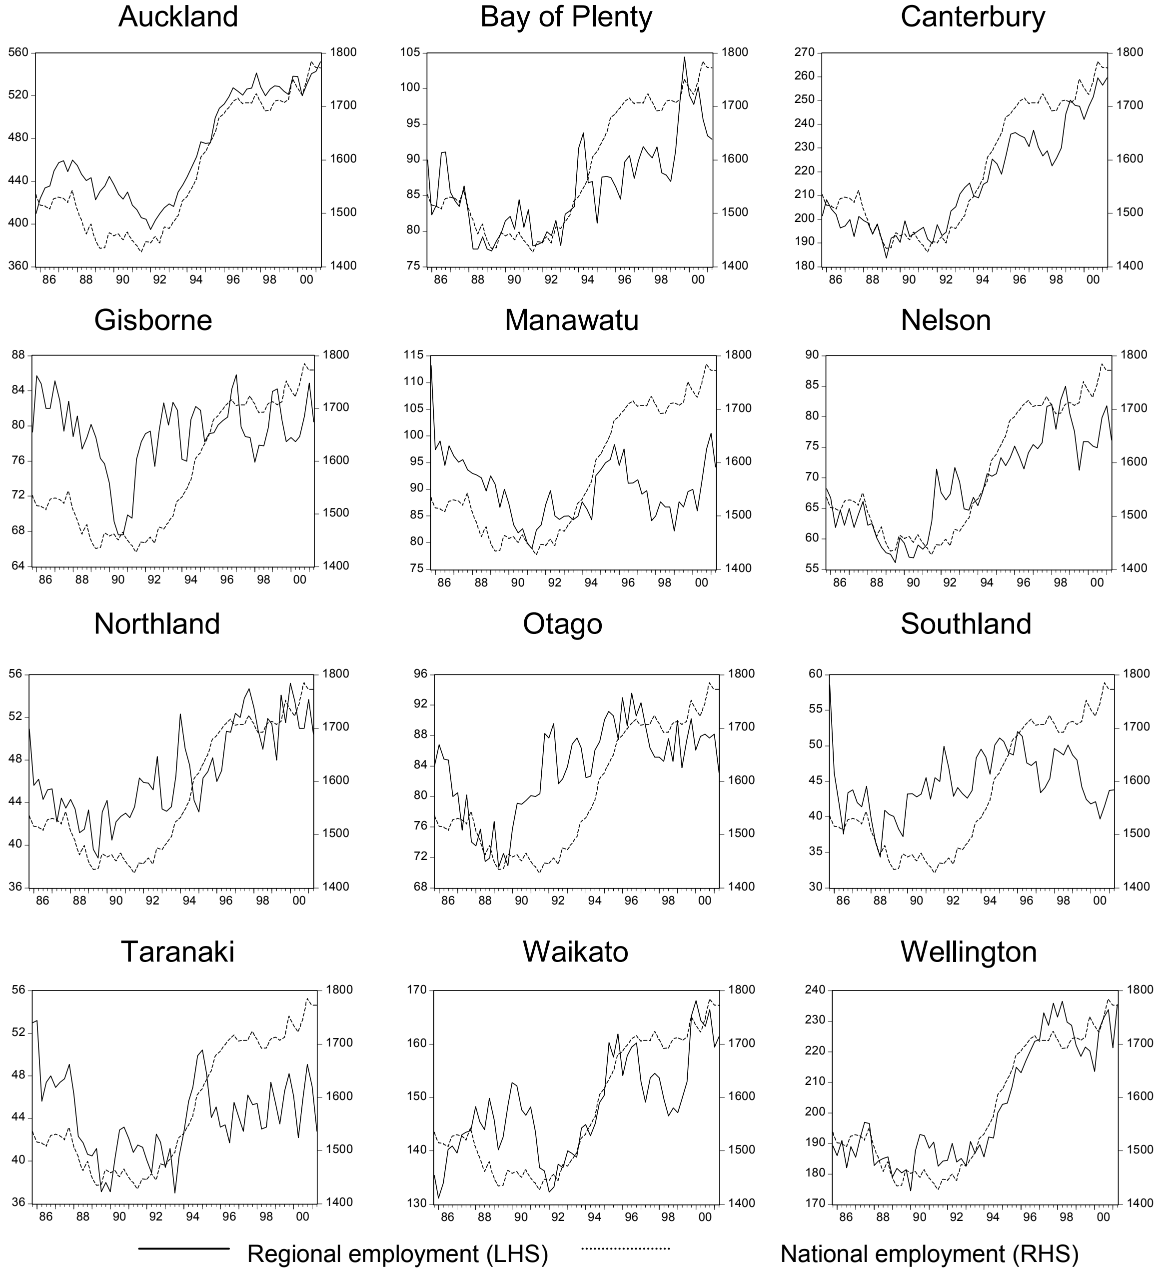


Source: Household Labour Force Survey

Note: The left-hand scale (LHS) differs for each region, whereas the right-hand scale (RHS) displays the corresponding national figure which stays the same for all regions.

***Figure 5 - Employment share by region***

Vertical Axis displays: Employment in each region as a percentage of total employment
Horizontal Axis displays: year

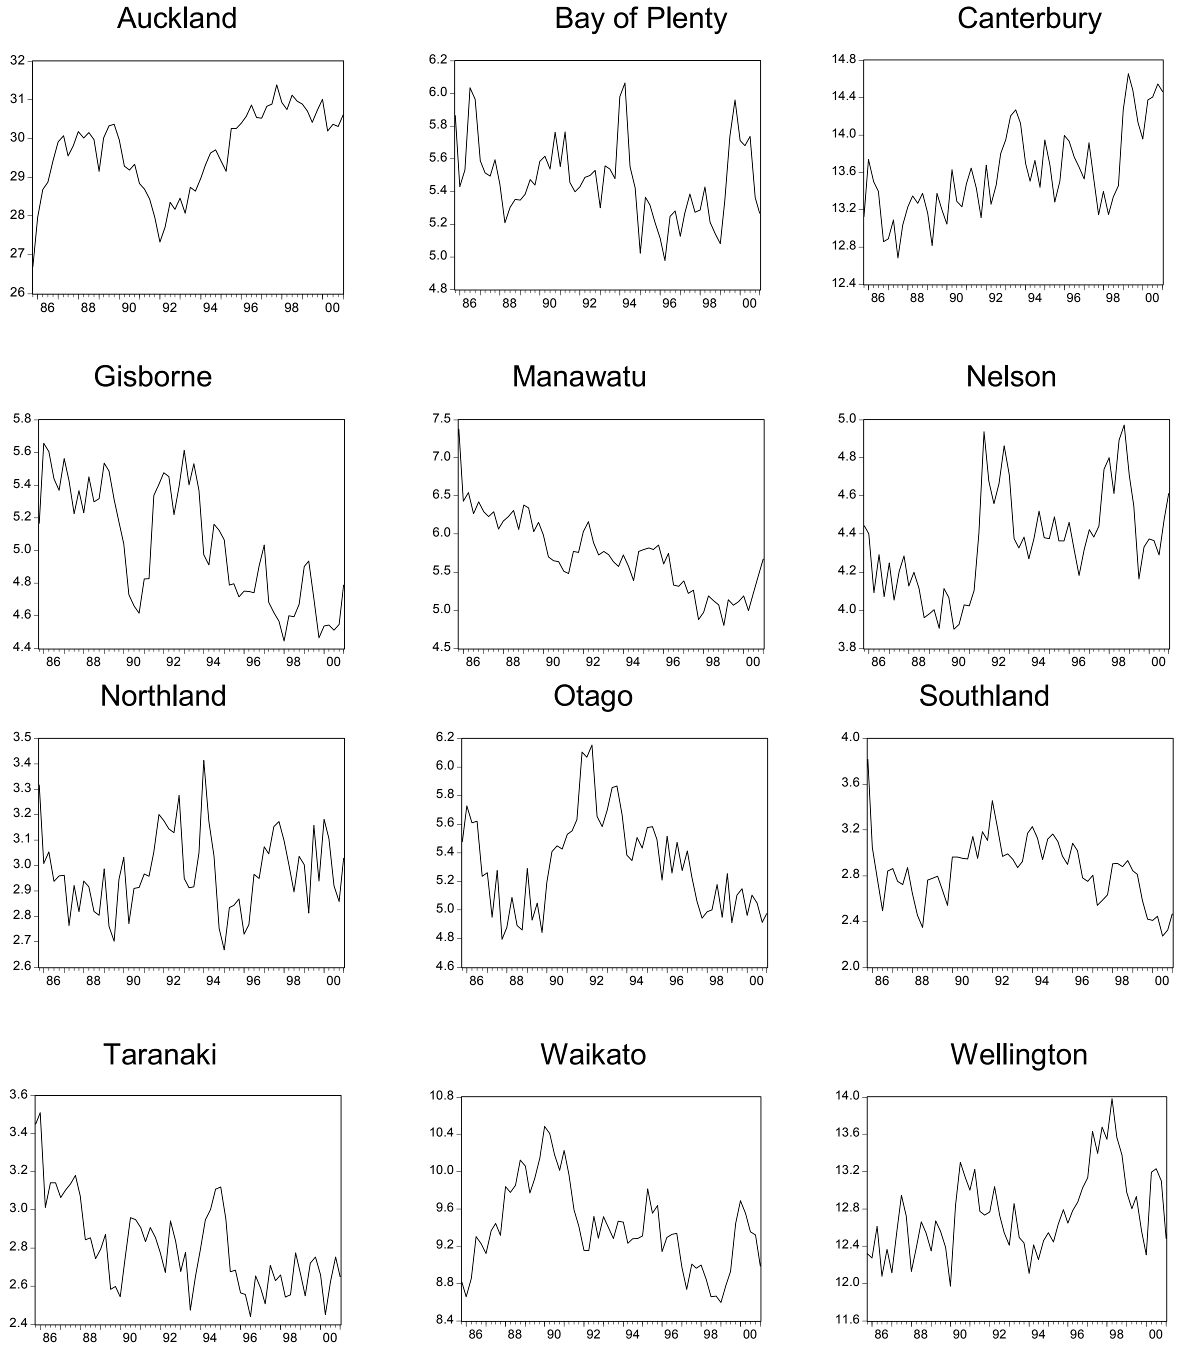


Source: Household Labour Force Survey and authors' calculations
Note: The scale differs for each region.

The fact that regions have different employment fortunes indicates that the different factors and shocks influencing employment are not all uniformly spread across regions. It is these phenomena that are of interest in this study, ie, shocks that only affect one region or the component of a shock that while affecting other regions is disproportionately felt in the region.

For a region to have an increasing (decreasing) employment share, the region must be growing faster (slower) than the national rate of employment growth. Clearly, the graphs above show that there have been differences in regional employment growth over time (otherwise, they would display a straight line). Next we examine whether these differences in regional employment growth persist over time.

Figure 6 shows the average annual employment growth for the regions over the periods 1986 to 1990 and 1996 to 2000. The regression line obtained has a slope of 0.31 (significant at the 10% level) and an $R^2$ of 0.28. Therefore, there seems to be some persistence in regional employment growth. Those regions that were growing relatively fast (compared to the national average growth rate) in the first five years of the sample continued to do so in the last five years to 2000.[8] Maré and Timmins (2001) also found a positive correlation (although not significant at the 5% level) between employment growth rates for the periods 1986 to 1991 and 1991 to 1996 using census data.

***Figure 6 - Regional employment growth rates in New Zealand (1986-2000)***

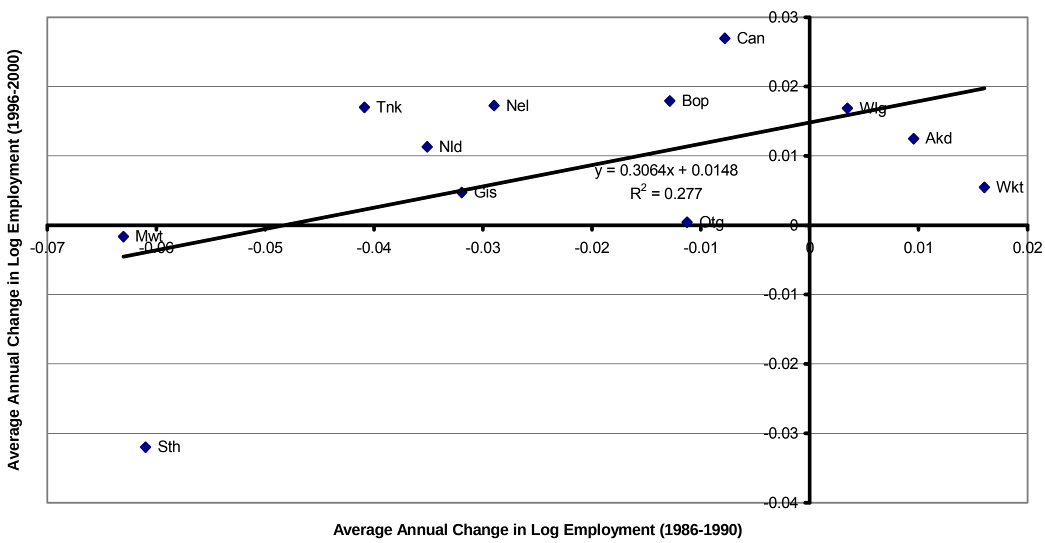


Source: Household Labour Force Survey and authors' calculations

Decressin and Fatas (1995) provide results from similar regressions examining the persistence of regional growth rates for both Europe (the European Economic Community) and the United States. For Europe, the slope of the regression line was found to be 0.55 ($R^2$ of 0.16), whereas for the United States, a slope of 0.25 ($R^2$ of 0.10) was found. New Zealand would therefore appear to fall between the United States and Europe but being much closer to the United States. Therefore, at a first glance, there is less persistence of employment growth in New Zealand and the United States than is the case for Europe. It is worth pointing out that confidence intervals are not provided for either the United States or Europe.

### 2.2.2 Unemployment rates

We now examine how the unemployment rates for different regions evolved over the study period. Figure 7 plots the unemployment rate for each region. The national unemployment rate is also plotted in each panel to aid an evaluation of relative performance. Figure A3 in Appendix A plots relative employment rates (in logs) for each

---

[8] We explored whether the picture of persistence changed if we used different end points (eg, 1986 to 1993 and 1994 to 2000). We find that the conclusion that there is persistence in regional employment growth is reasonably robust.

region, and therefore also enables an evaluation of relative employment rate (and consequently unemployment rate) performance.

***Figure 7 - Unemployment rates by region***

Vertical Axis displays: Unemployment rate
Horizontal Axis displays: year

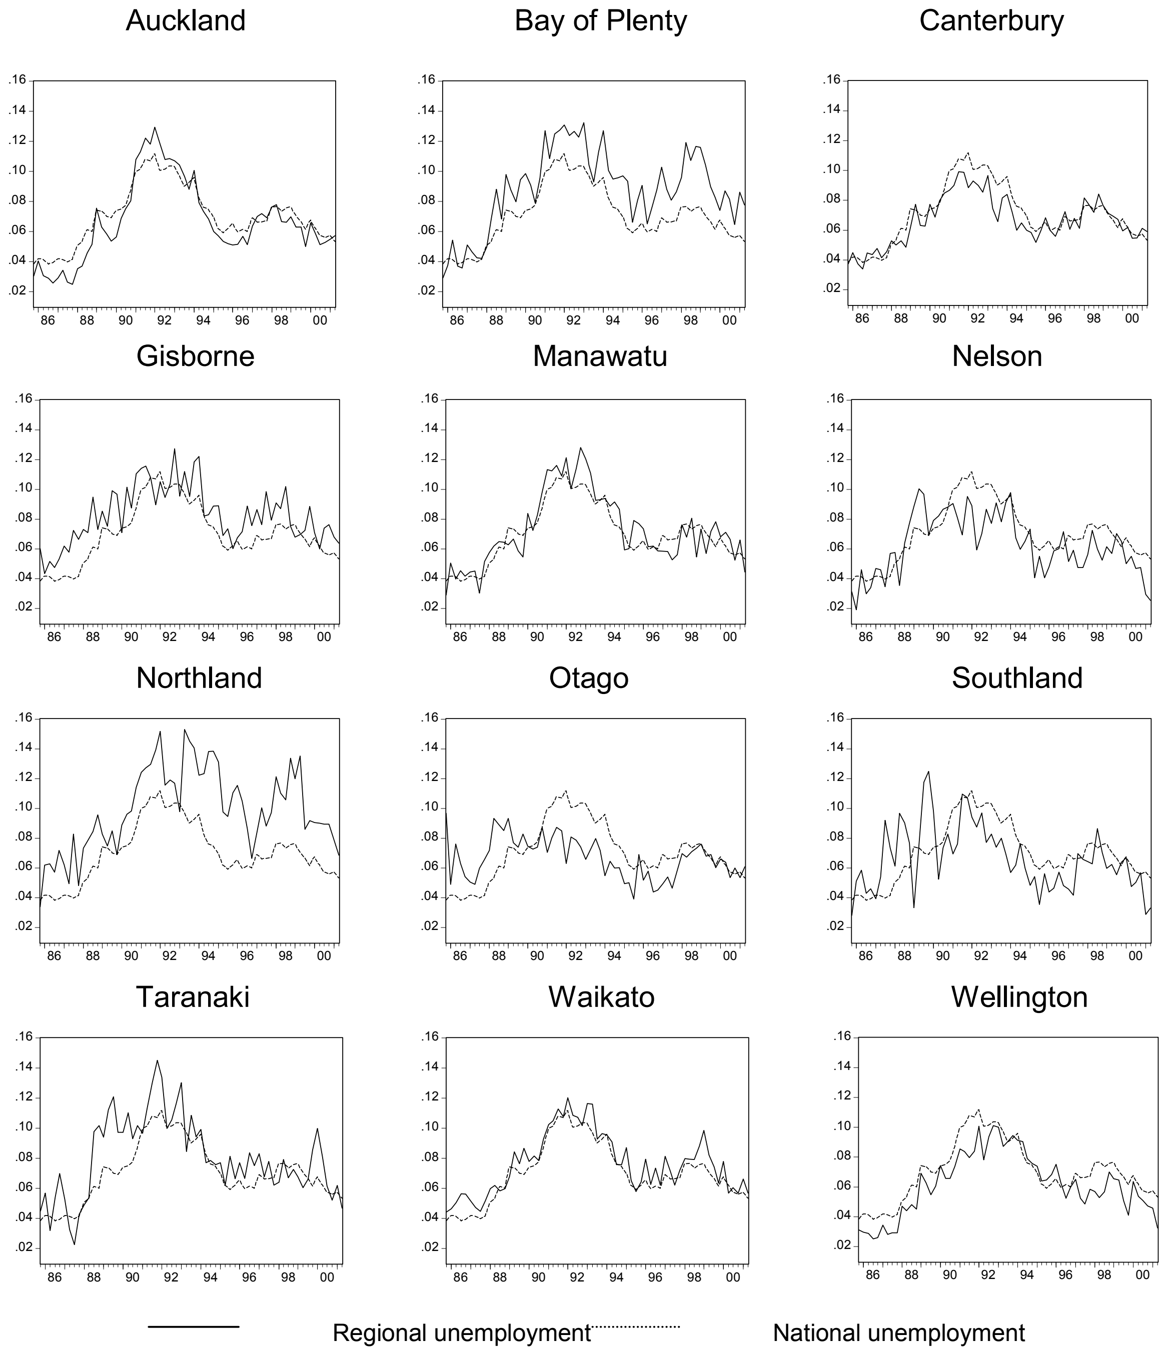


Source: Household Labour Force Survey

Throughout the study period, both Bay of Plenty and Northland have tended to have higher unemployment rates than the national average. On the other hand, Wellington has tended to have a lower unemployment rate than the national figure. Canterbury has tended to either track the national figure or be just below it. It can also be seen from the unemployment rate plots that the basic shape of some of the regional unemployment rate plots does not closely follow the national picture (see for example Otago). Thus, it appears that New Zealand regions can differ both in terms of their unemployment rates and also the pattern of evolution.

The observation that Bay of Plenty and Gisborne tended to have higher than average unemployment rates and Wellington lower than average rates would be consistent with there being persistence in unemployment rate outcomes. We investigate this by using the same method as used for the persistence of regional employment growth rates. According to Figure 8, there appears to be some persistence between a region's unemployment rate in 1986 and the corresponding figure 14 years later, as shown by a slope coefficient of 0.64 that is statistically significant at the 5% level. In other words, regions that begin with a high unemployment rate tend to end up with high unemployment rates later.[9] Our finding of persistence in regional unemployment rates is consistent with the findings reported in Maré and Timmins (2001) based on census data.

***Figure 8 - Persistence of unemployment rates across NZ regions (1986-2000)***

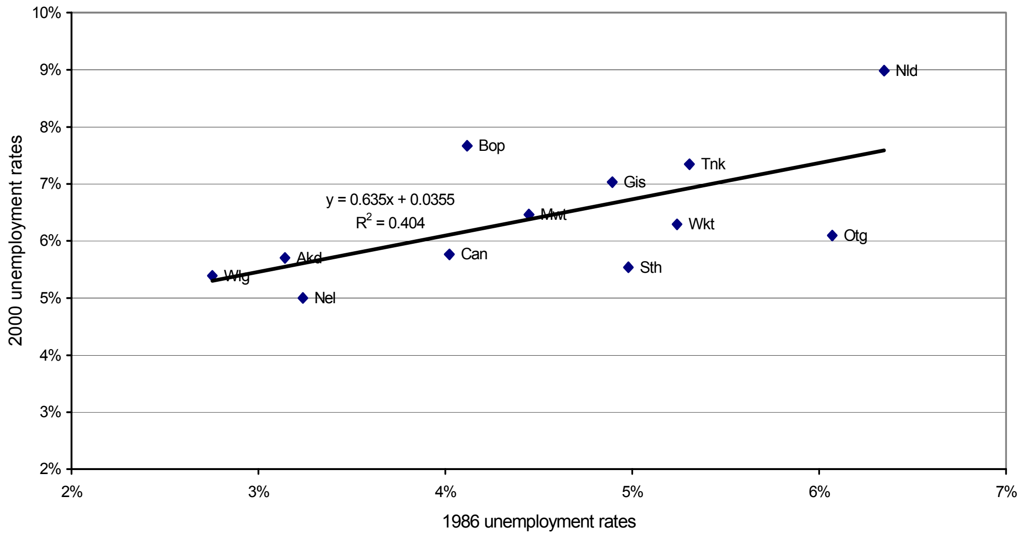


Source: Household Labour Force Survey and authors' calculations

It is possible to find examples of both persistence and no persistence in unemployment rates from overseas studies. Fredriksson (1999) quotes a 1999 OECD publication that found that the correlation between regional unemployment rates over time ranged from 0.46 to 0.92 for European countries (ie, evidence of persistence) but that the correlation between state unemployment rates was negative in the United States (ie, no persistence).

### 2.2.3 Labour force participation rates

The panels in Figure 9 show the participation rate in each region as well as the national participation rate. Participation rates are defined in this study as the proportion of the working age population (aged 15 to 64) that is part of the labour force. As can be seen from Figure 9, there are differences between regional participation rates in New Zealand. The participation rate in Manawatu is below the national figure throughout the study period. Northland and Otago also generally have a lower than average participation rate. On the other hand, Wellington and Nelson (post 1990) have participation rates that are above the national average.

---

[9] When we used different end dates, say 1986 against 1996, and 1990 against 2000, the same picture (ie, a positive slope) is obtained, although the different sample periods give different $R^2$, coefficients and t-statistics.

*Figure 9 - Labour force participation rates by region*

Vertical Axis displays: Participation rate
Horizontal Axis displays: year

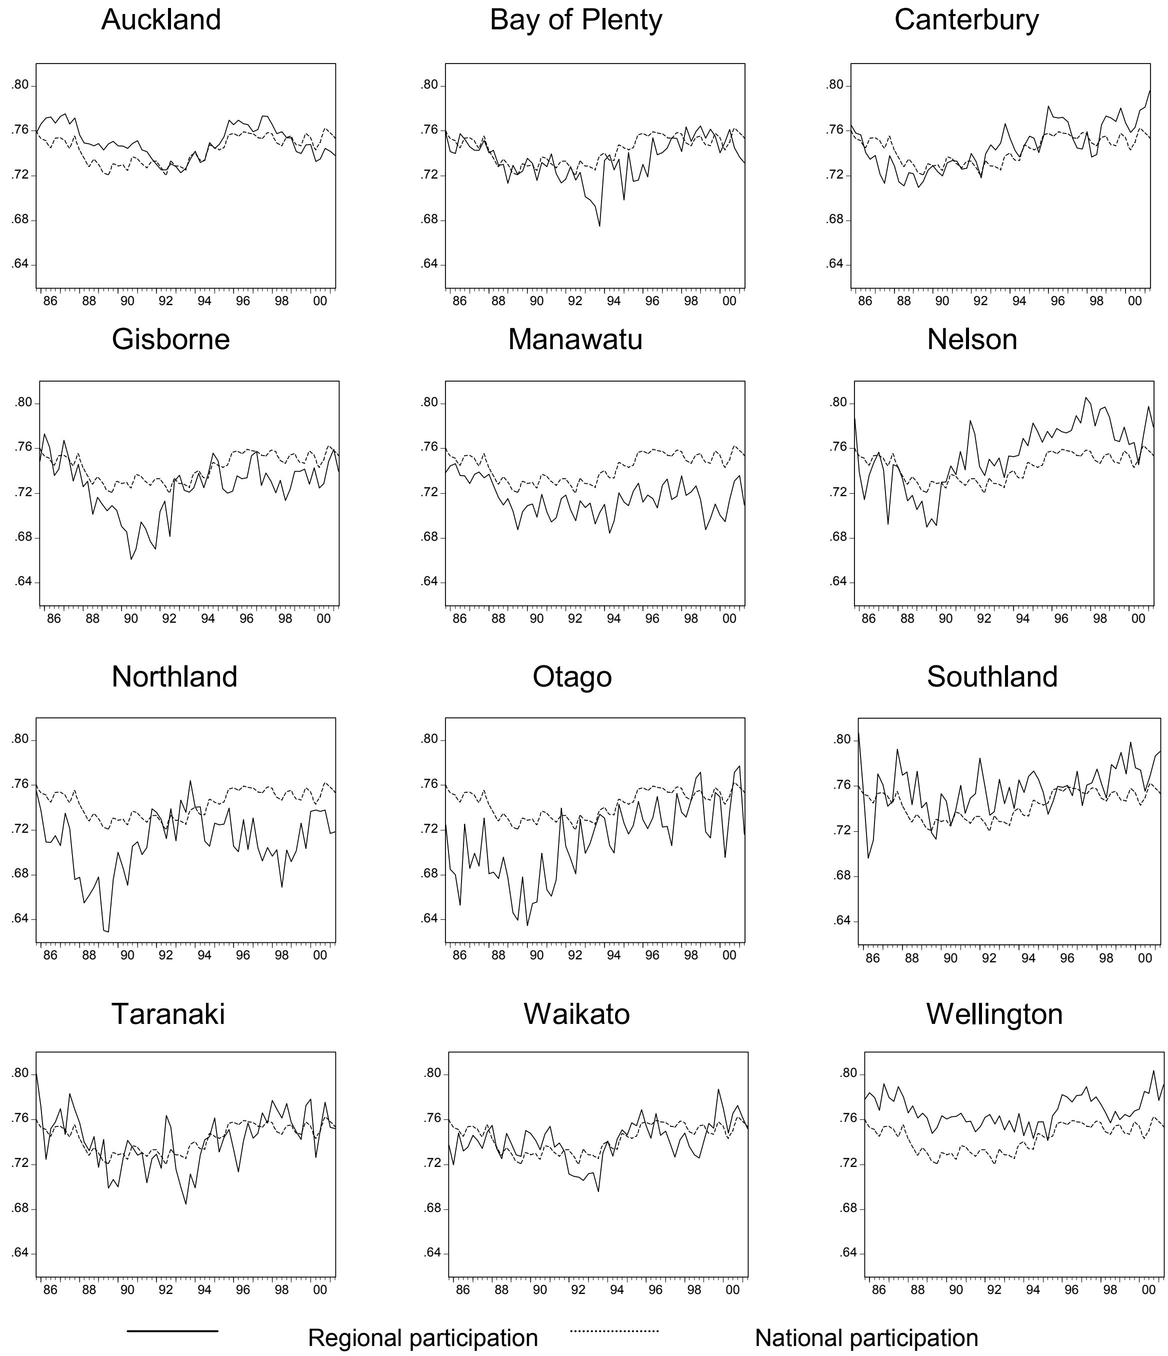


Source: Household Labour Force Survey

Figure 10 investigates whether there is persistence in labour force participation rates across New Zealand regions between 1986 and 2000. While the regression yields a positive slope coefficient of 0.27, the coefficient is not statistically significant at the 10% level. Maré and Timmins (2001) also find that their positive correlations between participation rates in 1986 and 1996 are not statistically significant. However, they do find statistically significant (at the 5% level) positive correlations between both 1986 and 1991 participation rates and 1991 and 1996 participation rates. Therefore, it would appear that participation rates are not strongly persistent over the longer term, but over the shorter term there can be persistence.[10]

***Figure 10 - Persistence of labour force participation rates across NZ regions (1986-2000)***

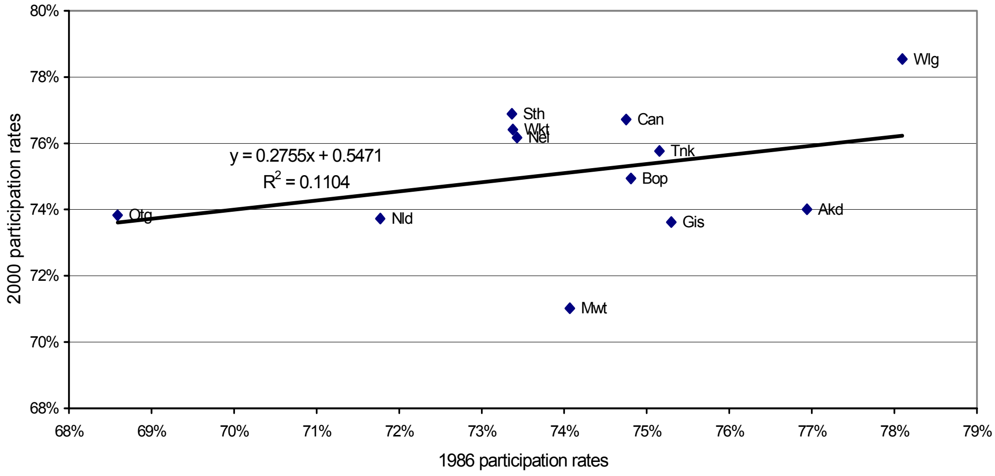


Source: Household Labour Force Survey and authors' calculations

### 2.2.4 Wages

In general, regional wage levels are either persistently higher or persistently lower than the national wage (see Figure 11). For example, wages in Wellington were approximately 10% higher than the national wages over the whole period. On the other hand, wages in Gisborne, Nelson and Manawatu are between 5 and 10% lower than the national wage level. Another interesting observation from Figure 11 is the stark seasonal pattern of wages in Southland and Otago.

Figure 12 confirms what was obvious from Figure 11 - there is persistence in rankings of wage levels across regions. It is, however, difficult to see from Figure 12 whether low-wage regions are "catching up" or otherwise with high-wage regions. Figure 13 is more informative in this regard – it plots the average rate of growth of hourly wages from 1990 to 2000 against their log value in 1990. It shows that low-wage regions have lower growth rates in wages, while high-wage regions continue to have high wage growth rates. Therefore, not only are low-wage regions not catching up, the wage differential is growing.

---

10 See Morrison (1999b) for a further discussion with regard to regional unemployment and labour force participation in the 1990s.

***Figure 11 - Relative wage levels by region (as a ratio to the national wage level)***

Vertical Axis displays: Average regional hourly wage as a proportion of the national average
Horizontal Axis displays: year

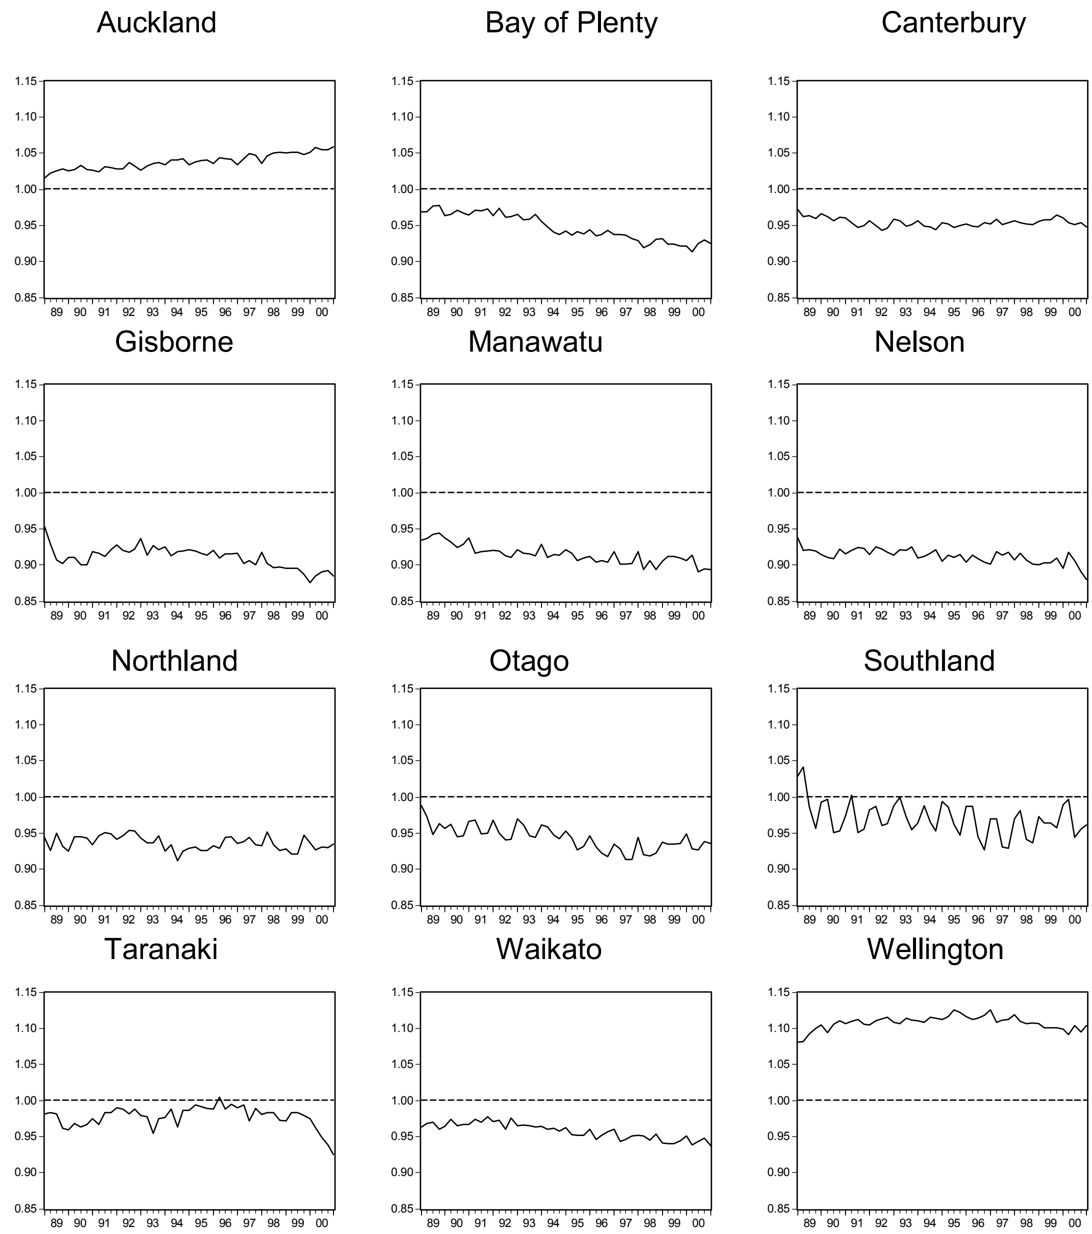


Source: Quarterly Employment Survey (QES)with authors' calculations and amendments[11]

---

[11] It is noted that the wages data used in this study have been amended slightly to account for changes in the QES data in 1999Q4, as discussed in Section 5.1.

***Figure 12 - Persistence of wages across NZ regions (1990-2000)***

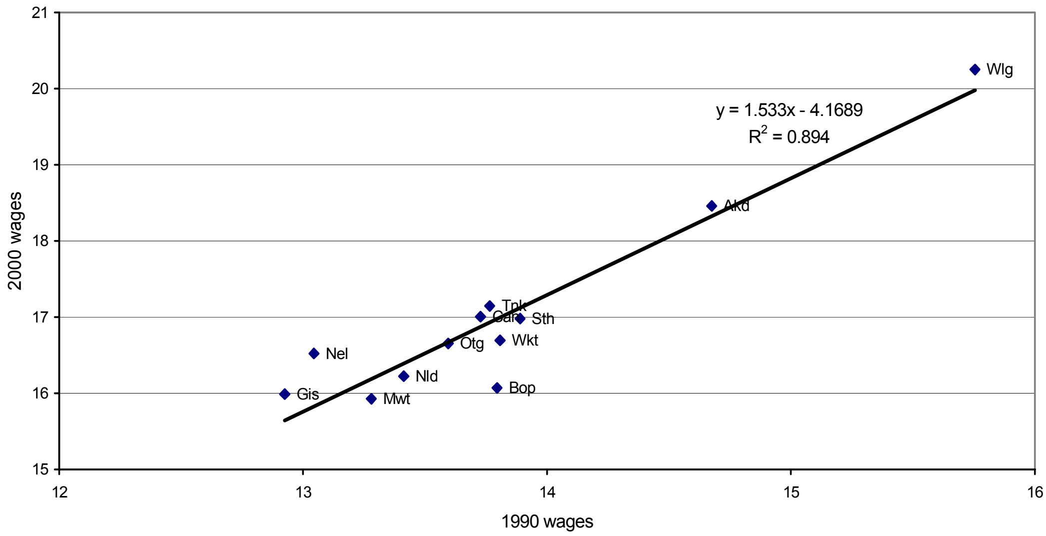


Source: Quarterly Employment Survey

***Figure 13 - Divergence of wages across NZ regions (1990-2000)***

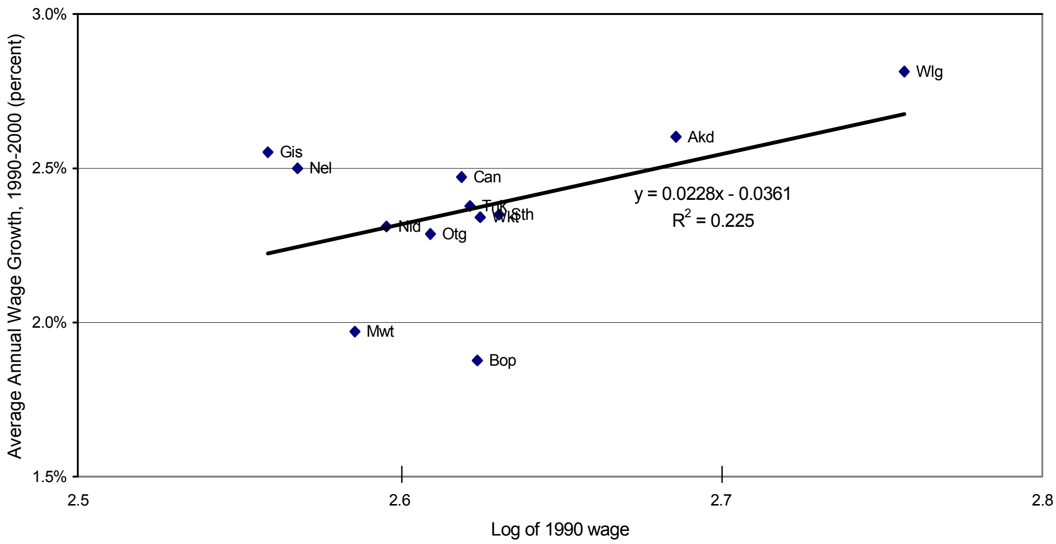


Note: The annual wage growth is measured by the average annual change in log wages over the 1990-2000 period.
Source: Quarterly Employment Survey

### 2.2.5 Regional labour market adjustment

As noted in Maré and Choy (2001), there is no universal agreement on what constitutes regional labour market adjustment. To many commentators, regional adjustment occurs when differences between regions become less. It is important to note that this view implies that the equilibrium is one when all regions are the same. An alternative viewpoint is that there may be stable long-run differences between regions meaning that regional adjustment entails restoring long-run relativities after a regional shock. We believe that it is possible to have permanent differences across regions, and thus we allow for this in our model, as discussed in Section 4.

In the previous subsections, we examined whether there was persistence and convergence/divergence for four different labour market variables. Simple tests like these are often used to provide preliminary evidence on the amount of regional labour market adjustment that occurs. Unfortunately, such simple tests do not fully answer the question of the extent of regional adjustment. A finding of little or no persistence would be consistent with quite fast and large labour market adjustments occurring in response to shocks.

As was noted when examining the time series history of the four different labour market variables (employment, the unemployment and participation rates, and wages), the preliminary evidence for New Zealand regions points toward there being persistence in possibly all of the variables.[12] These sustained differences between regions may suggest that adjustment is too slow to equivalise levels over the 15 years that we examined. Alternatively, it may be that for at least some (if not all) of the labour market variables that we are investigating, there are equilibrium differences in levels that do not necessarily require adjustment.

As our simple preliminary investigations into labour market adjustment do not adequately answer the question of the extent of labour market adjustment, the remainder of the paper explores the issue of regional labour market adjustment in New Zealand in more depth.

# 3 Literature review

In this study, we use a time series technique known as the Vector Autoregressive (VAR) approach.[13] Blanchard and Katz (1992) first applied this VAR approach to the investigation of labour market adjustment in the United States. This methodology has since been adopted by others such as Decressin and Fatas (1995), Debelle and Vickery (1999), Fredriksson (1999), and Jimeno and Bentolila (1998). This section briefly reviews the findings of some of these VAR studies. Further discussion of the VAR methodology used is provided in Section 4.

From the literature on regional labour market adjustment, it is clear that the speed of adjustment, and channels through which adjustment occurs, differ depending on the area of the world under investigation. Blanchard and Katz (1992) found that in the United States, migration plays a substantial role as a regional labour market adjustment mechanism. In other words, adjustment to labour demand shocks appears to occur mainly through migration flows (ie, laid-off workers leave depressed areas to find jobs elsewhere). Following a state-specific shock, the migration response is strong even in the first year after the shock. For example, if relative state employment falls by 10 workers, in the initial year, unemployment rises by 3 workers, participation falls by 0.5 workers, and 6.5 workers migrate out of the state. In the long run (after 7 to 10 years), employment falls by approximately 13 workers, all of whom have migrated to other states (as cited from Debelle and Vickery, 1999). Blanchard and Katz (1992) also conclude that wages decrease and dampen the employment response, but by relatively little. This evidence

[12] However, the evidence is not strong enough for one to obtain statistically significant measures of persistence in participation rates over the longer term (between 1986 and 2000).

[13] Different approaches have been used in the literature to assess regional labour market adjustment mechanisms and more particularly, the role of migration. The approaches differ according to the length of time series variation that is modelled, the range of variables that are modelled as changing, and how much theoretical structure is imposed (see Maré and Choy, 2001).

suggests that in the US, wages play a limited role as a regional labour market adjustment mechanism in response to economic shocks.

Applying Blanchard and Katz's methodology to Europe, Decressin and Fatas (1995) find that in European labour markets, labour force participation rate changes play a larger role in bringing unemployment back to trend after a region-specific shock, rather than migration. In other words, workers leave the labour force rather than migrate out of their region.

In addition to the finding above, Decressin and Fatas' (1995) study also illustrates that, even within a particular region, the exact nature of adjustment for a regional labour market may differ depending on whether the shock experienced is a region-specific or economy-wide (macro) shock.[14] When examining the univariate adjustment processes and focussing solely on how relative unemployment responds (or adjusts) when hit by a shock to itself, they note that "*regional relative unemployment rates in Europe return to their means fairly quickly*" (p. 1640).[15] Adjustment to US relative unemployment shocks is also quite fast.

Given the rigidities present in European labour markets, Decressin and Fatas note that the finding of fast adjustment to relative unemployment shocks could seem surprising. They therefore also examine what happens when *absolute* regional unemployment is shocked rather than *relative* unemployment.[16] The adjustment of US absolute unemployment to a shock is still fast whereas European adjustment takes a lot longer and points to there being permanent impacts on absolute unemployment. They therefore conclude that it is common or macro shocks which have permanent effects in Europe (as shocks to absolute regional unemployment include both region-specific and macro shocks). We compare the univariate adjustment processes of New Zealand regions to both region-specific and aggregate shocks in Section 5.

Debelle and Vickery (1998 and 1999) investigate Australian state labour market adjustment using the VAR methodology. They were particularly interested in the role of labour mobility across states. Debelle and Vickery (1999) "*find that interstate migration does play an important role in reducing differences in labour market conditions between states, although permanent (or very persistent) differences between state unemployment rates remain*" (p. 249). While migration does play an important role, they find that migration occurs slowly over a number of years with only about one third of the net migration that does occur taking place within two years of the shock.

Fredriksson (1999) also finds that migration between regions plays an important role in Swedish regional labour market adjustment. Fredriksson examines fears held by some that active labour market programs in high unemployment regions may be having adverse effects on adjustment by locking-in workers who have lost jobs in depressed regions. He finds little evidence in support of such fears with labour mobility being high and rapid compared to European standards. In the first year following a shock, 66% of the adjustment that occurs happens through migration. This rises to 87% in the second year.

---

[14] The terms macro, common, aggregate and economy-wide shocks are used interchangeably in this paper.

[15] Regional relative unemployment rates refer to the regional unemployment rate being measured as being relative to the national average (in Decressin and Fatas' case, the European average).

[16] Absolute regional unemployment is the regional unemployment rate with no adjustment for the European average rate of regional unemployment.

Given the extent of labour market integration between Australia and New Zealand (see Poot, 1995; and Bushnell and Choy, 2001), the results from our study have to be put into the trans-Tasman context. Several studies have compared New Zealand and Australian labour market adjustments, such as McCaw and McDermott (2000) and Aynsley (2000). McCaw and McDermott treat New Zealand as a state alongside Australian states and examine the role of migration compared to other labour market adjustment mechanisms using a similar methodology to that applied in this paper. Their results indicate that migration is an important adjustment mechanism for New Zealand, and is a more significant part of the adjustment process than for Australian states.

Aynsley (2000) compares the labour market adjustment mechanisms for Australia and New Zealand over the period 1978 to 2000. Aynsley examines labour market adjustment at the national rather than regional or state level, and focuses on aggregate shocks in her analysis. The findings suggest that an employment demand shock has a larger and more persistent impact on the New Zealand than the Australian labour market. However, international labour mobility appears to be a more important adjustment mechanism for New Zealand than Australia. The author suggests that the apparent difference in the role of international labour mobility may be due to the different size of the labour markets. Since Australia is larger than New Zealand, there are a greater number of local labour markets within the former, and thus there are likely to be greater opportunities for people to relocate within the country in response to an adverse shock in a local area. This is consistent with Debelle and Vickery's (1999) finding that labour mobility is an important inter-state labour market adjustment mechanism in Australia. In contrast, it is more difficult for people in New Zealand to find as great a range of alternative employment opportunities within the country. Aynsley (2000) also finds that changes to labour force participation is an adjustment mechanism for both countries, but not the real wage.

There are also a few other VAR-type studies that have used slightly different specifications. For example, in the US case, Bartik (1991) considers a once-and-for-all shock to local job growth, with subsequent growth unchanged from what it would have been (ie, a one-time growth shock). In contrast, Blanchard and Katz (1992) allowed the one-time shock to local job growth to affect subsequent growth (ie, a shock with readjustment). Not surprisingly, the studies reached different conclusions about local job growth's long run effects. Mauro, Prasad et al. (1999) examined the regional labour market adjustment process in Spain for the different skill groups within the Spanish population. They find evidence that suggests that the high-skilled in Spain are more likely to migrate than remain unemployed or drop out of the labour force, compared to the low-skilled. In other words, high-skilled workers migrate very promptly in response to a decline in local labour demand while low-skilled workers drop out of the labour force or stay unemployed for a long time.

Given that our study of New Zealand regional labour market adjustment uses a methodology that has been applied to a number of countries, we are able to compare New Zealand regional adjustment with that of other countries. We will also compare our results with other New Zealand studies using approaches other than the VAR technique, as discussed below.

In the New Zealand context, while a great deal of work has been done in the internal migration and in the regional labour market areas, there are few studies that investigate the interactions between the two. Nevertheless, such studies that focus on either area separately can still provide relevant insights for this study.[17]

Existing studies from the Waikato University Population Studies Centre provide a good summary of inter-regional migration patterns over the 1981-1996 period (for example, see Goodwin and Bedford, 1997, and Bedford, Goodwin et al., 1997). The main conclusion from these studies relevant to our work is that internal migration is an important source of population structure change for New Zealand regions. However, these studies do not explicitly investigate the links between migration flows and labour market outcomes.

There have also been studies that provide descriptions of local labour market conditions in New Zealand. For example, Morrison (1999a) finds that the 14 regional labour markets in New Zealand can be characterised in terms of four indicators, namely the labour force participation rate, unemployment rate, fulltime work rate and fulltime wage income. Other contributors to this line of work include the New Zealand Planning Council (NZPC (1989)) and Chapple (2000).

Chapple (2000) examines regional labour market adjustment by focusing on urban area units (which the author calls neighbourhoods). Chapple (2000) investigates the relationship between industry-based labour demand changes and local labour market adjustment (population growth, the unemployment rate, the participation rate and wages).[18] He finds that an increase in labour demand (ie, employment) has a number of effects on the neighbourhood labour market. In particular, a positive employment shock reduces the neighbourhood unemployment rate, raises labour force participation, and encourages in-migration. However, the migration response is much weaker than the impact on neighbourhood unemployment and participation.

Maré and Timmins (2001) examine the link between migration flows and regional labour market variables using simple gravity models. They find only a weak relationship between labour market changes and regional migration flows. However, Maré and Timmins do not examine what the strength of the relationship implies for regional labour market adjustment, or how migration compares with other forms of adjustment, or the speed of adjustment.

Another New Zealand study that examines the link between migration and labour market outcomes is Morrison, Papps et al. (2000). This study emphasises the role of regional migration in increasing competition between firms for labour inputs, thus reducing local monopsony power. They find that "openness" of a region, in terms of the rate of inter-regional migration, has a significantly positive effect on wages (especially for more mobile higher skilled workers). While this study focuses on the influence of migration on wage-setting behaviour (via the local monopsony power), it does not examine regional labour market adjustment *per se*, and the role of migration.

---

[17] While it is not possible to do justice to all this work in a small literature review section, it is worth noting at least a few references (see Maré and Choy (2001) for a more comprehensive review of the literature).

[18] To overcome the bias arising from the endogeneity of employment growth in the model, Chapple (2000) resorted to using instrumental variables estimation rather than ordinary least-squares. The labour demand instrument used is based on what the employment growth rate would have been, given the industry composition of the neighbourhood. In other words, each one-digit sector is assumed to grow at the overall urban national rate.

# 4 Methodology

To model regional labour market adjustment in New Zealand, we use a Vector Autoregressive approach. The Vector Autoregressive (VAR) technique models the systematic co-variation amongst the selected set of variables, and uses this to get predicted time paths for all of the variables. In other words, the VAR technique examines the joint fluctuations of the selected set of variables over time. This then forms the baseline results.

In using the VAR technique, there are some specification issues that need to be addressed. Some of these are similarly treated across different studies, whereas others have been handled in different ways (see Maré and Choy, 2001). Our treatment of the key specification issues are discussed below.

The first specification issue is the choice of what variables to include in the VAR model. Given that our objective is to investigate regional labour market adjustment, and based on the discussions in Sections 1 and 2, the set of variables selected are the employment level, employment rate[19], the participation rate and wages. While we had quarterly data on the first three variables for 1985Q4 to 2001Q2, we have data on wages only from 1989Q1 to 2001Q1. We therefore estimate a 3-variable VAR model including employment, the employment rate and participation rate for the longer period. A 4-variable VAR model (the first 3 variables plus wages) is estimated for the shorter time period.[20] Regardless of the number of variables included in the model, all variables enter the model in logarithmic form.

We are interested in investigating how regional labour markets adjust to shocks which have an impact on the relative status and attractiveness of a regional labour market compared with other regions. Not all shocks will change the relative attractiveness of a region. For example, if a slow-down in the global economy affects all regions evenly, then the relative attractiveness of a particular region would not change and therefore we would not expect a net migration response. Consequently, the literature in this area makes a distinction between regional shocks and aggregate shocks. A regional shock is one that is felt either specifically in the one region or for which a region experiences a disproportionate impact of a shock that also affects other regions in the economy. Examples of regional shocks include the closure of a large factory in the region or a downturn in an industry in which the region is relatively specialised.

The time series data that we have for regional employment, employment rate, participation rate and wages contain both the impacts of region-specific shocks (which we are interested in) as well as aggregate shocks (which we would like to remove). The approach that we use, and which is generally used in the literature, is to specify the regional values of the variables as deviations from the national mean. This is often done by expressing regional values as being relative to the national counterpart. Therefore all regional variables entering our model are first divided by the variable's national value and then the natural logarithm is taken. This is shown below where $z_{j,t}$ represents a variable

19 The employment rate for a region is specified as the employment level for a region divided by the labour force of the region. While we use the employment rate in the model we present our results in terms of the unemployment rate. This does not cause any problems due to the unemployment rate being equivalent to one minus the employment rate.

20 We also re-estimate the 3-variable VAR model over the shorter period as a simple test as to whether the choice of sample influences the results significantly.

in the form that it enters the VAR model (that is, relative to the national figure and in logarithms). $Z_{j,t}$ is the value of the variable (eg, the employment rate) for region j at time t and $Z_{NZ,t}$ is the national (total) value (eg, the national employment rate) at time t.

$$z_{j,t} = \ln\left(\frac{Z_{j,t}}{Z_{NZ,t}}\right) \Rightarrow z_{j,t} \cong \ln Z_{j,t} - \ln Z_{NZ,t}$$

We use the above method for removing common shocks from the data, that is we use simple log differences (the log of the regional value minus the log of the national value) to obtain the main results presented. This is consistent with the approach taken in Blanchard and Katz (1992) and Debelle and Vickery (1998).

It is worth highlighting that the method of taking simple log differences is a special case of a more general approach known as β-differencing. β-differencing can be written as:

$$z_{j,t} = \ln Z_{j,t} - \beta_j \ln Z_{NZ,t}$$

Therefore, imposing the constraint that β equals 1 results in β-differencing being the same as simple differencing. β can be thought of as measuring how responsive a regional value of a variable is to changes in the national value of the variable. In other words, β is an elasticity. Imposing the constraint that β equals 1 is therefore imposing both unitary elasticity and also that β does not vary across regions (ie, it is always 1).

Different regions may well differ in terms of their elasticity to common shocks (changes to the national value of a variable) depending on a range of factors, such as how the industry composition within the region compares to the national average, or how labour force characteristics differ across regions. If the actual value of β for a region differs from our assumption of β=1, then our estimate of the region-specific component of a shock will be biased. The size of this bias will be determined by how far the actual β is from our assumed β of 1. Some studies estimate the appropriate elasticity (value of β) for each region, for each variable and use these estimated values to obtain relative to national values for each variable using β-differencing. We have also obtained β-difference values for employment, employment rate and participation rate variables.[21] While we find that for a number of regions the estimated β is significantly different from 1, when we use these β-difference values in a 3-variable VAR model, the results differ little when compared with those obtained from a model with the same structure but using simple difference values (see Section 6.3).

The second major specification issue concerns *how* the variables should enter the VAR model. Should they enter in levels or difference form? This depends on the time series properties of the variables of interest. In particular, one needs to test whether the underlying stochastic process generating the series is invariant with respect to time, that is, is it stationary? This is because variables entering the VAR should be stationary. The

[21] Results from the regressions run to obtain the β values for each region are summarised in Appendix C.

general approach to test for stationarity involves testing for the presence of a unit root.[22] We conduct a number of unit root tests prior to deciding how to enter our chosen variables in to the VAR model. More information on the testing procedure applied and results obtained are provided in Section 5. It is sufficient to note at this stage that if a variable is found to be integrated of order one (I(1) or non-stationary), it has to be differenced or detrended prior to entering the VAR model. On the other hand, a variable that is integrated of order zero (I(0) or stationary) enters the VAR model in levels.[23]

Whether a variable enters the VAR model in levels or first difference is an important issue because it influences the dynamic adjustment process for the variable. When a shock is introduced to the model, a variable that is modelled in levels will return to its pre-shock level in the long run. On the other hand, a shock to the model will have a permanent effect on the level of a variable that is modelled in first differences.

Most studies that use the VAR methodology to examine regional labour market adjustment have modelled employment and wages in first difference, while the employment rate and participation rate have entered the model in levels. We also follow this specification for comparability reasons. But as will be discussed in Section 5, determining whether a variable is stationary or not can be more of an art than a science and therefore we also investigate slightly different model specifications (as discussed in Section 6).

The specification of our 4-variable model can thus be written as:

$$\Delta em_{j,t} = \alpha_{1,j} + \sum_{s=1}^{n} \beta_{1,s} \Delta em_{j,t-s} + \sum_{s=1}^{n} \delta_{1,s} er_{j,t-s} + \sum_{s=1}^{n} \phi_{1,s} pr_{j,t-s} + \sum_{s=1}^{n} \gamma_{1,s} \Delta w_{j,t-s} + \varepsilon_{1,j,t}$$

$$er_{j,t} = \alpha_{2,j} + \sum_{s=0}^{n} \beta_{2,s} \Delta em_{j,t-s} + \sum_{s=1}^{n} \delta_{2,s} er_{j,t-s} + \sum_{s=1}^{n} \phi_{2,s} pr_{j,t-s} + \sum_{s=1}^{n} \gamma_{2,s} \Delta w_{j,t-s} + \varepsilon_{2,j,t}$$

$$pr_{j,t} = \alpha_{3,j} + \sum_{s=0}^{n} \beta_{3,s} \Delta em_{j,t-s} + \sum_{s=1}^{n} \delta_{3,s} er_{j,t-s} + \sum_{s=1}^{n} \phi_{3,s} pr_{j,t-s} + \sum_{s=1}^{n} \gamma_{3,s} \Delta w_{j,t-s} + \varepsilon_{3,j,t}$$

$$\Delta w_{j,t} = \alpha_{3,j} + \sum_{s=0}^{n} \beta_{4,s} \Delta em_{j,t-s} + \sum_{s=1}^{n} \delta_{4,s} er_{j,t-s} + \sum_{s=1}^{n} \phi_{4,s} pr_{j,t-s} + \sum_{s=1}^{n} \gamma_{4,s} \Delta w_{j,t-s} + \varepsilon_{4,j,t}$$

With the variables above being defined as follows:

---

22 A series has a unit root if a shock to the series has a permanent (or long lasting) effect. There are two ways to test for stationarity. The first is to use a descriptive approach. For example, one can calculate correlation coefficients (between current and past levels) and plot persistence diagrams to assess whether regions have experiences sustained differences in the labour market variables of interest. Some of these persistence measures have already been illustrated in Section 2. A second and more formal method to test whether a series is stationary or otherwise, is to use unit root tests, such as the Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) tests. One can also use panel unit root tests, such as the one suggested by Im, Pesaran, and Shin (1997), for dynamic heterogeneous panels. These are further discussed in Section 5.

23 A series that is non-stationary or I(1) in levels, but is stationary in first difference is called a difference-stationary series. Conversely, a non-stationary series that is stationary once one allows for a deterministic trend is regarded as trend-stationary. For an introduction to VAR models, as well as the concept of stationarity, see Enders (1995) and other standard (time series) econometric textbooks.

$$em_j = \ln\left(\frac{\text{employment in region j}}{\text{total employment}}\right)$$

$$er_j = \ln\left(\frac{\text{employment in region j/labour force in region j}}{\text{total employment /total labour force}}\right)$$

$$pr_j = \ln\left(\frac{\text{participation rate in region j}}{\text{national participation rate}}\right) = \ln\left(\frac{\text{labour force in region j/working age population in region j}}{\text{total labour force/total working age population}}\right)$$

$$w_j = \ln\left(\frac{\text{wage in region j}}{\text{national wage}}\right)$$

In the model, each variable is modelled as being dependent on both lagged values of itself as well as lagged values all other variables. Equations 2, 3 and 4 in the model differ from the first in that they also include a contemporaneous employment (difference) term. The inclusion of the contemporaneous employment term in the last three equations follows the approach adopted by Blanchard and Katz (1992) and others, which assumes that unexpected movements in employment within the quarter reflect shocks to labour demand, rather than labour supply. Shocks to employment demand can be introduced to the model through the $\varepsilon_{1j,t}$ term in the first equation. The inclusion of the contemporaneous terms therefore enables shocks to employment demand to have an immediate impact on the employment rate, participation rate and wage rate. These changes are then allowed to feed back into the employment equation in the following quarters.

To obtain our main set of results, we estimated each equation in the model separately using unweighted pooled ordinary least squares (OLS) with data for 12 regions. We chose a common lag length of 4 quarters for each equation.[24] We allowed for the constant terms (commonly known as fixed effects) in each equation to differ for each region. This allows for fixed differences between regions in each of the variables. The use of pooled OLS meant that we imposed the rather strong assumption that the coefficients for each region, for each equation are the same. That is, with the exception of the intercept terms, we only obtained one value for each coefficient and not a different value for each region. What this means is that our pooled results should be interpreted as implying how the average New Zealand region responds to a labour demand shock.

In addition to the main *pooled* results, we also estimate the model using data for each region separately. This provides us with a separate set of results for each region rather than a focus on the average region. The trade off is that the results are estimated using a sample that has about a twelfth of the variation of the sample used to obtain the pooled results.

Having estimated the model and obtained estimates for all of the coefficients included in the model, one can use the estimated coefficients to obtain a projected path for all the variables in the model. Imposing a shock means changing the starting point and gives a different projection. Plotting the difference between the two scenarios gives what is known as an impulse response function (IRF).[25] Therefore, we were able to obtain IRFs

---

[24] One could use a range of information criteria to select the optimal lag length for the VAR model. The Schwarz criterion picked a lag length of 2, but we felt that a higher lag is needed to capture a richer picture of the dynamics, as discussed further in Section 6.3.

[25] Throughout this paper, we consider the effects of a *negative* shock to employment. However, the IRF methodology used assumes that the responses to positive and negative employment shocks are symmetric. In other words, the responses observed are in fact the period average of responses to positive and negative

for all the variables in the model. This means that we initially obtained IRFs for $\Delta$em, er, pr and $\Delta$w. To assist with the presentation of results, as well as to ensure consistency with other studies using this technique, a number of simple transformations of the initial set of IRFs are required. The first of these is to convert the relative employment and wage IRFs to level form rather than first difference form. That is, we want to know what happens to relative employment and wages rather than the first difference of relative employment and wages.

Rather than knowing what happens to the natural logarithm of the relative participation rate we want to present our results in terms of what happens to the relative participation rate. This is achieved by way of a transformation.[26] Also, rather than displaying what happens to the log of the relative employment rate we present what happens to the relative unemployment rate (not in logs).[27]

In Section 1, we signalled that one channel through which labour market adjustment may occur is through the migration of people away from the affected region. Migration is not included explicitly in the model that we estimate. We are able to derive an estimate of the impact of a shock on migration by utilising the following identity:

$$d\ln(P) = d\ln(E) - d\ln\left(\frac{E}{L}\right) - d\ln\left(\frac{L}{P}\right)$$

Here P represents the working age population within a region, E is regional employment and L the region's labour force. All variables are expressed as a proportion of the national level. dln(P) is the change in the log of the working age population (or approximate percentage change in the working population) and captures the migration response.

This method essentially equates the movement in the working age population to the migration of workers, based on the assumption that most of the changes in the working age population are due to migration, rather than to natural population increases (which arguably are relatively stable, particularly if one uses quarterly data). In addition, the way in which we measure working age population (as relative to national) removes aggregate changes from the data. Therefore, even if there are factors other than regional migration that are influencing the working age population, these will not materially affect our results unless there are significant differences in the importance of these other factors across regions.[28]

---

shocks. Therefore, had we prescribed a positive shock, the results obtained would have been symmetric but in the opposite direction. We do not differentiate between anticipated and unanticipated shocks. A shock is a change that would not be predicted by the model. Therefore, we do not deal with the question of whether people in the regions anticipate the change or otherwise.

26 The transformation $d\left(\frac{L}{P}\right) = \frac{L}{P} d\left(\ln\frac{L}{P}\right)$ is used where L is the labour force and P represents the working age population (all variables expressed as relative to the national figure).

27 This is done by utilising the transformation $d\left(\frac{U}{L}\right) = -\frac{E}{L} d\left(\ln\frac{E}{L}\right)$ where U is unemployment and E employment (again both relative to national).

28 Auckland seems to be the exception. As discussed later, labour supply may have a large impact on employment figures since a disproportionately large number of immigrants first settle in that region.

For each final IRF, we obtain standard error confidence bands. These are obtained by way of a bootstrapping procedure, as employed by Runkle (1987).[29] This procedure involves the following steps. First, the VAR model is estimated and the coefficients and fitted residuals are saved. Secondly, a random draw is taken from this set of residuals, to be used with the saved coefficients in constructing an artificial sample of data. We repeat this process until we get 1000 such simulated data sets. Third, the VAR model is estimated using each set of simulated data, and thus, 1000 (simulated) impulse response functions are obtained. Finally, a one-standard error band for the IRF paths is inferred from this range of IRFs.

Before moving on to presenting the results obtained from the methodology discussed in this section, Section 5 provides details on the data used, as well as information on the unit root tests which influence our decision with regard to model specifications.

# 5 Data and empirical issues

This section describes the data used in the analysis, and analyses the univariate stochastic properties of the data prior to presenting the VAR results in Section 6.

## 5.1 Data sources

Data on employment levels, the unemployment rate, the labour force participation rate, and the working age population were obtained from the *Household Labour Force Survey* (HLFS). HLFS data are available on a quarterly basis[30] for 12 regions – Auckland, Bay of Plenty, Canterbury, Gisborne (including Hawke's Bay), Manawatu, Nelson (including West Coast, Tasman, and Marlborough), Northland, Otago, Southland, Taranaki, Waikato, and Wellington. This data set includes data from the fourth quarter in 1985 through to the second quarter in 2001.

The working-age population comprises usually resident, non-institutionalised, civilian population of New Zealand aged between 15 and 64 years. The labour force consists of members of the working-age population who during their survey reference week are classified as employed or unemployed. Therefore, by definition, persons not in the labour force include any person of working age who is neither employed nor unemployed.[31]

The employed are defined as those who had: (i) worked for one hour or more, for pay or profit, in the context of an employee/employer relationship or self employment; (ii) worked without pay for one hour or more in work which contributed directly to the operation of a farm, business or professional practice owned by a relative; or (iii) had a job but were not at work due to various reasons. Therefore, the employed include both full-time and part-time workers. The unemployed include all persons in the working-age population who during their reference week were without a paid job and were available for

---

29 The idea behind bootstrapping is to obtain an estimate of the small-sample distribution of the VAR coefficients (and hence, the IRFs) without assuming that the error terms are Gaussian (see Hamilton, 1994, p. 337).

30 The data used have not been seasonally adjusted in any way.

31 This includes retired persons, persons attending educational institutions, persons with personal or family responsibilities such as child care, unpaid house work, persons permanently unable to work due to physical or mental handicaps, and persons who were not actively seeking work.

work, and had actively sought work in the past four weeks, or had a new job to start within four weeks.[32]

Employment rates are calculated as the number of employed expressed as a percentage of the labour force of the respective region. Unemployment rates are calculated by dividing the number of unemployed by the number of people in the labour force. Labour force participation rates are obtained by dividing the labour force in each region by the respective working-age population (15-64 years).

The weighting scheme for the HLFS may create a contemporaneous correlation between regional labour market aggregates within a region. This is a potential concern for our study because it may bias our estimation. HLFS sample weights are calibrated to produce reliable estimates of national aggregates, not regional aggregates, and this weighting may create artificial variation at the regional level.[33]

Although this weighting issue is a potential source of bias in our estimates of the relationship between regional labour market aggregates, we consider that it is unlikely to be driving our results. As the example in footnote 33 illustrates, the bias depends on the size of regional differences in response rate changes and population composition, and on particular interactions of these differences with different labour market outcomes for subgroups. We believe that the overall effect of these factors is unlikely to account for a large proportion of the co-variation of the regional labour market aggregates that we model - regional shares of working age population and employment. The potential impact will be smaller still for specifications where we model *changes* in working age population and employment shares. Unfortunately, without independent regional population estimates, it is not possible to estimate directly the impact of this potential bias.

Wage data were sourced from the *Quarterly Employment Survey* (QES). QES provides employment data on a place-of-work basis ("supply" of jobs from employers, number of jobs filled and weekly paid hours), using 15 categories derived from the Australia and New Zealand Standard Industry Classification (ANZSIC). The data we used were available for the same 12 regions as those for which we have HLFS data. The data we have are for employers of more than 2.5 full-time equivalent (FTE) workers. Wage data are only available from the first quarter in 1989 to the first quarter in 2001.[34]

One issue that should be noted is that the QES was redesigned in 1999. Statistics New Zealand notes that sample surveys (such as the QES) need be redesigned periodically to ensure that the sample adequately reflects the contemporary composition of the population. Up to and including August 1999, the QES was based on a statistical sample that was first surveyed in February 1989. Over time the sample design of the QES

---

32 A person whose only job search method in the previous four weeks was to look at job advertisements in newspapers is not considered to be actively seeking work.

33 An example can illustrate the mechanism. If the response rate drops for young males, the weight for this group will increase, and regions that have a higher proportion of young males will show a relative increase in estimated working age population, as well as a higher estimate of labour market aggregates for which young males are over-represented (ie, unemployment). If the response rate for young males drops in a region with many young males, the weight attached to responding young males will increase nationally. Because the impact of the region-specific response rate change is absorbed nationally rather than regionally, the estimated number of young males in other regions will increase, even if there is no actual change in numbers there.

34 The QES and HLFS data are collected slightly differently. The QES is a snapshot taken in the middle of the relevant quarter (eg, 15th February for the March quarter). The HLFS survey data are collected throughout the quarter, so the figures we have are an average for the (same) quarter. Therefore, we will be comparing averages (HLFS) with midpoints (QES), which is not too bad.

became less effective in representing the business population for reasons such as some industry groups flourishing while others declined.

The changes that were made to the QES from (and including) the November 1999 quarter (1999Q4) included:

- changes to the contribution of different industries
- extending the coverage of the survey population to include small businesses (those with fewer than 2.5 FTEs)
- completion of the transition from the New Zealand Standard Industrial Classification (NZSIC) to the Australia and New Zealand Standard Industry Classification (ANZSIC).

To try and minimise the impact of these changes, Statistics New Zealand provided us with data for November 1999 onwards, that are based solely on firms with more than 2.5 FTEs. This means that all our QES data relate to firms with 2.5 plus FTEs.

Statistics New Zealand notes that the changes to the QES design may result in some discontinuity in the statistical time series. A preliminary time plot of relative log wages suggests that such a discontinuity may be present in the original data. This is corrected for by making the assumption that the change in the QES design results in a level-shift (ie, a jump). To remove this level shift, we impose the assumption that the change in wages between the 2nd and 3rd quarters in 1999 was zero (essentially removing the jump).[35]

In Section 2, we presented a number of time series plots and charts displaying the persistence of each variable over time. The next subsection builds on this by examining the stochastic properties of the variables of interest (univariate analysis), before we proceed to the VAR modelling stage (multivariate analysis).

## 5.2 Stochastic properties and univariate IRF estimates

As discussed in Section 4, whether a variable enters the VAR model in levels or first difference is an important issue because it influences the dynamic adjustment process for the variable. This section presents the unit root test results and the univariate impulse response functions (IRFs) on the constructed variables to determine how they should enter the VAR model.

To test for the stationarity of the variables, we performed unit root tests on the *panel* (ie, all the regions as a whole) and on *individual regions*.[36] The panel unit root test results are discussed first.

We adopted the panel unit root test recommended by Im, Pesaran and Shin (1997) (hereafter, IPS). Essentially, the IPS panel unit root test combines information on the stationarity or non-stationarity characteristics of the time series data for each region in the cross-section to give a conclusion for the entire panel. The test procedure and the results

---

[35] Although the results presented in this paper use the corrected wage data, earlier modelling efforts using the original wage data gave basically the same results. Therefore, our treatment for the discontinuity in the wage data does not in any way drive the results.

[36] The unit root test results presented in this paper are for the regional relative variables obtained through simple log-differencing (see Section 4).

are discussed in detail in Appendix D. The results indicate that all the variables (ie, logs of relative employment, relative employment rate, relative participation rate and relative wages) are stationary at the 1% significance level.

We also conducted the usual unit root tests, namely the Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) tests, for each regional relative variable. These two tests are well-known in the literature, and thus will not be discussed here.[37] The results are summarised in Table 1 below.

***Table 1 - Summary of ADF and PP unit root tests: Number of cases that are I(0)***

| Relative variable in logs | ADF | | PP | |
|---|---|---|---|---|
| | Without trend | With trend | Without trend | With trend |
| Employment | 8 | 9 | 8 | 9 |
| Employment rate | 10 | 9 | 12 | 11 |
| Participation rate | 9 | 9 | 10 | 12 |
| Wages | 3 | 6 | 7 | 11 |
| ∆Employment | 12 | 12 | 12 | 12 |
| ∆Employment rate | 12 | 12 | 12 | 12 |
| ∆Participation rate | 12 | 12 | 12 | 12 |
| ∆Wages | 12 | 12 | 12 | 12 |

Note: The table above shows the number of cases (or regions) out of 12 that are stationary or I(0) at least at the 10% significance level. The detailed test statistics for each regional relative variable and the lags used are reported in Appendix D.

It is clear from Table 1 that the results do not unanimously support or reject the unit root null hypothesis for all variables in levels. The results are particularly mixed for relative employment and relative wages. For example, according to the ADF test, log relative wages is stationary (at the 10% significance level) for only 3 out of 12 cases (without a trend term), or half the cases (with a trend term). Using the ADF test, the log of employment share is found to be stationary in eight (without trend) or nine (with trend) out of 12 regions.

If one places more emphasis on the *panel* unit root results, one would go ahead and model all the variables in levels. On the other hand, the individual regions' results are less clear-cut.[38] In most overseas studies of this type, employment share and wages are modelled in first difference. For comparison purposes, we will allow for employment share to enter the VAR in levels, as well as in first difference. As for wages, faced with the choice of a deterministic or stochastic trend, we find the latter more appealing and thus model wages in first difference.[39]

---

[37] The two unit root tests have their own respective strengths and weaknesses. For a discussion of their relative merits, see Hamilton (1994, pp. 515-516) and Pesaran and Pesaran (1997, p.213).

[38] The less clear-cut nature of the unit root results are not unique to our study. For example, Fredriksson (1995) found that the unit root null hypothesis was rejected in only 13 cases out of 24 for their (regular) employment variable, but modelled the variable in first difference.

[39] If wages *in fact* have a deterministic trend, one possible treatment would be to remove a linear trend before modelling the variable in levels. We tried this and it did not change the results significantly. Furthermore, as will be discussed in Section 6, wages do not play a major role in the adjustment process.

In addition to the unit root tests, one can consider how each variable adjusts when the particular variable is shocked. We model the univariate adjustment process for each variable for a typical region, allowing for four lags, as follows:

$$y_{j,t} = \alpha_{0,j} + \alpha_1 y_{j,t-1} + \alpha_2 y_{j,t-2} + \alpha_3 y_{j,t-3} + \alpha_4 y_{j,t-4} + \varepsilon_{j,t}$$

where y represents each individual variable entering the VAR in turn. Since the unit root test results were not strongly conclusive on the relative employment and relative wage variables, we estimate the equation above in both levels and first difference for these two variables.

The univariate process was estimated using pooled OLS techniques, which allowed for fixed effects for each region.[40] No form of weighting was used. From these estimated coefficients, we derive the associated univariate impulse response functions (IRFs), which give the response of the variable (in levels) to a 1% shock in $\varepsilon_{j,t}$ (ie, a shock to itself). These univariate IRFs give us a further insight in to the persistence, or otherwise, of the impact of a shock to each variable. We consider both aggregate and region-specific shocks, and so model the univariate adjustment processes of both absolute and relative variables separately. These univariate IRF results are shown in Figures 14 to 17.

Figure 14 shows that when employment is modelled in levels (panel (a)), the estimated effect of a shock is less persistent than when employment is modelled in first difference (panel (c)). Employment takes a longer time (about 35 quarters, as shown in Figure 14(a)) to reach its post-shock long run level when modelled in levels than is the case when it is modelled in first difference (approximately 7 quarters for relative employment to reach its (permanently lower) long run level, as shown in Figure 14(c)). The fact that the estimated adjustment process takes a long time in the levels case may point towards relative employment containing a unit root. This coupled with the mixed results from the individual regions' unit root tests results in us being less certain that the panel unit root test conclusions of stationarity are totally appropriate. Hence, our approach of allowing for two different model specifications; one of which treats relative employment as I(1) and the other as I(0).

40 That is, $\alpha_{0,j}$ is allowed to differ across regions.

***Figure 14 - Univariate impulse response of absolute and relative employment***

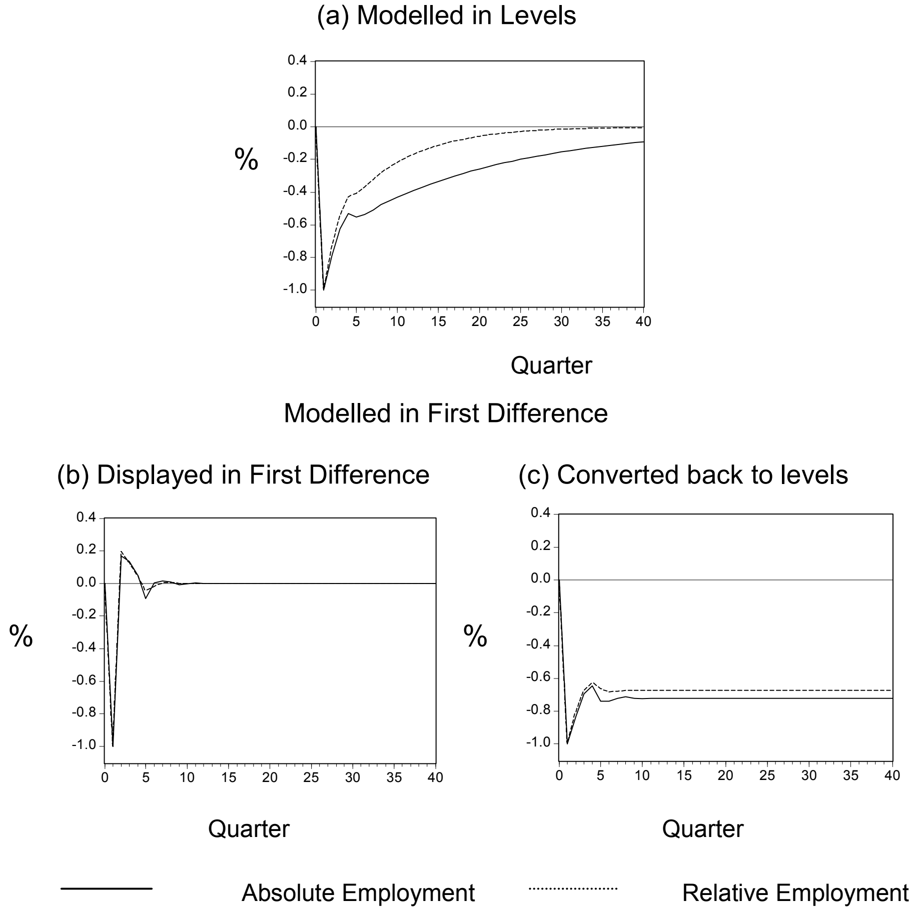


A second point to note from Figure 14 is that *common* or *aggregate* shocks have more lasting effects than *region-specific* shocks. This can be seen from Figures 14(a) and (c) where a shock to absolute employment (which includes both region-specific and aggregate effects) has either a larger or longer-lasting impact than is the case for a shock to relative employment. The intuition for this observation is that a macro shock affects the entire economy and therefore there is relatively little to be gained from shifting regions. On the other hand, when there is a negative shock to the particular region in which one lives, the relative attractiveness of other regions increases and therefore it may be possible to improve one's welfare by moving to a relatively more attractive region. Consequently, there is likely to be more adjustment in response to a region-specific shock and hence, the shock is more likely to be dissipated.

Figures 15(a) and (b) show the IRFs for employment rates and unemployment rates respectively. It is clear that the choice between modelling employment rates or unemployment rates is not of crucial importance, as shown by the virtually identical IRFs.[41] A second point to note is that regional employment rates return to their equilibrium faster after one strips off the effects of aggregate or common shocks. In other words, adjustment to a shock to the relative employment (or unemployment) rate is more rapid than is the case of a shock to the absolute employment (or unemployment) rate.

[41] It can be shown that $\ln[(1-UR_j)/(1-UR_{NZ})] \approx UR_{NZ} - UR_j$, whereas $\ln(UR_j/UR_{NZ})=\ln(UR_j) - \ln(UR_{NZ})$.

***Figure 15 - Univariate impulse response of absolute employment (unemployment) rates and relative employment (unemployment) rates***

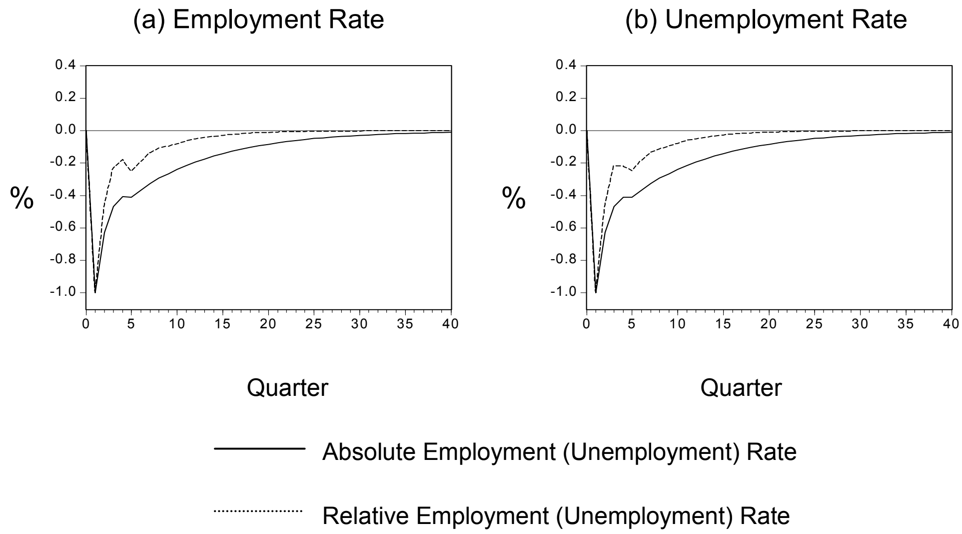


The finding that *common* or *aggregate* shocks to unemployment rates in New Zealand have more lasting effects than *region-specific* shocks can be compared with findings for other countries. In Europe, the effects of an aggregate shock to unemployment rates are more persistent than the effects of a region-specific shock. However, the European adjustment takes a lot longer than the New Zealand case.[42] In contrast, the impact of state-specific shocks to unemployment rates in the US is more persistent.

Figure 16 shows that the adjustment to a shock to *relative* participation rates is faster than the adjustment to a shock to *absolute* participation rates.

***Figure 16 - Univariate impulse response of absolute and relative participation rates***

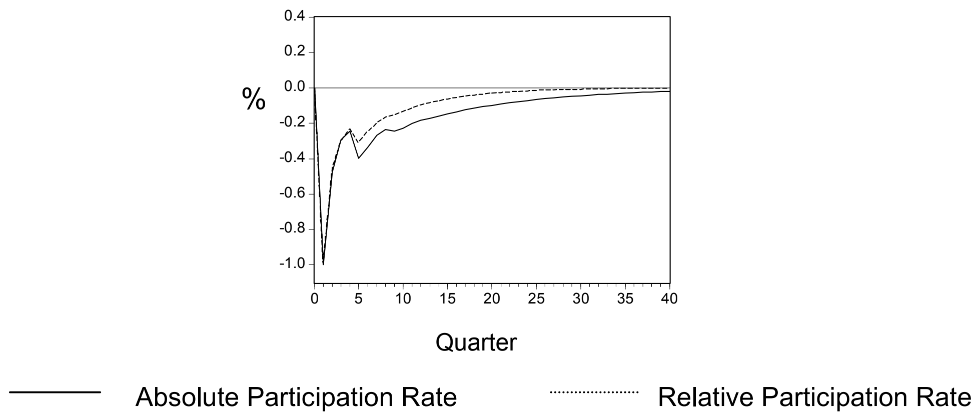


[42] For example, the impact of an aggregate shock to unemployment is fully reversed in New Zealand after about 10 years, whereas for Europe over half of the initial impact of the shock remains.

***Figure 17 - Univariate impulse response of the wage level and relative wages***

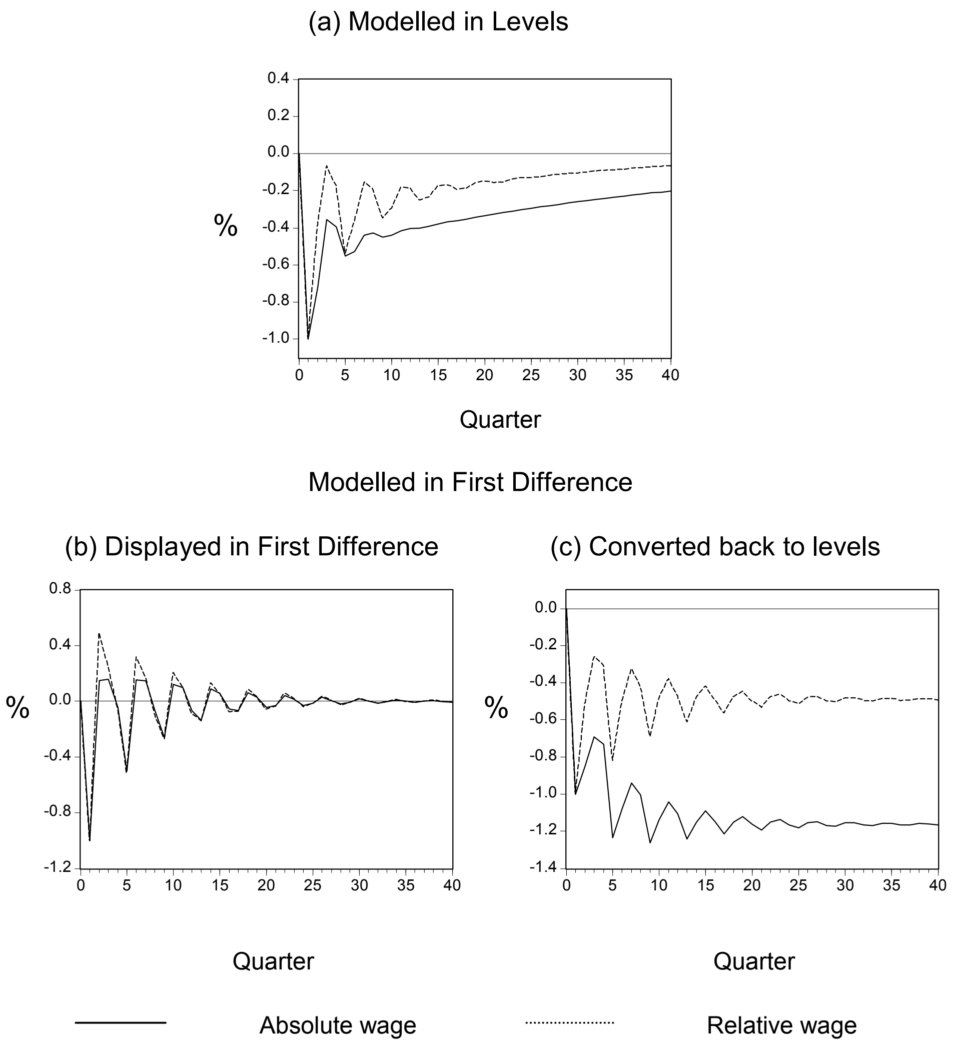


The univariate adjustment process for wages is shown in Figure 17. When wages are modelled in levels (panel (a)), the effect of a shock is less persistent than when wages are modelled in first difference (panel (c)). In terms of the speed of adjustment, when relative wages are modelled in levels, adjustment to a shock is very slow. Nearly 10 years after the shock, about 5% of the initial impact of the shock is still present (see Figure 17(a)). On the other hand, when relative wages are modelled in first difference, it takes approximately 7 to 8 years for relative wages to reach its (permanently lower) long run level (see Figure 17(c)). The apparent lack of, or very slow, adjustment in the levels case suggests that relative wages may contain a unit root. This coupled with the mixed results from the individual regions' unit root tests results, makes us more convinced that we should model relative wages in first difference. A second point to note from Figure 17 is that *common* or *aggregate* shocks have more lasting effects than *region-specific* shocks, similar to the other variables.

# 6 Empirical results

This section presents the results from estimating the VAR model using various model specifications. As mentioned in Section 4, as well as obtaining results for the average or typical region using pooled regression techniques, we also produced unique results for each region by estimating the 3-variable VAR model separately for each region. We begin this section by focusing on the results obtained using the pooled estimation technique. The availability of wage data for a shorter time span than the data availability for the other variables of interest resulted in us choosing to estimate both a 3-variable VAR model (excluding wages), for the longer time period, as well as a 4-variable model (including wages) for the shorter time period.

Due to the ambiguous nature of the unit root tests for the employment share variable, as well as the precedents set in earlier studies, we chose to estimate the 3-variable model using two different specifications. The first specification involved modelling employment share in first differences. We call this scenario our baseline case as it enables comparisons with many other international studies that have also treated employment share as I(1) (hence modelled in first differences). The second specification involved modelling employment share in levels. When expanding our model to the 4-variable case, we only present results for the specification in which employment share (as well as wages) are modelled in first differences.[43]

Table 2 summarises the different model specifications for which we present results in the remainder of this section.

***Table 2 - Summary of cases for which the results are presented***

| Set 1: Pooled results | |
|---|---|
| 3-variable VAR (1985Q4-2001Q2) | 1. Relative employment entering in first difference (the baseline results)<br>2. Relative employment entering in levels |
| 4-variable VAR (1989Q1-2001Q1) | 3 Relative employment and wages entering in first difference |
| Set 2: Individual regions' results | |
| 3-variable VAR (1985Q4-2001Q2) | 4. Relative employment entering in first difference<br>5. Relative employment entering in levels |

## 6.1 Results for the average region (pooled results)

In this subsection, we present the results for 3 different model specifications where the models were estimated using pooled OLS. These results can be interpreted as applying to the typical (or average) region.

43 As will be seen, the inclusion of wages in the model does little to change the adjustment processes of the other 3 variables. That is, wages do not seem to play a significant role in the adjustment process. Therefore, extending the 3-variable VAR model, in which employment share is modelled in levels, to a 4-variable model including wages, would provide little additional information over and above the 3-variable specification.

### 6.1.1 Results for a 3-variable VAR model (employment share in first difference)

The first model for which we present results is the 3-variable VAR model in which employment share is modelled in first differences. This forms what we call our baseline case.

***Table 3 - Summary of VAR Coefficients for the 3-variable VAR (employment share in first difference) (the baseline results)***

| Sum of coefficients on lags of: | Equation | | |
|---|---|---|---|
| | Change in Employment | Employment rate | Participation rate |
| Change in Employment | -0.39 | 0.16 | 0.37 |
| | [0.002] | [0.000] | [0.000] |
| Employment rate | -0.17 | 0.70 | 0.23 |
| | **[0.496]** | [0.000] | [0.027] |
| Participation rate | -0.31 | 0.02 | 0.81 |
| | [0.001] | **[0.121]** | [0.000] |
| Adjusted $R^2$ | 0.0728 | 0.6437 | 0.7667 |

Table 3 summarises the coefficients obtained when estimating the model. In the table, the first number displayed is the sum of the lagged coefficients of the particular variable, and the second number is the p-value for the F-test that the lagged variables are jointly significant. Those cases where the lags are jointly not significant have been bolded. More detailed results (ie, the coefficient on each lagged term) are provided in Appendix E.

While the VAR coefficients may be of interest to some, it is difficult to interpret them by themselves given that there is such a large number of coefficients. For example, in our baseline case (a 3-variable VAR model with 4 lags), there are 38 coefficients. To present the information contained in the large number of coefficients in a more accessible form, we follow convention by constructing IRFs from the estimated coefficients.

***Figure 18 - Impulse response to a negative employment shock (3-variable VAR) (1985:4-2001:2) (employment share in first difference)***

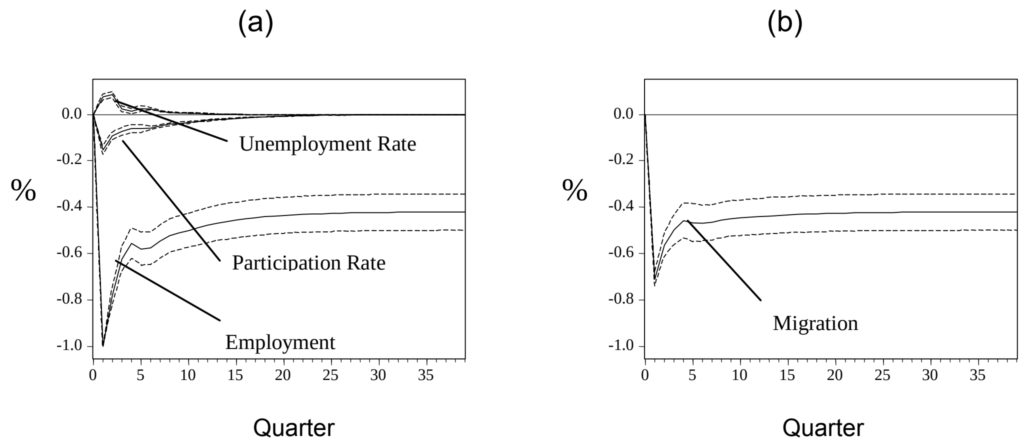


According to Figure 18(a), a 1% adverse shock to employment is associated with a rise in the unemployment rate of about 0.08 percentage points, which then falls back to its pre-shock or equilibrium level. The participation rate falls initially by about 0.16 percentage points before rising back towards its baseline value. However, relative employment does not return to its initial level. Instead, the long run impact of a negative 1% shock to employment is a permanent reduction in employment of 0.42%. The fact that employment does not return to its baseline value in the long run is an artefact of the way in which the variable has been modelled. In this baseline specification, employment has been treated as being I(1), that is, a shock can have permanent effects on relative employment.

Figure 18(b) plots the migration response to the negative 1% shock to employment. The migration response is computed from the time path of the employment, employment rate and participation rate variables that entered the VAR model.[44] As the figure shows, migration plays a substantial role in the adjustment process with the working age population falling by 0.7% in the period of the shock and then recovering partially in subsequent periods. The permanent reduction in employment mentioned earlier flows through to migration with the shock also having a permanent effect on the working age population (that is, a reduction in the working age population of 0.42%).

We can simplify things even further by showing what our IRFs imply in terms of the number of people.[45] This is shown in Table 4 below.

***Table 4 - The impact of a shock - a people story (employment share in first difference)***

| | Initial quarter | 1 year after shock | 4 years after shock | 6 years after shock |
|---|---|---|---|---|
| A: Net impact of change in employment on: | | | | |
| Working Age Population (migration) | 71 | 47 | 43 | 42 |
| Unemployment | 6 | 2 | 0 | 0 |
| Non-Labour Force Participants | 22 | 9 | 2 | 0 |
| Employment Response to Shock | 100 | 58 | 45 | 43 |
| B: Migration's impact on: | | | | |
| Unemployment | 6 | 4 | 3 | 3 |
| Non-Labour Force Participants | 27 | 18 | 16 | 16 |
| Employment | 71 | 47 | 43 | 42 |
| Migration Response to Shock | 104 | 69 | 63 | 62 |

Note: The sums of some of the figures in the table do not tally with the total due to all numbers being rounded to the nearest whole number.

In Table 4, we have assumed that the shock to employment shown in the IRFs is equivalent to a loss of 100 jobs in the region. Panel A of Table 4 suggests that when regional employment falls by 100, then of those people who lost their jobs, 71 leave the region in the quarter of the shock. The remaining 28 stay in the region, of which 22 exit the labour force and 6 become unemployed.[46] It is noted that we have been rather loose in the interpretation of these numbers. We do not actually know who are the ones

44 As discussed in Section 4, it is obtained from the identity: dln(P)=dln(E)-dln(E/LF)-dln(P/LF). P is the working age population, E is employment and LF the labour force.

45 IRF is the standard way of presenting the results in terms of responses to a shock. However, given that the different IRFs shown above have different denominators (that is, the unemployment rate response is a proportion of the labour force, while the participation rate and migration responses are specified as a proportion of the working age population), it is also worth interpreting them in terms of changes in numbers of people.

46 The reason 71, 22 and 6 do not sum to 100 is due to all numbers being rounded to the nearest whole number.

moving; whether they are the same people who directly lost their jobs as a result of the shock. All that the IRFs tell us is what was the net impact on region's employment (or whatever variable we are considering) as a result of the shock. Having said this, we shall continue to use these rather loose wordings for convenience.

The shock is felt not just by those who would have had a job in the absence of the shock. This can be seen from the overall impact of the shock on migration. Migration is measured as the change in working age population of the region. The total migration response (ie, people leaving the region due to the shock) is 104 people. In addition to the 71 people mentioned above, a further 33 people who were not employed immediately prior to the shock also leave the region. Using knowledge of the pre-shock unemployment and participation rates, we can separate the expected number of people (of the 33) who were unemployed from those who were not participating in the labour force prior to the shock. We find that 27 people who would have been classified as non-labour force participants leave, as do 6 people who would have been unemployed if they remained.

One year after the shock, the impact on the region's employment has diminished significantly due to a recovery in employment. Employment in the region is now 58 jobs lower than it would have been without the shock. Of the 58 people who do not have a job due to the shock, 47 have left the region, 9 have dropped out of the labour force and 2 remain in the region but are unemployed. The impact of the shock on migration has also diminished one year out, with the shock causing working age population to be 69 lower than what it was prior to the shock. What this means is that in the intervening year since the shock, a net 35 people (104 minus 69) have migrated into the region (these may or may not be the same individuals who left when the shock hit). The return of people to the region is mainly due to the partial recovery in employment.

The long run picture is reached about five to six years after the shock. In the long run, the region is left with 43 fewer jobs than would have been the case had the shock never occurred. So clearly, not all of the initial fall in employment is reversed over time. The long run migration response as a result of this is for the working age population to be 62 people fewer. This 62 people is made up of 42 people who would have jobs in the absence of the shock, 3 people who would have been unemployed had they remained in the region and 16 people who would not have been in the labour force had they stayed in the region.[47]

We can summarise the baseline results as follows. Most of the adjustment to a shock occurs in the first few quarters following the shock and is via regional migration. The next most important adjustment channel is participation rate changes (people dropping out of the labour force). Changes to unemployment play a small role. There is a permanent negative impact on the region's employment (about 40% of the initial shock), which is fully accounted for by a reduced working age population (the migration response).

In the next subsection, we look at the results for the specification in which the employment share variable is modelled in levels (instead of first difference).

---

[47] Again, the reason 3, 16 and 42 do not sum to 62 is due to all numbers being rounded to the nearest whole number.

### 6.1.2 Results for a 3-variable VAR model (employment share in levels)

Table 5 summarises the coefficients obtained when estimating the 3-variable VAR model with the employment share variable in levels. Again, only the sum of the coefficients for the lagged variables entering each equation are shown, with more detailed results being available in Appendix E.

***Table 5 - Summary of VAR coefficients for the 3-variable VAR (employment share in levels)***

| Sum of coefficients on lags of: | Equation | | |
|---|---|---|---|
| | Employment | Employment rate | Participation rate |
| Employment | 0.86 | 0.01 | 0.03 |
| | [0.000] | [0.000] | [0.000] |
| Employment rate | 0.04 | 0.67 | 0.16 |
| | **[0.532]** | [0.000] | **[0.239]** |
| Participation rate | -0.18 | 0.01 | 0.77 |
| | [0.064] | **[0.153]** | [0.000] |
| Adjusted $R^2$ | 0.9974 | 0.6448 | 0.7619 |

To aid in the interpretation of our results, IRFs based on the estimated coefficients are again presented.

***Figure 19 - Impulse response to a negative employment shock (3-variable VAR) (1985:4-2001:2) (employment share in levels)***


As Figure 19(a) shows, modelling employment share in levels (rather than differences) has little impact on the response of the unemployment and participation rates, when compared to the baseline case. Both rates move back towards their pre-shock levels. The difference comes down to the response in relative employment and even then, in the short run, the responses are very similar. Under this specification, employment share returns to its initial level in the long run, whereas the baseline results imply a permanent reduction in relative employment in the long run. Again, this is not surprising given that all the variables have been treated as I(0).

As Figure 19(b) shows, migration plays a substantial role in the adjustment process with the working age population falling by 0.7% in the period of the shock and then recovering in subsequent periods. Four quarters after the shock, the effect on the working age population has reduced to a fall of 0.30%. The fact that employment returns to its pre-shock level translates to the working age population also returning to its initial level. In other words, the negative shock does not have a permanent effect on the working age population. Complete adjustment takes about 10 years, although the migration response has fallen to about a seventh of its initial effect after about 13 or 14 quarters.

Similar to the baseline case, one can interpret the IRFs in terms of numbers of people (see Table 6). In the period in which the shock occurs, the numbers story is almost identical to that given above for the baseline case. The impact of the shock differs to the baseline case after the initial period and this is due to the employment IRF returning to its initial level. One year after the shock, the shock's impact on the region's employment has diminished significantly due to a recovery in employment. The recovery was sufficient to reclaim nearly 70% of the job losses that occurred at the time of the shock.

***Table 6 - The impact of a shock - a people story (employment level model)***

| | Initial quarter | 1 year after shock | 4 years after shock | 6 years after shock |
|---|---|---|---|---|
| A: Net impact of change in employment on: | | | | |
| Working Age Population (migration) | 71 | 30 | 7 | 3 |
| Unemployment | 7 | 1 | 0 | 0 |
| Non-Labour Force Participants | 23 | 8 | 1 | 0 |
| Employment Response to Shock | 100 | 39 | 8 | 3 |
| B: Migration's impact on: | | | | |
| Unemployment | 6 | 2 | 1 | 0 |
| Non-Labour Force Participants | 27 | 11 | 3 | 1 |
| Employment | 71 | 30 | 7 | 3 |
| Migration Response to Shock | 103 | 43 | 11 | 4 |

The impact of the shock on migration has also diminished one year out with the shock causing the working age population to be 43 lower than what it would have been in the absence of a shock. What this means is that in the intervening year since the shock, 60 people (in *net* terms) have migrated into the region. The return of people to the region is mainly due to the partial recovery in employment.

Six years after the shock, most of the shock's impact on the region has dissipated. Employment is 3 jobs fewer than what it would have been in the absence of the shock six years previously. And these 3 people, who would have had a job in the absence of the shock, are no longer living in the region. The overall impact on the working age population of the shock is a reduction of 4.

To summarise, the difference between the results here (with employment modelled in levels) and the baseline case (with employment modelled in first difference) is the *long run* employment track (and consequently, the *long run* migration response). However, there is a great deal of similarity between the two cases in the *short run*, with most of the adjustment occurring via migration, particularly in the first few quarters.

### 6.1.3 Results for a 4-variable VAR model (with wages)

We now examine whether allowing for wage adjustment changes the baseline results significantly. Table 7 summarises the estimated coefficients obtained when we extend the baseline case (3-variable VAR model with employment entering in first difference) by also including wages (in first differences) in the VAR model. It is noted, however, that adding wages comes at the expense of having a shorter sample period (1989Q1 to 2001Q1).

Incorporating wages into the model allows us to explore another potential adjustment channel for when regional labour markets are hit by employment demand shocks.

***Table 7 - Summary of VAR coefficients for the 4-variable VAR (employment share and relative wages in first difference)***

| Sum of coefficients on lags of: | Equation | | | |
|---|---|---|---|---|
| | Change in Employment | Employment rate | Participation rate | Change in Wages |
| Change in Employment | -0.39 | 0.08 | 0.12 | 0.00 |
| | [0.006] | [0.000] | [0.000] | [0.002] |
| Employment rate | -0.49 | 0.53 | 0.15 | 0.02 |
| | [0.012] | [0.000] | **[0.442]** | **[0.545]** |
| Participation rate | -0.43 | 0.01 | 0.74 | -0.02 |
| | [0.000] | [0.051] | [0.000] | **[0.244]** |
| Change in Wages | 1.23 | -0.18 | 0.08 | -0.95 |
| | **[0.108]** | **[0.093]** | **[0.437]** | [0.000] |
| Adjusted $R^2$ | 0.0887 | 0.7148 | 0.7541 | 0.4948 |

***Figure 20 - Impulse response to a negative employment shock (4-variable VAR) (1989:1-2001:1)***

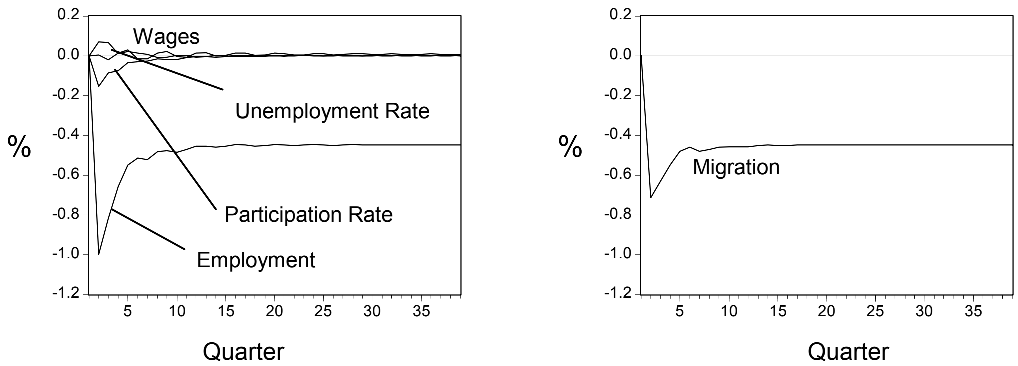


Figure 20 displays IRFs from the 4-variable VAR model. Comparing the IRFs shown in Figure 20 with those of the baseline case (Figure 18) shows that the inclusion of wages in the model does not alter the adjustment process of the other three variables much at all. The shock to employment also does not appear to have much impact on wages, with the wage IRF oscillating closely around the zero line.

To summarise the results thus far from the different specifications, adjustment seems to occur in response to region-specific shocks. In other words, the effects of the region-specific shock get dissipated across regions and over time. Migration appears to be the primary adjustment mechanism, and occurs rapidly as well. Changes in participation is the next most important adjustment mechanism, followed by unemployment changes. On the other hand, wage changes contribute very little to the adjustment process.

## 6.2 Results for individual regions

This section presents the IRFs for individual regions. These IRFs summarise the results obtained from estimating the VAR model for each region separately (see Figure 21). Given that wages are not a significant regional labour market adjustment mechanism, we shall only consider the 3-variable VAR model here. The discussion here focuses on how these IRFs for individual regions compare with the pooled baseline results (ie, employment share modelled in first difference).[48]

The unemployment response for all the regions are broadly similar to that of the pooled results, except that it is smaller in Bay of Plenty, Manawatu, Taranaki and Waikato. Similarly, the participation response is larger in Canterbury. As for the employment variable, the exceptions are Auckland, Gisborne and Otago. The most striking difference is the Auckland case. When employment is shocked negatively by 1%, it more than recovers. That is, employment actually rises above its initial or pre-shock level. We have not determined why this occurs, but it could be due to the fact that a large proportion of international migrants coming into New Zealand settle, at least initially, in Auckland. In Gisborne's case, the employment level initially drops but then overshoots slightly before returning to its new equilibrium level. The overshooting observation is a peculiar one. In Otago, the employment level does not seem to have stabilised even after 40 quarters. The peculiarities in the employment responses (mainly) flow on to the migration responses. Therefore, the exceptions, not surprisingly, are Auckland, Gisborne and Otago. For example, in Auckland, a negative 1% shock to employment results in an increase in the working age population in the long run.

However, it is worth highlighting that in the short-run, all the responses for each region appear to broadly match those of the pooled results. For example, the short run migration response (in the initial quarter) for all the regions lies between 0.58 and 0.9% of the working age population. The comparable estimate for the pooled VAR is about 0.75%. This consistency in the results for the short-run is expected because assumptions about stationarity and how the variables should enter the VAR influence the long-run evolutions more than the short run evolutions.

---

48 Appendix F shows the regional IRFs when employment is modelled in levels.

***Figure 21 - Individual regions' IRFs (baseline case)***

Vertical axis displays: percentage

Horizontal axis displays: quarter

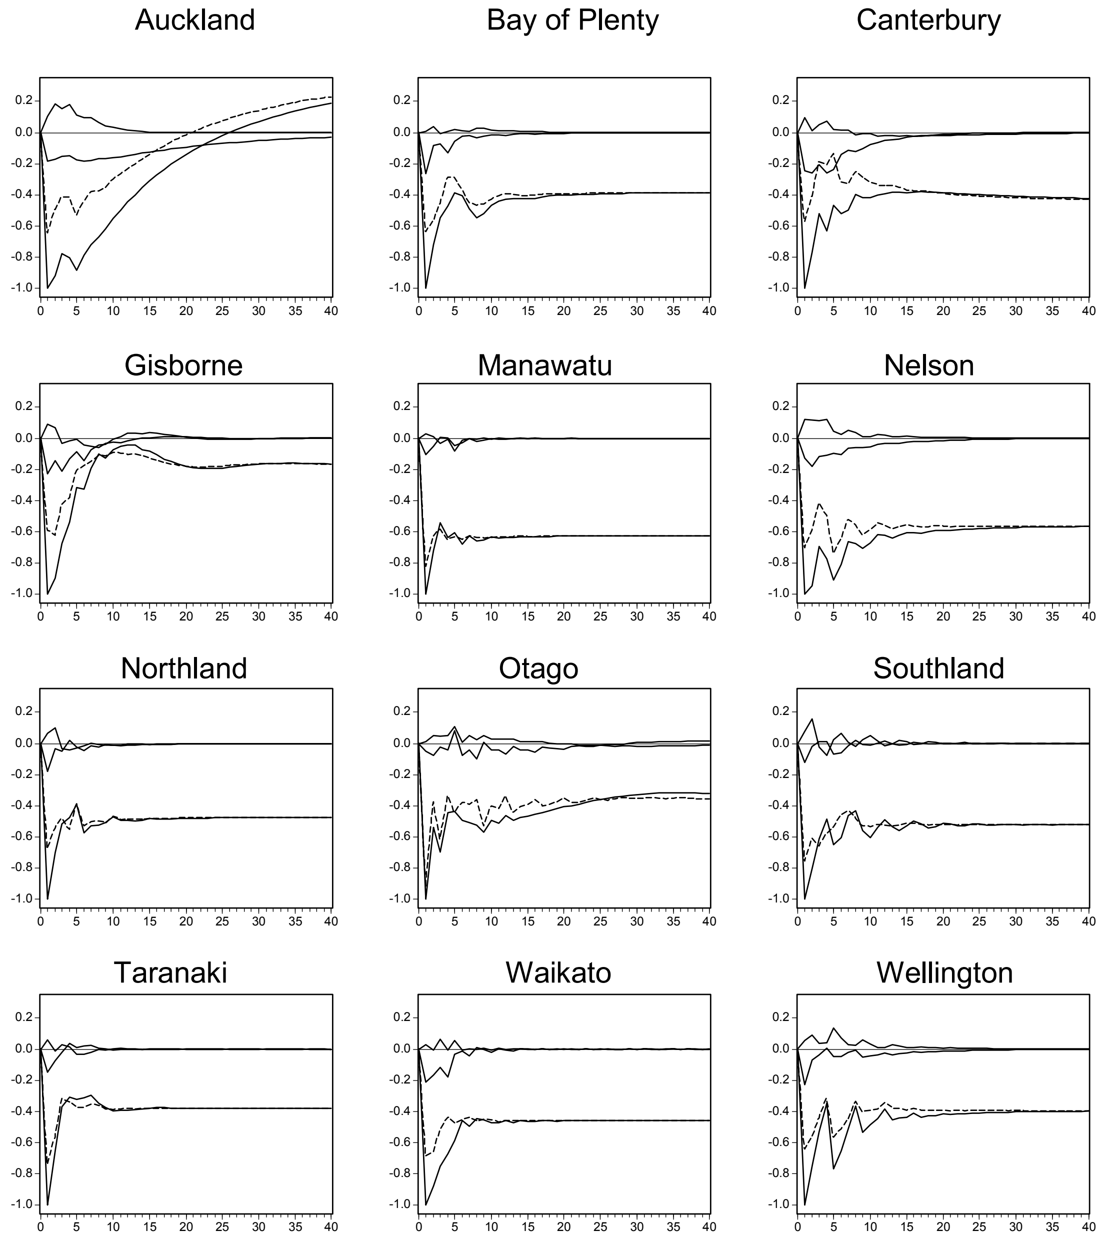


Note: For the three solid lines (from the left of each panel), the upper most line is the unemployment rate response, the middle line is the participation rate response, and the lowest line is the employment response. The dotted line displays the migration response.

## 6.3 Sensitivity analyses

We have found that migration is the primary adjustment channel with changes to participation and unemployment playing lesser roles. We have necessarily had to make a number of choices in the process of estimating the VAR model. To what extent do these choices drive our results? Put another way, how sensitive are our results to changes in the choices made? This section investigates if and how the baseline results change when alternative choices are made.

### Different lag lengths

When estimating VAR models, one has to select the optimal lag length of the right-hand-side terms. As explained in Section 5, the Schwarz criterion chose 2 lags as the optimal lag length for our VAR models, but we estimated the VAR models using 4 lags to capture a richer dynamics story. We experimented with the use of different lag lengths in the VAR model, ranging from 2 to 6 lags. The results using different VAR lag lengths do not appear to change much from the baseline case, except that the IRFs appear "smoother" when higher lags are included.

### Choice of derived variable

As discussed in section 4, the log of the working age population is a linear combination of log employment, log employment rate and log participation rate. Therefore we can independently estimate up to three of the variables of interest. The results shown thus far (for the 3-variable case) are based on a model that includes employment, employment rate and participation rate. The migration response (ie, change in working age population) is derived as a function of the responses for the modelled variables. Alternatively, one could have modelled working age population explicitly but exclude any one of the other three variables. Provided one makes stationarity assumptions that are consistent under both strategies, one can then compare the results to see if the choice of the derived variable drives the results. We experimented with substituting working age population for the employment rate in the model to be estimated. We find that the migration response is similar under both approaches, and consequently the relatively high migration response does not appear to be an artefact of treating migration as a residual (that is, attributing everything that is not captured by the three explicitly modelled variables to migration).

### Different sample periods

While we did not systematically test whether there are any structural breaks in the data set, we did estimate the same 3-variable VAR model over two different sub-periods, namely 1985Q4 to 2001Q2 (the full period) and 1989Q1 to 2000Q1 (the period over which wages data are available). Comparing the baseline results for the longer period with the results from the shorter period suggests that the change in the sample period does not change the baseline results much, as shown in Figure 22.[49]

[49] Figure 22 only shows that the IRFs for employment, unemployment rate and participation rate are not all that sensitive to a change in the sample period. While not shown, the migration response is just a function of the three IRFs and therefore, is also not sensitive to a change in the sample period.

*Figure 22 - A comparison of IRFs for different sample periods*

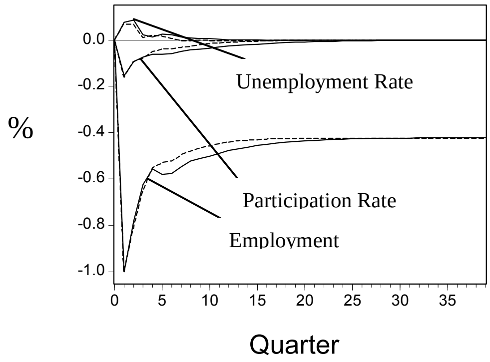


Note: Solid line displays: IRFs for the sample period 1985Q4 to 2001Q2.

Dotted line displays: IRFs for the sample period 1989Q1 to 2001Q1.

### Constructing regional relative variables differently (β-differencing vs. simple differencing)

As mentioned in Section 4, the unitary elasticity null hypothesis is rejected for some regions. Up to this point, we have shown only the results which are based on relative regional variables that were constructed using simple log differences (ie, $\beta=1$ for all regions). The rejection of unitary elasticity for some regions suggest that changes to aggregate variables may impact differentially on different regions. Therefore, it may have been more appropriate to have used β-differencing in constructing the regional relative variables.

However, as Figure 23 below shows, the baseline IRFs do not change significantly when we estimate the VAR model using β-differenced relative to national variables, instead of simple log differenced ones.

*Figure 23 - IRFs using alternative methods to construct regional relative variables*

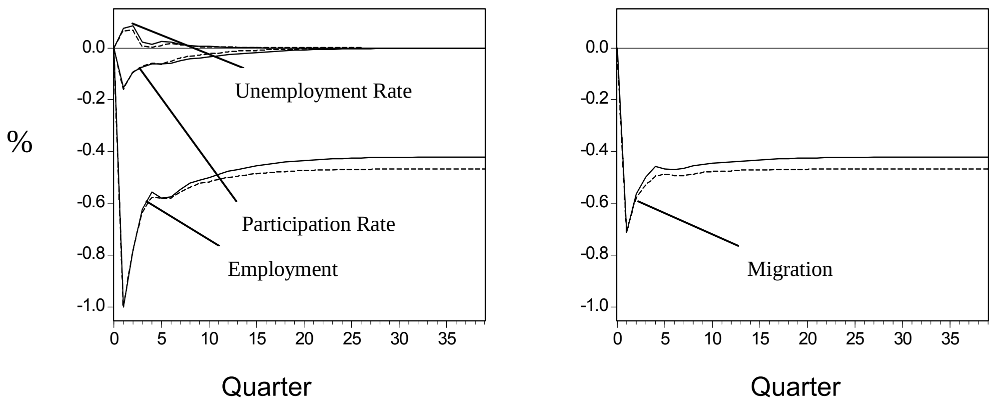


Note: Solid line displays: IRFs simple log differenced relative variables.
Dotted line displays: IRFs using β-differenced relative variables.

### Weighting/Different sub-groups of regions

Our main (pooled) results are based on unweighted regressions, where each region is given an equal weighting. This results in us interpreting our pooled results as applying to the average, or typical, region. An alternative way of estimating the model would be to use weighted regressions based on the population size of a region. The results obtained from using this alternative method would be interpreted as the adjustment process faced by the region in which the average person lives. We experimented with this by doing some weighted least squares (WLS) estimation. Preliminary analysis suggests that there is little change to the baseline results. Therefore, despite putting more weighting on the outlier Auckland case, the overall picture does not change much.

One might even suggest omitting Auckland explicitly. Although we have not done this, we suspect that this would not make much difference, given that altering the weighting from one-twelfth (ie, equal weighting) to a third (Auckland's appropriate weighting) does not change the baseline results significantly.

### Diagnostic tests

There may be a concern that there could still be information in the pooled residuals that have not been explicitly modelled, such as correlations across regions, across equations, and over time. The penalty of omitting such information from our model is that our results (ie, the point estimates and the standard errors) may not be accurate and believable. To investigate the importance of these issues we have performed a number of diagnostic tests on the residuals from the baseline case (ie, the pooled 3-variable VAR with employment entering in first difference).

*Autocorrelation*

We performed the Breusch-Godfrey test for first-order autocorrelation in our pooled residuals. There was no statistically significant evidence of a first-order autocorrelation in the pooled residuals and this suggests that the lag structure used in the VAR model is appropriate.[50]

*Cross-region heteroskedasticity*

Using the Breusch-Pagan-Godfrey test, we found that the residual variances are statistically different across regions. Therefore, the F-tests (that the lagged variables are jointly significant) in Tables 3, 5, and 7 (Section 6.1) as well as the standard errors and t-test statistics reported in Appendix E are White-heteroskedastic consistent.

*Cross-equation correlation*

In Section 4, we discussed the inclusion of a contemporaneous employment term in all equations except the (first) employment equation. The inclusion of this term essentially allows the shock to employment to impact on the other variables in the period of the shock. This is arguably appropriate given that preliminary checks of the correlation matrix of the residuals from the reduced-form model (ie, without the contemporaneous employment term in the second and third equations) suggest there are cross-equation correlations. These correlations are reduced with the inclusion of the contemporaneous employment term.[51]

---

50 The $\chi^2$ test statistic for the employment, employment rate and participation rate equations are 1.16, 0.98 and 3.60 respectively. Given that the critical value with one degree of freedom is 3.84, the null hypothesis of no first-order autocorrelation cannot be rejected for any of the equations.

51 Under the reduced-form model, the correlation coefficients are: $\rho_{E,ER} = 0.25$; $\rho_{E,PR} = 0.40$; and $\rho_{ER,PR} = -0.06$, where $\rho$ represents the correlation coefficient, E the employment equation, ER the employment rate equation and PR the participation rate equation. Under our "semi-reduced form" model, the corresponding correlation coefficients are: $\rho_{E,ER} = 0.00$; $\rho_{E,PR} = 0.00$; and $\rho_{ER,PR} = -0.18$.

*Cross-region correlation*

We have also performed some preliminary checks to see if there are contemporaneous correlations across regions. There is evidence of such correlations and this means that the residuals for each region obtained from the pooled VAR model are not independent of each other. This could *potentially* influence the standard error bands for the impulse response functions because the bootstrapping procedure used to obtain these bounds assumes that the residuals across regions are both independent and come from the same distribution (ie, we effectively have 12 independent draws). However, this does not appear to be the case. Preliminary investigations suggest that the standard error bands turn out to be very similar even when we allow for correlations across regions' residuals.

In summary, the baseline results seem to be reasonably robust. Adjustment seems to occur in response to region-specific shocks, or put another way, the effects of the shock get dissipated across regions and over time. Migration appears to be the primary adjustment mechanism, and occurs rapidly as well. Changes in participation are the next most important form of adjustment, followed by unemployment rate changes. Wages, in contrast, do very little of the adjustment.

# 7 Conclusions, policy implications and future research

## 7.1 Summary

Looking at a number of labour market variables over the study period, it is clear that regions experienced differing fortunes. There is some indication of persistence in most of the labour market variables we investigated. This finding of persistence could point to a lack of adjustment in New Zealand regional labour markets or signal that there are quite significant equilibrium differences in levels that do not necessarily require adjustment. This paper investigates the extent to which adjustment does occur in response to a region-specific shock, and through which channels. In defining the shock, the identifying assumption is that unexpected movements in employment reflect movements in labour demand.

We find that adjustment occurs primarily through migration and that this migration response is rapid. A region which experiences a negative employment shock equivalent to 100 job losses would, on average, experience a net migration loss of 104 people in the quarter of the shock. In the long run, our baseline results suggest the shock has a negative impact on the region's working age population of 62 people. Changes in participation and unemployment also play a role in the adjustment process, with participation rate changes being the more important of the two. Wages, in contrast, do very little of the adjustment.

However, our results capture only the responses of the labour force as a whole. This study does not investigate how the adjustment processes may differ for groups with different labour market characteristics eg, high and low skilled workers. For example, Morrison (1999a) found that lower skilled groups were relatively immobile and therefore their labour market outcomes were more influenced by regional shocks.

Before discussing how our results compare to those of other relevant studies, it is worth noting two points highlighted by the insignificance of wages in the adjustment process. First, regardless of the model specification chosen, the long run impact on the region's employment is lower than the initial impact of the shock. This means that in the periods after the shock, there must be some recovery (whether partial or full) in the region's employment level (ie, job creation). With wages not adjusting, what other forces would entice people to create jobs in the region? It must be based on some other costs of inputs for businesses. For example, housing and land prices could have fallen and thus, attract businesses to relocate from other regions, or spur new ventures within the region. However, this study does not model these other factors. Secondly, the lack of adjustment in wages could be viewed as surprising given the aims of the Employment Contracts Act (1991).

Changes in industrial structure can be seen as a cause of a shock. The more specialised the region is, the more likely it is to observe region-specific shocks. There has been significant change in New Zealand's industrial structure since 1985 (see Buckle, Haugh and Thomson, 2001). For example, the shares of primary, services and manufacturing have changed significantly. These changes in production shares are to a large extent also reflected in changes in employment shares. Indeed, one reason why the beta coefficient might vary across regions is likely to be differences in industrial structure. Similarly, differences in the composition of the labour force may also generate different regional responses to aggregate shocks.

How does the process of adjustment for New Zealand compare with that of other international studies? The significance of migration varies significantly across international studies. Fredriksson (1999) decomposes the contribution of three channels of adjustment (changes in unemployment and participation rates, and the migration response) in response to a shock in employment growth in Europe, US, Spain and Sweden. We build on this by providing similar estimates for Australia and New Zealand, as shown in the Figure 24.

In the European Union, changes in participation rates are the primary response to shocks in employment. Three quarters of the job losses are accounted for by a fall in participation rates in the first year (ie, the period of shock), and about 50% in the second year. In the US, the response of participation is much lower but it is compensated for by the large response of migration, which by the second year, accounts for about 54% of the adjustment. The main feature in the Spanish pattern is that the joint response of migration and participation to shock is relatively smaller than the US and overall European cases. A large proportion of the adjustment is borne by unemployment rate changes (39% one year after the shock).

In Australia, participation rate changes *initially* form the largest response – about 70% of the adjustment. One year after the shock, the migration response is the primary adjustment channel, accounting for 52% of the adjustment to the shock. Indeed, out-migration from a state that experienced a negative shock occurs slowly and steadily, with most of the migration taking place within four years. Conversely, migration has been the principal response to job destruction in Sweden for all periods. In the period of the shock, the migration response accounts for about 66% of adjustment to the shock, and this figure rises to 87% one year after the shock.

*Figure 24 - Decomposition of the response to a shock in employment growth (percentages of the change in log employment)*

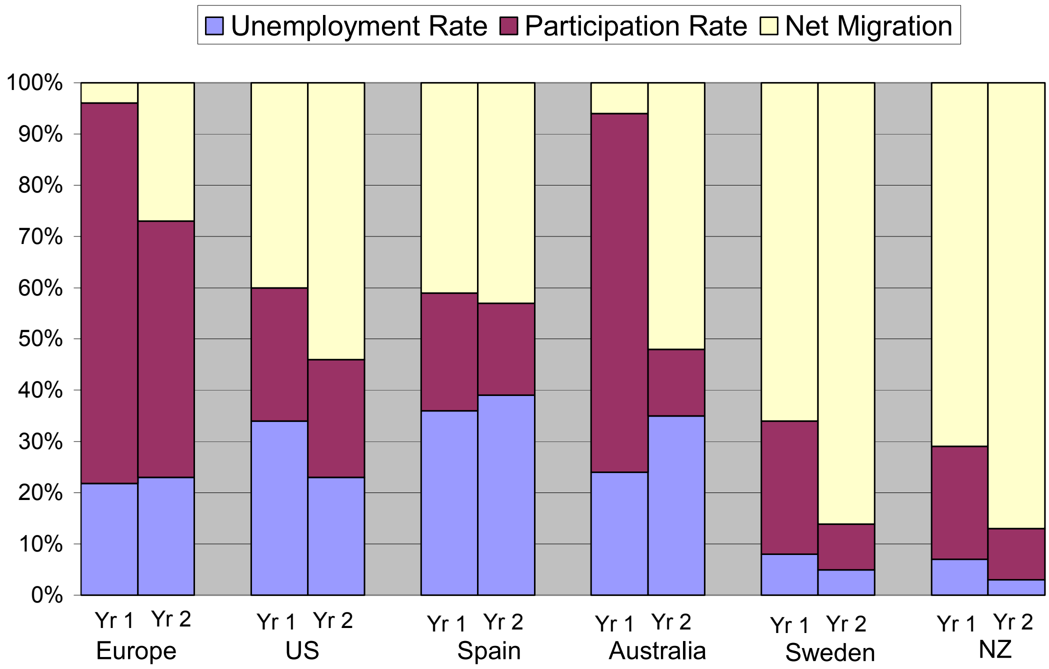


Notes: The figures for Europe, US, Spain and Sweden were obtained from Table 1 (p. 636) from Fredriksson (1999), who in turn sourced the measures for Europe, US, and Spain from Decressin and Fatas (1995), Blanchard and Katz (1992), and Jimeno and Bentolila (1995, 1998) respectively. We calculated the Australian estimates using figures from Debelle and Vickery (1999), and the New Zealand estimates are from our "baseline" impulse response f unctions, that is, using a 3-variable VAR with employment share modelled in first difference (for comparability purposes).

Where does New Zealand fit in? We find that New Zealand's results make it closer to the Swedish experience than that of other countries. A typical New Zealand region responds to a labour demand shock mainly via migration. In particular, the migration response accounts for 71% of the response to the shock in the period of the shock, and 87% one year later.

Given the comparisons above, one might ask if our results are surprising. On the one hand, New Zealand's migration response *can* appear extremely large, especially when compared to the European, US, Spanish and Australian cases. On the other hand, given the size of our regions, it is hardly surprising.[52] The regions in our study are much smaller in scale compared to regions as defined in the other studies. For example, the average population size of a region in Europe is over 6.8 million people, and the average population size of a US state is close to 5.3 million people. The average population size of a state or region in Australia and Spain is approximately 2.3 million people. On the other hand, the average population size for the New Zealand regions used in our study is 320,000 (about 1/17th the size of a US state).

Given this, one would expect to observe more inter-regional migration in New Zealand in response to negative region-specific shocks than is the case for Europe, US, Spain and Australia. This is because there are fewer alternative labour market opportunities *within* a (smaller) typical New Zealand region, and therefore people are more likely to have to look outside their own regions when there is a relative downturn in their regional labour market.

[52] This point was also noted in Frederiksson (1999).

In addition, the physical size of a region directly influences what is counted as a move. If a region is physically large, then people may move a substantial distance or even change urban centres without necessarily leaving the "region", and thus not contribute to migration figures. Within New Zealand there is a similarity in culture, language and regulations across regions. This is likely to lead to greater mobility than when regions differ in these important features. For example, there are significant cultural and language differences across many European regions.

The more comparable study in terms of the size of regions would be Frederiksson's (1999) Swedish study. In this study, the average population size for Swedish regions is about 370,000, not too far away from the 320,000 figure for New Zealand regions. Indeed, the variation in the size of regions is quite similar to ours. The Swedish regions range from the county of Stockholm, which has close to 1.8 million people (20% of the Swedish population) to the county of Gotland with just under 60,000 people (about 0.7%). If one accepts that the Swedish case is a more relevant benchmark, then our adjustment channels closely resemble theirs – migration is the primary adjustment mechanism.

In the previous discussion, we have been focusing on the mix of adjustment mechanisms. Another aspect of the dynamics is the speed of adjustment. When there is an adverse employment shock to the local labour market, how long does it take for the adjustment mechanisms to occur?

In the US, net migration plays a substantial role in the first year following an employment shock. After five to seven years, the employment response consists entirely of worker migration (Blanchard and Katz, 1992). In Europe, it takes about three years for the effect on the labour force participation rate and four years for the effect on the unemployment rate to disappear (Decressin and Fatas, 1995). In Australia, most of the migration takes place, on average, within four years. In particular, approximately one-third of the out-migration occurs within two years, roughly two-thirds of the net migration takes place within three years of the shock, and then the rate of out-migration flattens out. The process of adjustment is complete after seven years (Debelle and Vickery, 1999).

Regional labour market adjustment in Sweden and New Zealand have been comparatively rapid. In the case of New Zealand, most of the adjustment occurs within two to three years. The process of adjustment is complete after about five years. In the Swedish case, the employment response consists almost entirely of out-migration of workers two years after the shock.

It is also worth discussing briefly how our results compare with the findings of previous New Zealand studies. The relatively large migration response is consistent with the picture of high migration rates as found by Bedford, Goodwin et al. (1997) and others. Indeed, internal migration in New Zealand is high when compared to other OECD countries (see OECD, 2000, and Greenwood, 1997).

Our results appear to be different from those of Chapple (2000), where he finds that the migration response is much weaker than the impact on the area's unemployment and participation. However, it is worth noting that Chapple used a different methodology, a different definition of a labour demand shock (industry-based), a different data set (1981 census to 1996 census), and different geographical units of analysis (*urban* area units).

## 7.2 Implications for policy in New Zealand

Our pooled results describe what happens to a typical region in response to a region-specific shock. In at least two cases, government policies can be characterised as employment shocks. First, regional development policies aiming to stimulate employment in a particular region can be thought of as positive shocks to regional employment. If policymakers are primarily trying to improve employment outcomes of the existing regional population, then it is important to consider the size of the likely migration response to such a (positive policy) shock. If the migration response is high, then the policy could end up benefiting new entrants to the region rather than the initially targeted regional population or community. This may result in a lower proportion of the initially targeted population receiving the benefits envisaged, and/or the cost of providing assistance to the originally targeted population being higher than allocated.

Another example of government policies acting as region-specific shocks is when non-spatial policies have spatial effects. Some policies, although not directly aimed at particular regions, may nevertheless have regional implications. For example, changes in conservation policy such as banning native logging, although applied nation-wide, would have a greater impact on regions that have a proportionately high level of native logging. In this case, the ban would constitute a negative shock on labour demand for those regions, and thus induce an out-migration of people. Again, this may or may not be a good thing, but it is important that policymakers understand fully the dynamic effects of the policy proposed.

Similarly, industry assistance or industry protection policies such as import licensing, tariffs, export subsidies, can have differential regional impacts. Previous studies (as cited in Gibson, 1993) have found that industry protection policies may have been an important cause of internal migration patterns in the past, from provincial areas to metropolitan areas such as Auckland, Lower Hutt, Wellington and Christchurch because the latter were favoured by having a high concentration of protected, import substituting manufacturing.

While the results suggest that migration is the primary channel for regional labour market adjustment, we do not know *who* actually moves in response to the shock, and thus, it is not all that clear who (if anyone) policymakers should be trying to assist. Different policy interventions may be warranted depending on who actually moves. For example, it may predominantly be skilled workers who are able to migrate out of a region that is experiencing a negative shock, as in the Spanish case (see Mauro, Prasad et al. (1999)). Such circumstances may suggest that one should focus on assisting less mobile workers, for example by investing in education and training. However, as discussed further below, this study does not address this question.

Furthermore, it is difficult to know whether people should move more or less than they do. On the one hand, the more flexible the labour force is, the more efficient the economy would be in the allocation of its resources. On the other hand, there are large fixed costs and externalities associated with migration, which may mean that overall welfare would be improved if people moved less, and businesses moved more.[53] Therefore, when a region experiences a negative shock, it is not clear whether migration out of the depressed area should be encouraged, or whether firm migration should be encouraged into the depressed area.

---

[53] The externalities work in a number of directions: there is the benefit of agglomeration, the cost of overcrowding, the cost of emigration on those remaining in terms of property prices and local demand.

Perhaps a more challenging question is whether locally targeted policies can be effective in improving regional outcomes. The answer to this question depends on a range of factors, one of which is the time horizon considered. Our results suggest that the benefits of such policies dissipate quite rapidly across regions, and thus, benefits to the region could be only *short run* in nature. Whether the targeted region can benefit in the long run is unclear from our study. Our study focuses on the short run dynamics and so we had to make assumptions about the long run consequence – ie, whether employment is I(0) or I(1). To address the long run consequences of shocks, one would need to take a different approach.

Our findings complement those of other studies that investigated the responsiveness to aggregate shocks. This study finds that regional migration is an important adjustment channel in response to *regional* shocks. This finding, together with Aynsley's (2000) finding that international migration plays an important role in the adjustment process in response to *national* shocks, point towards the importance of labour migration as an adjustment mechanism.

The importance of worker mobility for New Zealand labour markets has implications for the currency union debate in New Zealand (for a discussion of the economic issues related to the currency union debate, see Coleman, 1999; Hargreaves and McDermott, 1999). One of the issues raised in the debate is how our economy would adjust to macroeconomic shocks if the exchange rate channel of adjustment were closed (a consequence of a currency union). McCaw and McDermott (2000) discuss in more detail these alternative channels, one of which is labour mobility.

In their study, McCaw and McDermott treat New Zealand as a state alongside Australian states and examine the role of migration compared to other labour market adjustment mechanisms using a similar methodology to that applied in our paper. Their results, which are consistent with Aynsley's (2000), indicate that migration is an important adjustment mechanism for New Zealand, and is a more significant part of the adjustment process than for Australian states. The authors suggest that if New Zealand were to join a currency union with Australia, trans-Tasman migration flows would increase and therefore, it would be important that labour flows remain unfettered under a currency union. In view of this point, a similar point was made by Gregory (2001), although the latter notes that the new social security rules in February 2001 could be interpreted as signs that trans-Tasman labour mobility may be more restricted in the future.

## 7.3 Future research

This section outlines four particular areas for future research which may provide further insights into the labour market adjustment processes of New Zealand regions.

The VAR model helps us understand the relative size of the migration response to regional labour market changes, compared to other adjustment channels, as well as the speed of adjustment. However, it does not help us understand *who* actually moves. For example, we do not know if those who move out of the region that experienced the negative shock are those who are more skilled or less skilled. Other dimensions that may be of interest include ethnicity and age. One option would be to estimate a different VAR model for different subgroups of the population, bearing in mind potential small sample size problems. A second option would be to use a spatial interaction model and link internal migration movements with labour market changes, controlling for things such as demographic composition, etc. (see Maré and Timmins, 2001). The spatial interaction model is something that we expect to be able to address when the 2001 census data become available.

This study focuses on labour market adjustment for a *typical* region or an *average* region. This is especially true for the pooled results. However, even the IRFs for individual regions show the response for the region as a whole. The adjustment channels may well be different for different towns within the region. This suggests that one could perhaps examine lower levels of aggregation, for example Territorial Local Authorities, or even local labour markets. However, at this stage, there are still questions about the availability of such data.

Another option for uncovering the labour market adjustment channels at lower levels of aggregation would be to conduct case studies. Case studies could focus on the way that labour market adjustment affected particular communities, particularly capturing more qualitative information. However, we might not be able to generalise the results identified for New Zealand as a whole. In other words, case studies can capture a lot more on a great deal less because they focus on selected regions or towns.

While this study goes some way to answer the question of *why* people move, it is only a partial answer. This paper shows that people move because of differing employment opportunities, and not because of differing wage levels. However, this study does not consider other factors for migration such as family reasons, and a host of other non-labour market reasons. The spatial interaction model mentioned above can help us understand a bit more about why people move. It is an indirect way understanding what motivates movers. A more direct way to deduce the motives of movers would be to conduct surveys. There is currently work underway in this area. For example, Richard Bedford and others at the University of Waikato have just sent out survey questionnaires to people who have either moved into or moved out of the Western Bay of Plenty area, asking about, *inter alia*, their reasons for moving. These migrants were identified through the change of address notification lodged with New Zealand Post.

A fourth area that one could develop further is international linkages. This study has not modelled international migration explicitly, and changes in the working age population as a result of the shock are attributed to regional migration. More importantly, international migrants (particularly, immigrants) are not symmetrically distributed across the regions. Therefore, the migration responses across different regions should be interpreted with some caution.[54] It is also worth highlighting that this study focuses on region-specific shifts, and thus the working age population (WAP) measure for each region is taken as relative to the national WAP. This means that this study abstracts from changes to the relative attractiveness of New Zealand as a destination for potential international migrants.

---

[54] For example, one could argue that, in response to immigrants settling in Auckland, wages in Auckland should decline and unemployment there should rise. This would in turn discourage in-migration from other regions, and thus maintain the equality of long-term employment growth rates between the regions. Alternatively, local residents and immigrant workers might be complements rather than substitutes, in which case these implications do not follow.

# References


Aynsley, M. (2000) "Labour Market Adjustment in Australia and New Zealand." Department of Economics, University of Melbourne.

Banerjee, A. (1999) “Panel Data Unit Roots and Cointegration”, *Oxford Bulletin of Economics and Statistics*, Special Issue: 607-629.

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* (Kalamazoo, MI, W. E. Upjohn Institute for Employment Research).

Bedford, R. D., J. E. Goodwin, E. Ho, and J. Lidgard (1997) "Migration in New Zealand 1986-1996: a Regional Perspective." *New Zealand Journal of Geography:*16-43.

Blanchard, O. J. and L. F. Katz (1992) "Regional Evolutions." *Brookings Papers on Economic Activity* 1:1-75.

Bushnell, P. and W. K. Choy (2001) "Go West, Young Man, Go West!” *People and Place* 9(3): 61-81.

Chapple, S. (2000) "Spatial inequality and labour market adjustment in urban neighbourhoods." 5th draft. New Zealand Department of Labour, *mimeo*.

Coleman, A. (1999) "Economic Integration and Monetary Union." Treasury Working Paper 99/6, New Zealand Treasury. <http://www.treasury.govt.nz/workingpapers/1999/99-6.asp>

Debelle, G. and J. Vickery (1998) "Labour Market Adjustment: Evidence on Interstate Labour Mobility." Research Discussion Paper 9801, Reserve Bank of Australia.

Debelle, G. and J. Vickery (1999) "Labour Market Adjustment: Evidence on Interstate Labour Mobility." *Australian Economic Review* 32(3):249-263.

Decressin, J. and A. Fatas (1995) "Regional Labour Market Dynamics in Europe." *European Economic Review* 39: 1627-1655.

Enders, W. (1995) *Applied Econometric Time Series*. (New York, John Wiley & Sons, Inc).

Fredriksson, P. (1999) "The Dynamics of Regional Labor Markets and Active Labor Market Policy: Swedish Evidence." *Oxford Economic Papers* 51(4):623-648.

Gibson, J. (1993) "Regional Consequences of New Zealand Industry Assistance Policies." *New Zealand Journal of Geography* April:2-6.

Glass, H. and W. K. Choy (2001) "Brain Drain or Brain Exchange?" Treasury Working Paper 01/22, New Zealand Treasury. <http://www.treasury.govt.nz/workingpapers/2001/01-22.asp>

Gregory, B. (2001) "Does Moving Away from Labour Market Integration Matter? Implications of Recent Changes in Australian and New Zealand Immigration Policy." Paper presented at the symposium "Off the Map in the Global Economy? Implications of economic geography for small and medium-size economies at peripheral locations", Victoria University of Wellington.

Goodwin, J. E. and R. D. Bedford (1997) "Inter-regional Migration in New Zealand, 1986-1996: A Preliminary Analysis." Briefing Paper for the Population Conference: Communities, People and Growth, No. 4, Wellington NZ, 12 to 14 November.

Greenwood, M. J. (1997) "Internal Migration in Developed Countries." In M. R. Rosenzweig and O. Stark eds *Handbook of Population and Family Economics* 1B. (Amsterdam, Elsevier Science):647-720.

Hamilton, J. D. (1994) *Time Series Analysis*. (New Jersey: Princeton University Press.)

Hargreaves, D. and C. J. McDermott (1999) "Issues relating to optimal currency areas: theory and implications for New Zealand." *Reserve Bank of New Zealand Bulletin* 62:16-29.

Im, K., H. Pesaran and Y. Shin (1997) "Testing for Unit Roots in Heterogeneous Panels." Working Paper No. 9526, Department of Applied Economics, Cambridge University.

Jimeno, J. F. and S. Bentolila (1998) "Regional unemployment persistence (Spain, 1976-1994)." *Labour Economics* 5(1): 25-51.

MacKinnon, J. G. (1991) “Critical Values for Cointegration Tests”,Chapter 13 in Engle, R. F. and C. W. J. Granger eds *Long-Run Economic Relationships: Readings in Cointegration* (Oxford University Press, Oxford).

Maloney, T. and J. Savage (1996). Labour Markets and Policy. In: Silverstone, B., A. Bollard and R. Lattimore eds (1996) *A Study of Economic Reform: The Case of New Zealand.* (Elsevier Science: Amsterdam): 173-214.

Maré, D. C. and W. K. Choy (2001) "Regional Labour Market Adjustment and the Movements of People: A Review." Treasury Working Paper 01/08, New Zealand Treasury. <http://www.treasury.govt.nz/workingpapers/2001/01-8.asp>

Maré, D. C. and J. Timmins (2001) "Internal Migration and Regional Adjustment: Some Preliminary Issues." Paper presented at the 9th Conference on Labour, Employment and Work, Institute of Geography, Victoria University of Wellington.

Mauro, P., E. Prasad, et al (1999) "Perspectives on Regional Unemployment in Europe." IMF Occasional Paper No. 177, Washington DC, International Monetary Fund.

McCaw, S. and J. McDermott (2000) "How New Zealand adjusts to macroeconomic shocks: the implications for joining a currency area." *Reserve Bank of New Zealand Bulletin* 63:35-51.

Morrison, P. S. (1999a) "Regional Labour Markets in New Zealand, 1991." Occasional Paper 1999/2. Labour Market Policy Group, New Zealand Department of Labour.

Morrison, P. S. (1999b) "Unemployment and Labour Force Participation in the 1990s: The New Zealand Experience." In: K. O'Connor, K. ed (1999) *Housing and Jobs in Cities and Regions: Research in Honour of Chris Maher*. (University of Queensland Press for the Australian Housing and Urban Research Institute):223-236.

Morrison, P. S. (2001). Employment. In: Willis, R.ed (2001) "New Zealand in the Nineties." A special issue of *Asia Pacific Viewpoint* 42(1):85-106.

Morrison, P. S., K. L. Papps and J. Poot (2000) "Local Labour Markets, Migration and Wage Determination: Theory and Evidence for the Wage Curve in New Zealand." GSBGM Working Paper 1/00, Graduate School of Business and Government Management, Victoria University of Wellington.

Newell, J. and K. L. Papps (2001) "Identifying Functional Labour Market Areas in New Zealand: A Reconnaissance Study using Travel-to-Work Data." Occasional Paper 2001/6, Labour Market Policy Group, New Zealand Department of Labour. <http://www.lmpg.govt.nz/publications/op2001-6.pdf>

NZPC (1989) *Diversity and Change: Regional Populations in New Zealand*. (New Zealand Planning Council.)

OECD (2000) Disparities in Regional Labour Markets. *OECD Employment Outlook:* 31-78.

Pesaran, M. H. and B. Pesaran (1997) *Working with Microfit 4.0: Interactive Econometric Analysis*. (Oxford, Oxford University Press.)

Poot, J. (1995) “Do Borders Matter? A Model of Interregional Migration in Australasia.” *Australasian Journal of Regional Studies* 1(2):159-182.

Runkle, D. E. (1987) “Vector Autoregressions and Reality.” *Journal of Business & Economic Statistics* 5(4): 437-454.

White, H. (1980) “A Heteroscedasticity Consistent Covariance Matrix Estimator and a Direct Test of Heteroscedasticity.” *Econometrica* 48:817-818.

# Appendix A: Plots of the relative variables entering the VAR (using simple log differences)

Appendix A plots the variables in their various forms entering the VAR. It is noted that these variables are defined relative to national (simple log-differences) and are in logarithms. It is noted that for all the figures in this appendix, the scales on the vertical axis differ for each region. The reader may find it useful to look at these figures in conjunction with the unit root test results in Appendix D.

### *Figure A1 - Log of employment share*

Vertical axis displays: natural logarithm of region's employment share

Horizontal axis displays: year

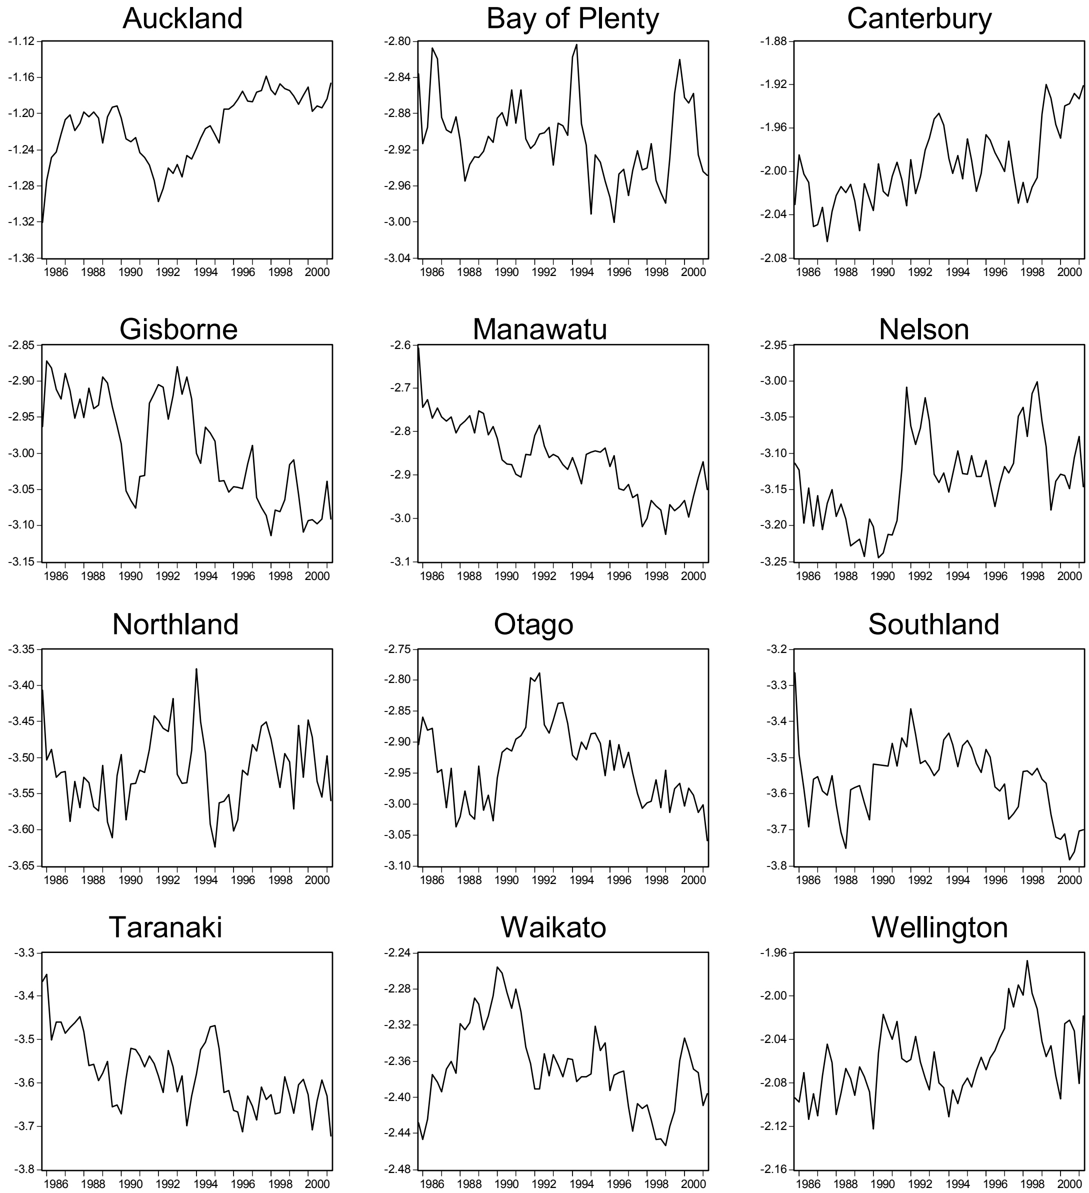

### *Figure A2 - Log of employment share in first difference*

Vertical axis displays: natural logarithm of region's employment share in first difference

Horizontal axis displays: year

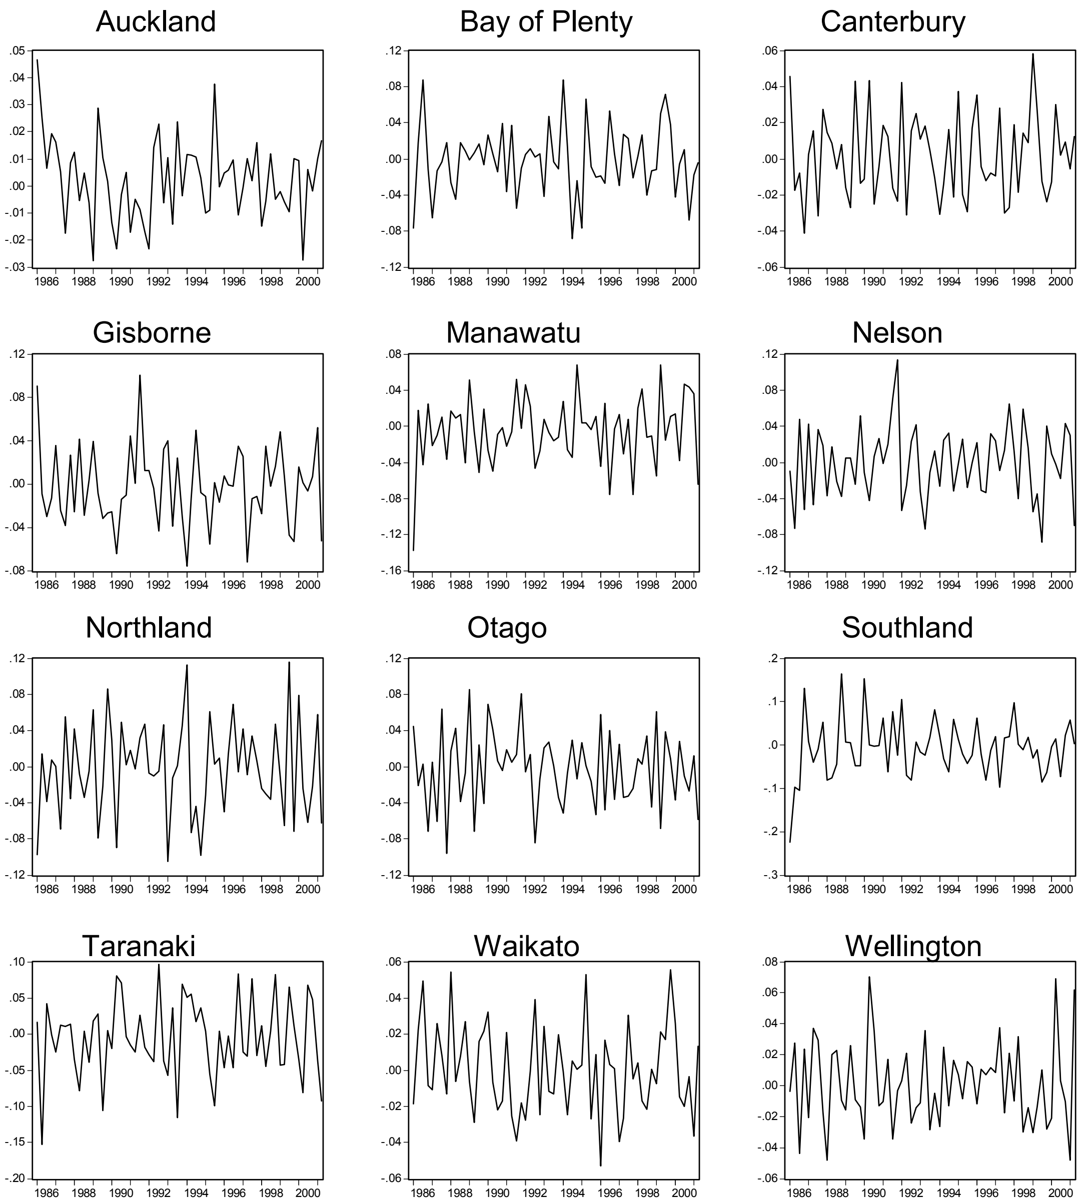

## *Figure A3 - Log of relative employment rate*

Vertical axis displays: natural logarithm of region's relative employment rate

Horizontal axis displays: year

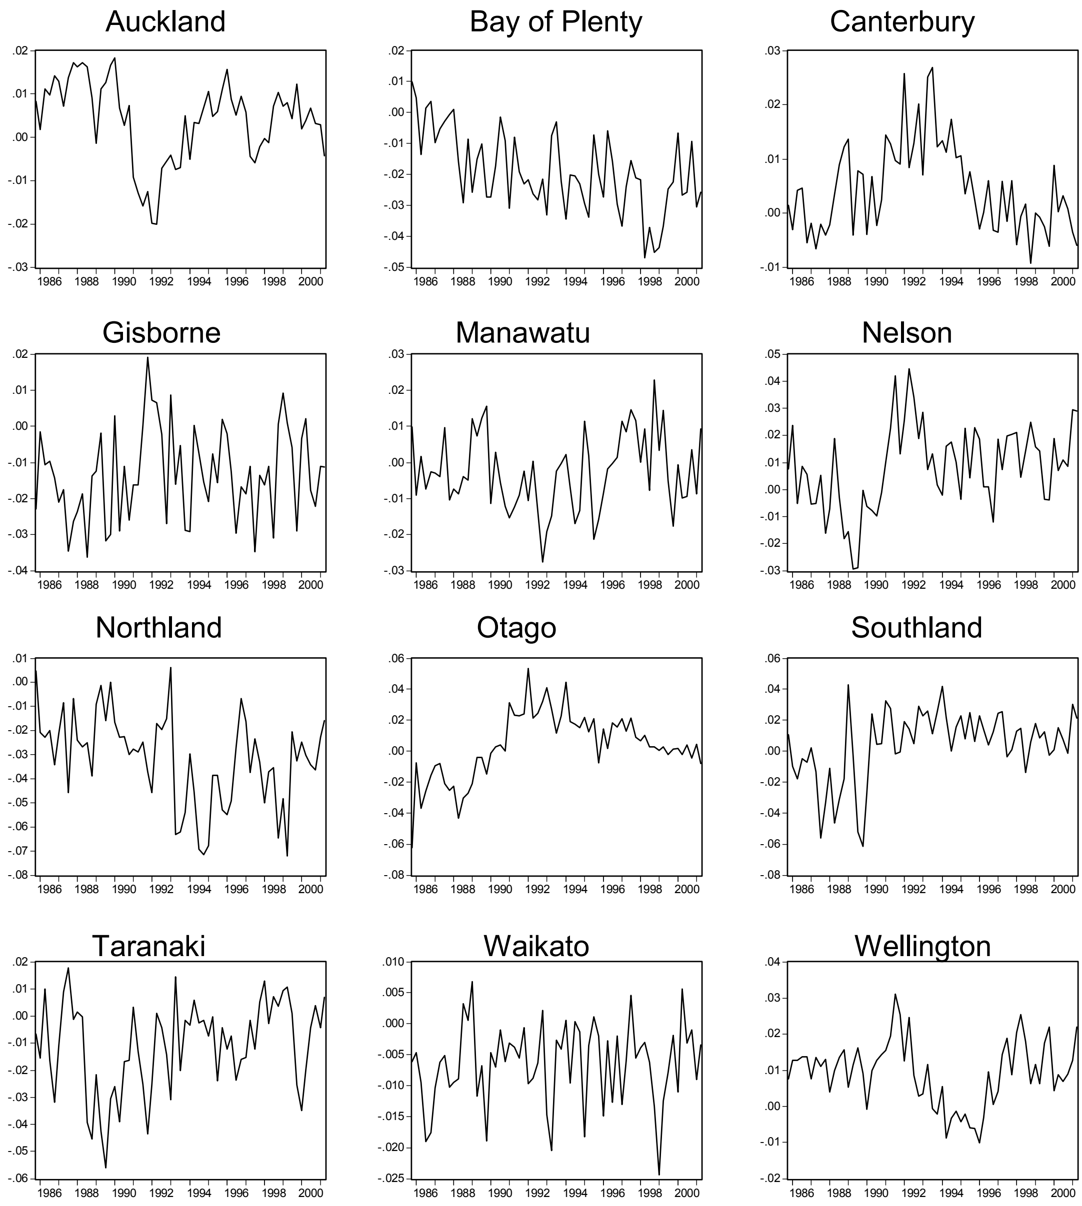

## *Figure A4 - Log of relative participation rates*

Vertical axis displays: natural logarithm of region's relative participation rate
Horizontal axis displays: year

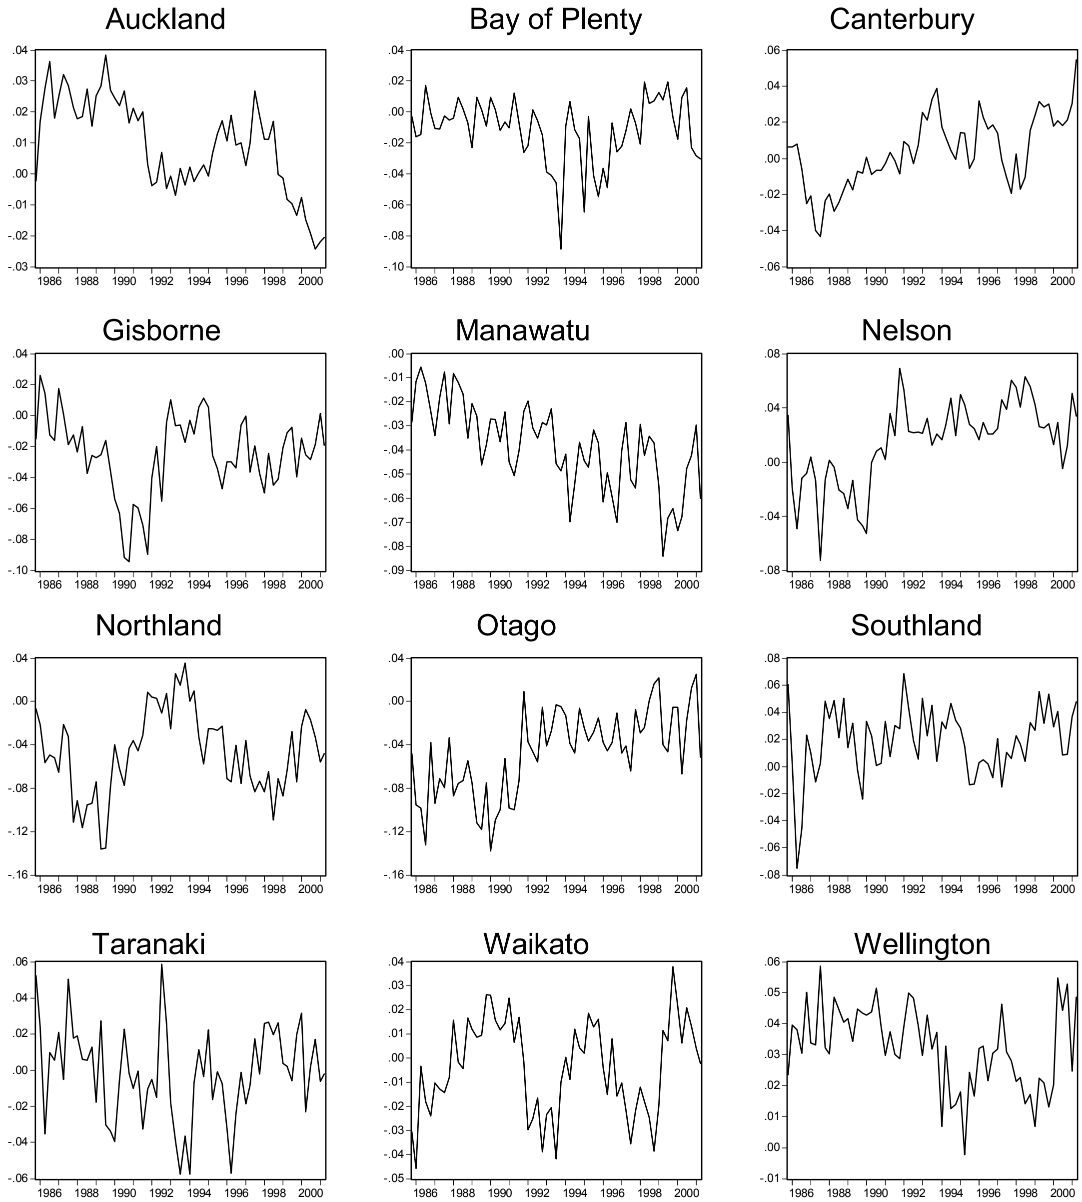

*Figure A5 - Log of relative wages in first difference*

Vertical axis displays: natural logarithm of region's relative average hourly wage
Horizontal axis displays: year

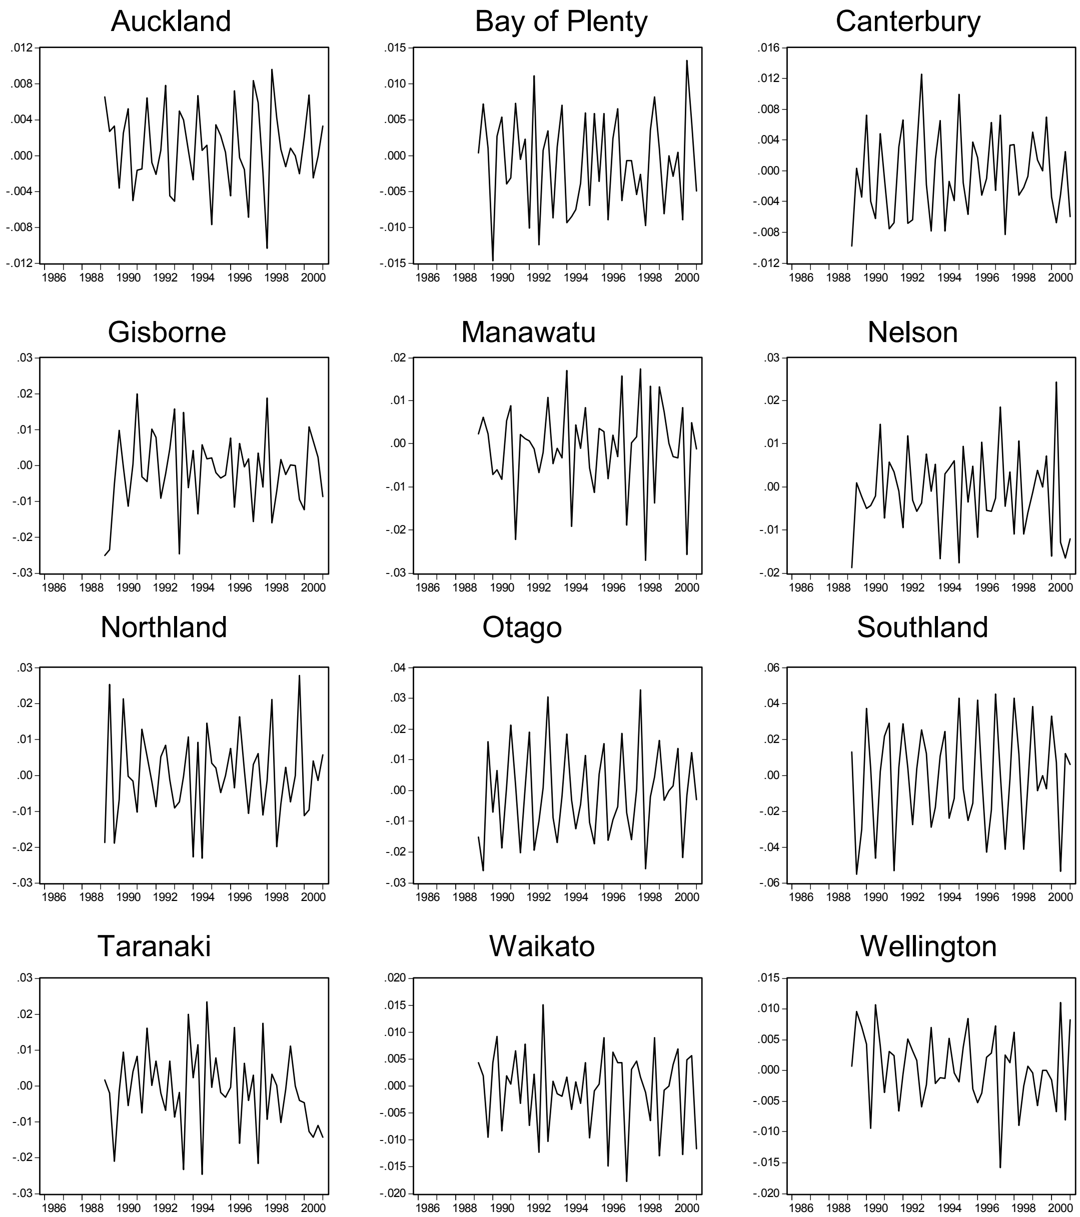

# Appendix B: Plots of the relative variables entering the VAR (using β–differences)

Appendix B, similar to Appendix A, plots the relative variables in their various forms entering the VAR. However, the β here used to construct the relative variables are obtained from regressions for each region. It is noted that for all the figures in this appendix, the scales on the vertical axis differ for each region.

***Figure B1 - Log of employment share (β-differences)***

Vertical axis displays: natural logarithm of region’s employment share
Horizontal axis displays: year

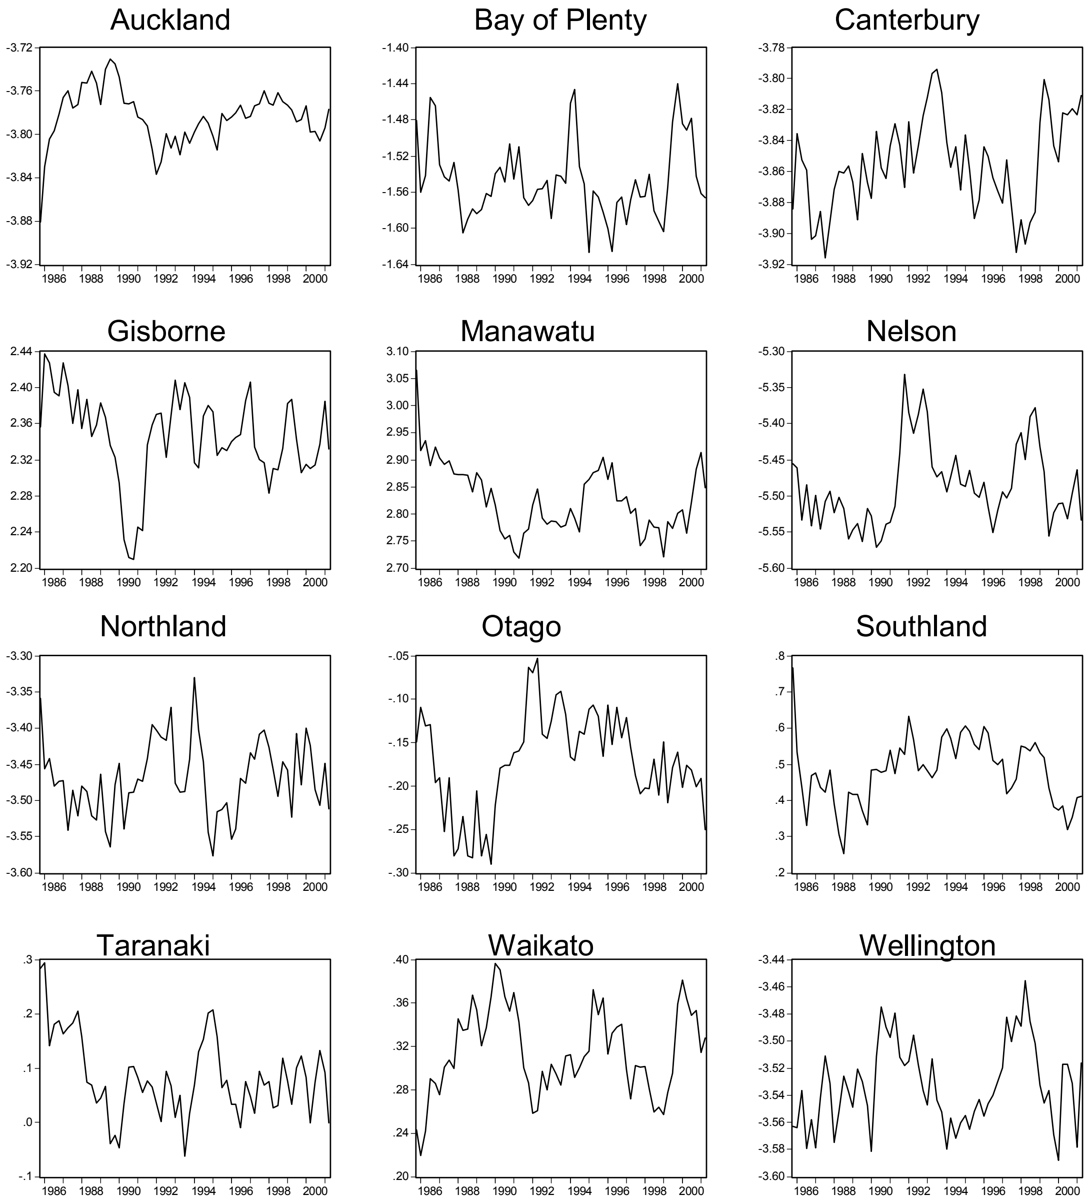

## *Figure B2 - Log of employment share in first difference (β-differences)*

Vertical axis displays: natural logarithm of region's employment share in first difference
Horizontal axis displays: year

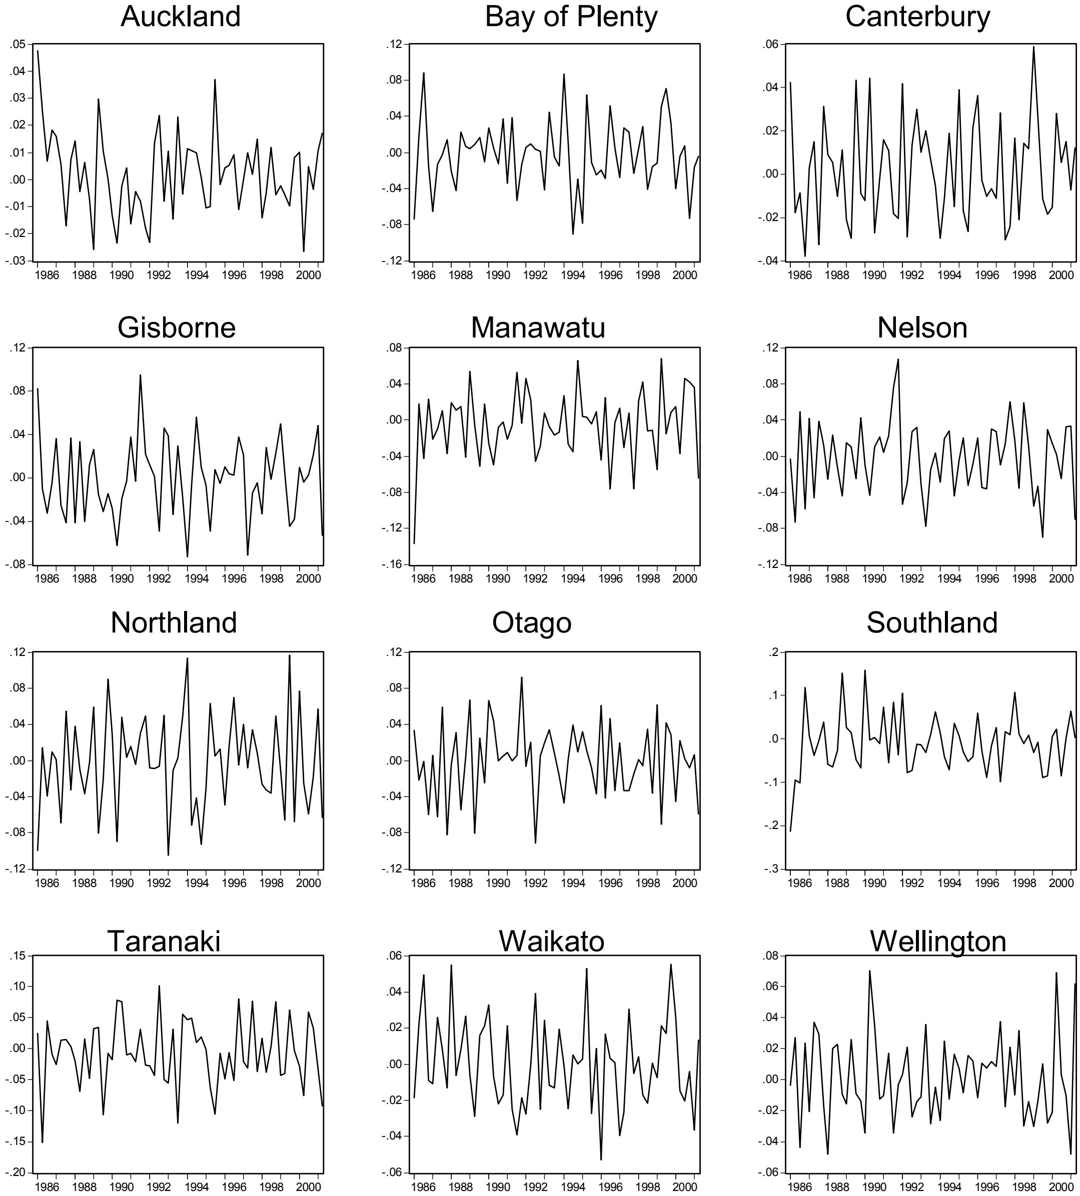

## *Figure B3 - Log of relative employment rate ($\beta$-differences)*

Vertical axis displays: natural logarithm of region's relative employment rate
Horizontal axis displays: year

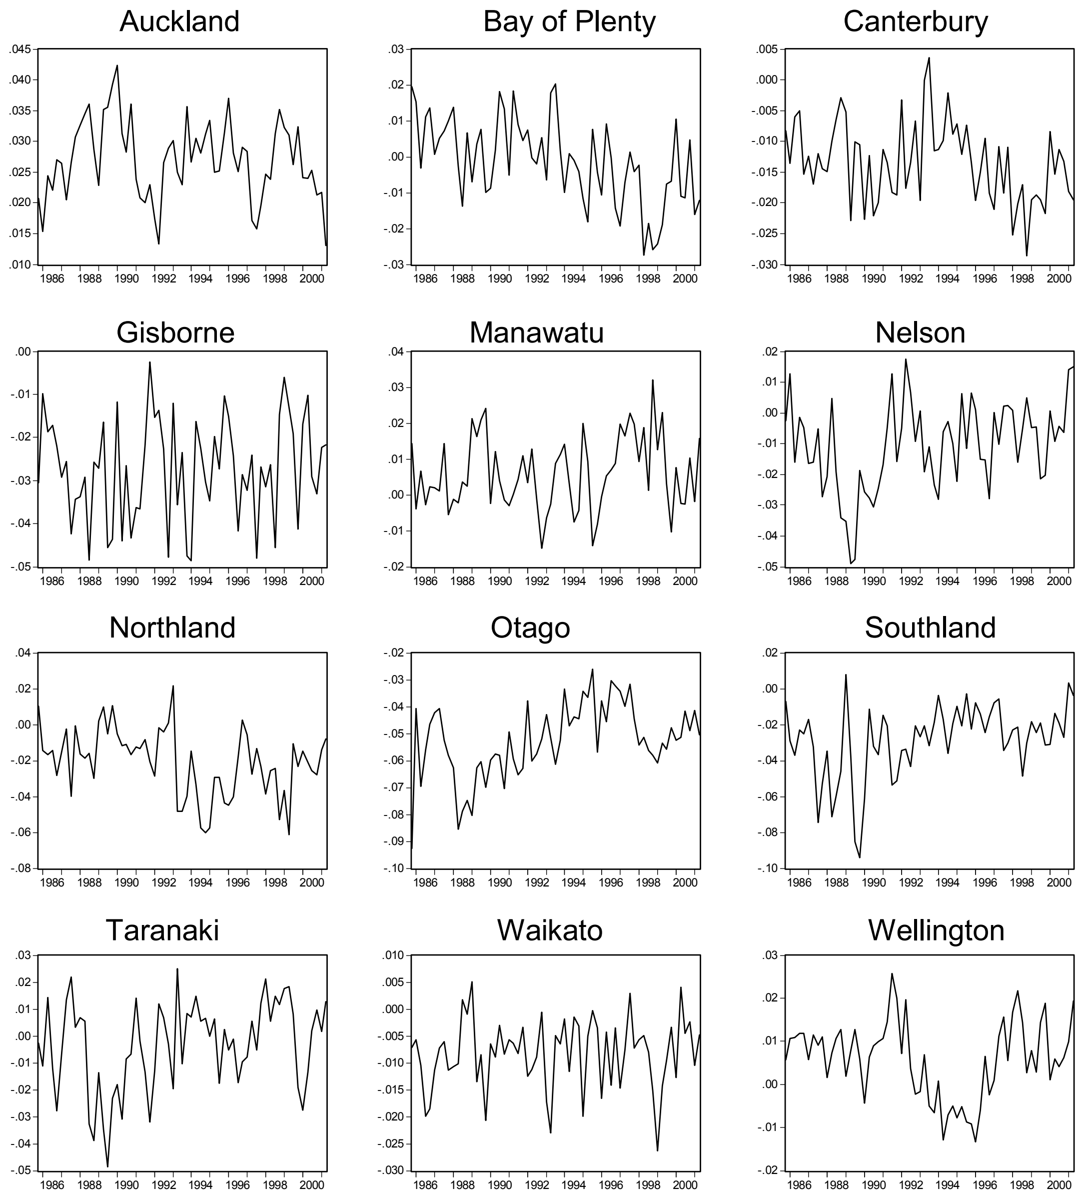

## *Figure B4 - Log of relative participation rate ($\beta$-differences)*

Vertical axis displays: natural logarithm of region's relative participation rate
Horizontal axis displays: year

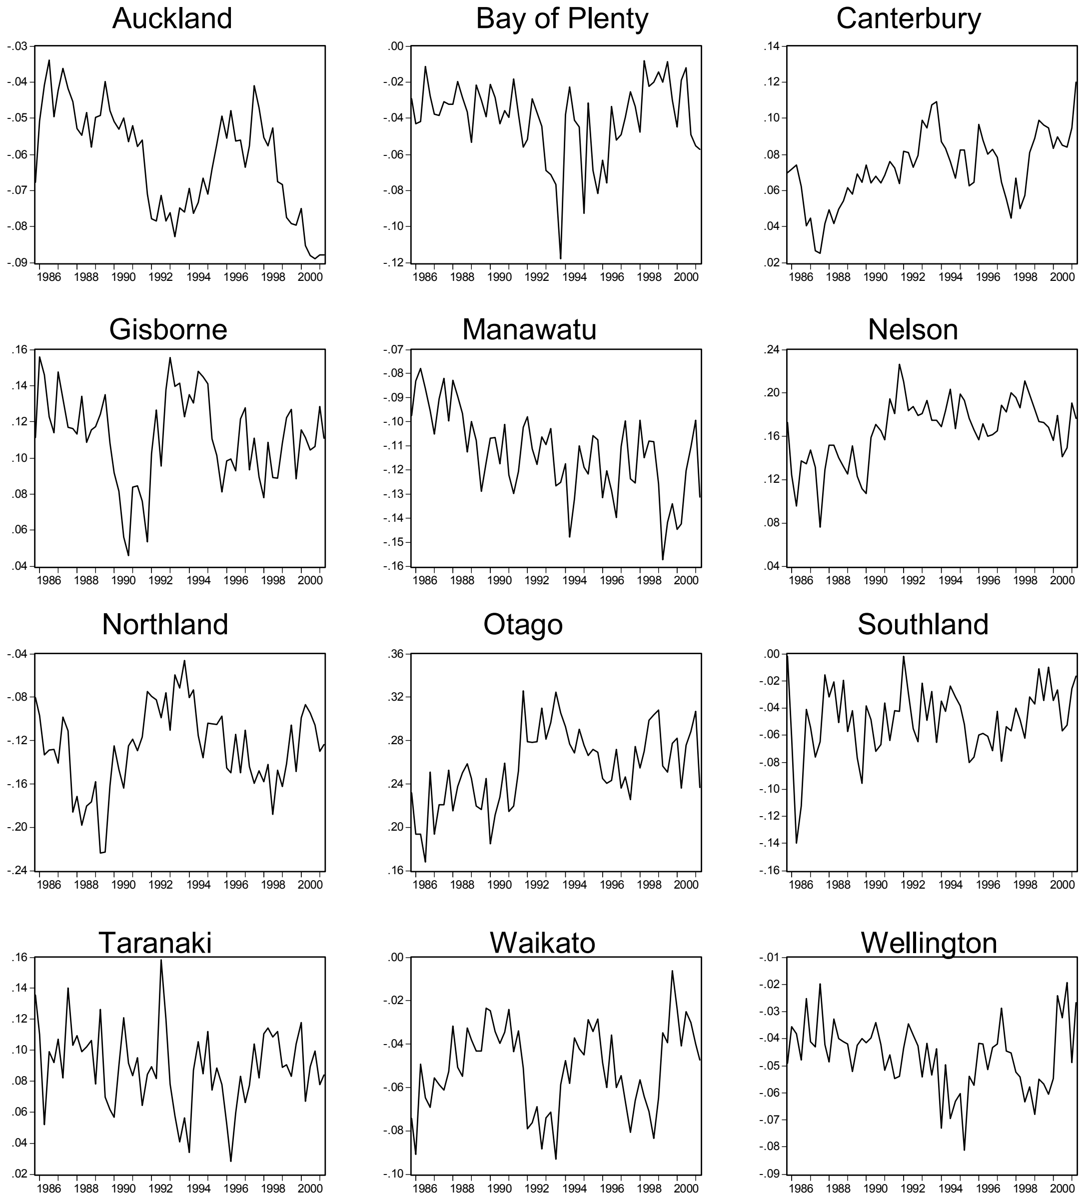

# Appendix C: Influence of an aggregate shock on regional variables

This appendix investigates the extent to which aggregate shocks get distributed across regions, as explained in Section 4. If the null hypothesis that β=1 is rejected, then one should, strictly speaking, construct the regional relative variables using beta-differencing. Conversely, if the null is not rejected, then one can use simple log-differencing.

***Table C1 - Results from the beta-regression (employment)***

*Estimated equation:* $\Delta \log(E_{jt}) = \alpha_{1j} + \beta_{E,j} \Delta \log(E_{NZt}) + \varepsilon_{1jt}$ [55]

| Region j | $\beta_{E,j}$ | t-stat ($H_0$: $\beta_{E,j}$ =1) | Adjusted $R^2$ |
|---|---|---|---|
| AKD | 1.0831 | 0.4779 | 0.3828 |
| BOP | 1.2319 | 0.5286 | 0.1014 |
| CAN | 0.7499 | -0.9222 | 0.0983 |
| GIS | 0.3511 | -1.5707 | -0.0046 |
| MWT | 1.0850 | 0.1939 | 0.0774 |
| NEL | 1.4744 | 1.0279 | 0.1311 |
| NLD | 0.8124 | -0.3114 | 0.0132 |
| OTG | 0.1141 | -1.8832* | -0.0157 |
| STH | 1.9503 | 1.2442 | 0.0830 |
| TNK | 1.6589 | 1.0605 | 0.0913 |
| WKT | 1.0162 | 0.0566 | 0.1593 |
| WLG | 0.9972 | -0.0090 | 0.1309 |

Note: * Reject $H_0$:β=1 at the 10% significance level.

According to the results above, the null hypothesis that β = 1 cannot be rejected at the 5% significance level for all regions; although it can be rejected at the 10% significance level for 1 region. Therefore, we will not lose much by imposing the restriction that β = 1 for constructing relative employment growth (ie, *dlnempn*).

[55] In terms of notation, $E_j$ stands for the employment level in region j. The variable with the NZ subscript denotes the corresponding national variable.

*Table C2 - Results from the beta-regression (employment rate)*

Estimated equation: $\log(ER_{jt}) = \alpha_{2j} + \beta_{ER,j} \log(ER_{NZt}) + \varepsilon_{2jt}$ [56]

| Region j | $\beta_{ER,j}$ | t-stat ($H_0$: $\beta_{ER,j}$=1) | Adjusted $R^2$ |
|---|---|---|---|
| AKD | 1.3152 | 8.1857*** | 0.9495 |
| BOP | 1.2451 | 3.4343*** | 0.8303 |
| CAN | 0.7547 | -6.3070*** | 0.8583 |
| GIS | 0.8083 | -2.6529*** | 0.6669 |
| MWT | 1.1177 | 1.9126* | 0.8414 |
| NEL | 0.7430 | -2.9370*** | 0.5343 |
| NLD | 1.1455 | 1.2989 | 0.6254 |
| OTG | 0.2296 | -9.3058*** | 0.0974 |
| STH | 0.5469 | -3.7097*** | 0.2351 |
| TNK | 1.1028 | 1.0274 | 0.6601 |
| WKT | 0.9760 | -0.6115 | 0.9085 |
| WLG | 0.9535 | -0.8713 | 0.837 |

Notes:
*** Reject $H_0$:$\beta$=1 at the 1% significance level.
* Reject $H_0$:$\beta$=1 at the 10% significance level.

The null hypothesis that $\beta$ = 1 can be rejected at the 1% level for 7 regions, and at the 10% level for 1 region. While this may suggest that one should use region-specific beta coefficients to constructing the regional relative employment rates, we have used simple log differences (assuming $\beta$ = 1) to obtain the main results in our paper. This assumption does not drive the results (see Section 7).

[56] $ER_j$ stands for the employment rate in region j. The variable with the NZ subscript denotes the corresponding national variable.

*Table C3 - Results from the beta-regression (participation rate)*

Estimated equation: $\log(PR_{jt}) = \alpha_{3j} + \beta_{PR,j}\log(PR_{NZt}) + \varepsilon_{3jt}$ [57]

| Region j | $\beta_{PR,j}$ | t-stat ($H_0$: $\beta_{PR,j}$=1) | Adjusted $R^2$ |
|---|---|---|---|
| AKD | 0.7615 | -2.0817** | 0.4106 |
| BOP | 0.9043 | -0.5850 | 0.3228 |
| CAN | 1.2319 | 1.5254 | 0.5104 |
| GIS | 1.4599 | 2.4156*** | 0.4825 |
| MWT | 0.7482 | -1.8608* | 0.3228 |
| NEL | 1.5061 | 2.2094** | 0.4052 |
| NLD | 0.7318 | -0.8963 | 0.0744 |
| OTG | 2.0191 | 3.6974*** | 0.4593 |
| STH | 0.7734 | -1.1574 | 0.1907 |
| TNK | 1.3045 | 1.5530 | 0.4110 |
| WKT | 0.8405 | -1.0635 | 0.3289 |
| WLG | 0.7344 | -2.7692*** | 0.4818 |

Notes:
*** Reject $H_0$:$\beta$=1 at the 1% significance level.
** Reject $H_0$:$\beta$=1 at the 5% significance level.
* Reject $H_0$:$\beta$=1 at the 10% significance level.

Given the results above, the null hypothesis that $\beta = 1$ is rejected at the 5% significance level for 5 out of 12 cases (6 cases at the 10% significance level). Therefore, there is mixed evidence in terms of whether we can assume unitary elasticity for all regions, or whether we should construct the regional relative participation rates using region-specific beta coefficients. As for the employment rate variable, we have assumed unitary elasticity in the main results in our paper. Nevertheless, the results in Section 7 suggest that the alternative way of constructing the relative variables do not make much difference to the results.

[57] $PR_j$ stands for the participation rate in region j. The variable with the NZ subscript denotes the corresponding national variable.

# Appendix D: Unit root test results

There are two parts in this appendix. The first part shows the results from the *panel* unit root tests. The second provides the unit root test results for each region *separately.* The variables tested for stationarity are those that enter the VAR (ie, in logs, and constructed as relative to national).

We applied the panel unit root test suggested by Im, Pesaran and Shin (1997) (hereafter, IPS). Essentially, the IPS panel unit root tests combine information on the stationarity or non-stationarity characteristics of the time series data for each region in the cross-section to give a conclusion for the entire panel. IPS proposed two panel unit root tests based on the mean of individual unit root test statistics: the LM-bar test and t-bar test. We have chosen to use the latter. This approach involves estimating an Augmented Dickey-Fuller (ADF) regression for each variable and region, and comparing a statistic based on the average t-statistic on the lagged level variable against a critical value, to test the null hypothesis of a unit root.

The following test statistic is employed:

$$z = \frac{\sqrt{N}[\bar{t}_{NT} - a_{NT}]}{\sqrt{b_{NT}}} \overset{a}{\sim} N(0,1)$$

where N is the number of time series in the panel, $\bar{t}_{NT}$ is the average t-statistic from a series of univariate ADF tests on each variable (where the number of lags in the regression is determined using the Schwarz criterion), and $a_{NT}$ and $b_{NT}$ are the expected small sample mean and variance of the distribution of t-statistics under the null hypothesis of non-stationarity (found through Monte-Carlo simulation, as obtained from Table 2 of IPS). This statistic is asymptotically distributed as a standard normal variable. The appropriate hypothesis is a one-sided test of z=0 against the alternative z<0. The exact sample critical values of the $\bar{t}_{NT}$ statistic is obtained from Table 4 of IPS. For a more detailed explanation of the test procedure, see Im, Pesaran and Shin (1997) and Banerjee (1999).

It is noted that in performing the IPS test procedure, an assumption has to be made about the deterministic components of each variable. In our case, the ADF regression included a constant, but no time trend for all the variables, with the exception of the employment share and relative wages, for which a time trend was included (only when both were tested in levels). The following table summarises the IPS panel unit root test results.

***Table D1 - Panel unit root test results***

| Variable | z |
|---|---|
| Log of Employment share (lnempn) | -6.0322*** |
| Change in Log of Employment share (dlnempn) | -27.9024*** |
| Log of Employment Rate (lnern) | -10.5782*** |
| Change in Log of Employment Rate (dlnern) | -29.2253*** |
| Log of Participation Rate (lnprn) | -7.2701*** |
| Change in Log of Participation Rate (dlnprn) | -30.2638*** |
| Log of Wages (lnwtotn) | -5.4664*** |
| Change in Log of Wages (dlnwtotn) | -24.6975*** |

Note:
*** = A unit root null hypothesis is rejected at the 1% significance level

In addition to performing the IPS panel unit root test, we have also performed the Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root tests for the variables for each region. These two tests are already well known in the literature and will not be covered here. Notes on how the optimal lag length was chosen are made under the table. For both tests, we allow for both a specification with and without a trend term.

***Table D2 - Individual unit root tests – ADF and PP***

| | Augmented Dickey Fuller (ADF) | | Phillips-Perron (PP) | |
|---|---|---|---|---|
| | Without trend | With trend | Without trend | With trend |
| em_akd | -2.8550* (0) | -2.8107 (0) | -2.9195** | -3.0032 |
| em_bop | -3.8130*** (0) | -3.9045** (0) | -3.7873*** | -3.8981** |
| em_can | -2.4233 (0) | -3.8295** (0) | -2.3029 | -3.8134** |
| em_gis | -1.7627 (0) | -3.3708* (0) | -1.7628 | -3.5905** |
| em_mwt | -3.1869** (0) | -4.6805*** (0) | -3.1902** | -4.8695*** |
| em_nel | -2.7112* (0) | -3.2633* (0) | -2.7845* | -3.3911* |
| em_nld | -4.8482*** (0) | -4.9216*** (0) | -4.9091*** | -4.9477*** |
| em_otg | -1.4669 (1) | -1.5564 (1) | -2.1170 | -2.3250 |
| em_sth | -3.6232*** (0) | -3.5749** (0) | -3.8561*** | -3.8520** |
| em_tnk | -3.1741** (0) | -4.1153*** (0) | -3.0295** | -4.0640** |
| em_wkt | -2.2234 (0) | -2.8006 (0) | -2.2570 | -2.7602 |
| em_wlg | -3.2854** (0) | -3.9362** (0) | -3.2140** | -3.9932** |
| Δem_akd | -7.8111*** (0) | -7.6928*** (0) | -7.9221*** | -7.7778*** |
| Δem_bop | -8.4680*** (0) | -8.3910*** (0) | -8.8797*** | -8.7745*** |
| Δem_can | -9.4046*** (0) | -9.4111*** (0) | -9.8753*** | -9.9802*** |
| Δem_gis | -8.0562*** (0) | -7.9856*** (0) | -8.1395*** | -8.0633*** |
| Δem_mwt | -10.9565*** (0) | -10.8837*** (0) | -11.3781*** | -11.3511*** |
| Δem_nel | -7.1319*** (1) | -7.0372*** (1) | -8.3742*** | -8.2821*** |
| Δem_nlD | -10.3198*** (0) | -10.2079*** (0) | -11.6116*** | -11.4560*** |
| Δem_otg | -11.9392*** (0) | -11.8545*** (0) | -12.1507*** | -12.0718*** |
| Δem_sth | -8.4125*** (0) | -8.3114*** (0) | -9.1616*** | -9.0662*** |
| Δem_tnk | -8.8288*** (0) | -8.7819*** (0) | -9.0820*** | -9.0374*** |
| Δem_wkt | -7.8902*** (0) | -7.9674*** (0) | -7.9022*** | -7.9901*** |
| Δem_wlg | -3.6963*** (3) | -3.6610** (3) | -10.1666*** | -10.0520*** |

| | Augmented Dickey Fuller (ADF) | | Phillips-Perron (PP) | |
|---|---|---|---|---|
| | Without trend | With trend | Without trend | With trend |
| er_akd | -2.6574* (0) | -2.7094 (0) | -2.6607* | -2.7296 |
| er_bop | -4.8822*** (0) | -6.0031*** (0) | -4.8380*** | -5.9585*** |
| er_can | -1.4937 (2) | -1.6258 (2) | -4.2952*** | -4.3453*** |
| er_gis | -6.5049*** (0) | -6.5434*** (0) | -6.5783*** | -6.5948*** |
| er_mwt | -5.4046*** (0) | -5.4547*** (0) | -5.5187*** | -5.5499*** |
| er_nel | -4.0675*** (0) | -4.4897*** (0) | -4.1511*** | -4.5678*** |
| er_nld | -4.8595*** (0) | -5.0324*** (0) | -4.8757*** | -5.0639*** |
| er_otg | -1.5404 (4) | -1.1572 (4) | -3.9365*** | -3.7472*** |
| er_sth | -4.1789*** (0) | -4.8794*** (0) | -4.0948*** | -4.7848*** |
| er_tnk | -4.4402*** (0) | -4.7154*** (0) | -4.5000*** | -4.7346*** |
| er_wkt | -6.8382*** (0) | -6.8338*** (0) | -6.7923*** | -6.7843*** |
| er_wlg | -3.5043** (0) | -3.4584* (0) | -3.5631*** | -3.5208** |
| Δer_akd | -9.3461*** (0) | -9.2700*** (0) | -9.4081*** | -9.3234*** |
| Δer_bop | -9.7929*** (1) | -9.7241*** (1) | -13.3336*** | -13.3661*** |
| Δer_can | -10.1419*** (1) | -10.1626*** (1) | -15.9182*** | -16.1567*** |
| Δer_gis | -7.6694*** (2) | -7.5990*** (2) | -17.0067*** | -16.8469*** |
| Δer_mwt | -12.2498*** (0) | -12.1503*** (0) | -14.4803*** | -14.3396*** |
| Δer_nel | -11.3267*** (0) | -11.2978*** (0) | -12.5126*** | -12.5118*** |
| Δer_nld | -10.9788*** (0) | -10.9220*** (0) | -13.7231*** | -13.7228*** |
| Δer_otg | -4.0534*** (3) | -4.2033*** (3) | -15.8251*** | -16.4787*** |
| Δer_sth | -5.3566*** (5) | -5.3063*** (5) | -10.4166*** | -10.2721*** |
| Δer_tnk | -10.2038*** (0) | -10.1242*** (0) | -11.3409*** | -11.2509*** |
| Δer_wkt | -8.0640*** (2) | -7.9871*** (2) | -15.3227*** | -15.1611*** |
| Δer_wlg | -8.6527*** (1) | -8.6120*** (1) | -10.8187*** | -10.7737*** |
| pr_akd | -1.7549 (0) | -3.7553** (0) | -1.7246 | -3.9103** |
| pr_bop | -4.8846*** (0) | -4.8745*** (0) | -4.9469*** | -4.9422*** |
| pr_can | -1.6808 (0) | -3.1004 (0) | -1.7701 | -3.1928* |
| pr_gis | -3.4160** (0) | -3.3830* (0) | -3.3898** | -3.3530* |
| pr_mwt | -3.4162** (0) | -5.4839*** (0) | -3.3254** | -5.4230*** |
| pr_nel | -3.1270** (0) | -4.7139*** (0) | -3.0464** | -4.8661*** |
| pr_nld | -3.3114** (0) | -3.3363* (0) | -3.2523** | -3.2700* |
| pr_otg | -1.4766 (4) | -2.5446 (4) | -4.0129*** | -6.2065*** |
| pr_sth | -5.3653*** (0) | -5.5770*** (0) | -5.5871*** | -5.7970*** |
| pr_tnk | -5.5445*** (0) | -5.4778*** (0) | -5.5818*** | -5.5149*** |
| pr_wkt | -3.6033*** (0) | -3.5520** (0) | -3.5977*** | -3.5433** |
| pr_wlg | -2.7463* (1) | -2.9823 (1) | -4.9457*** | -5.4564*** |
| Δpr_akd | -9.8971*** (0) | -9.9099*** (0) | -10.4453*** | -10.5133*** |
| Δpr_bop | -4.6970*** (5) | -4.6295*** (5) | -13.8463*** | -13.7009*** |
| Δpr_can | -7.6271*** (0) | -7.6899*** (0) | -7.6144*** | -7.6865*** |
| Δpr_gis | -9.9333*** (0) | -9.9203*** (0) | -10.5761*** | -10.6653*** |
| Δpr_mwt | -9.1671*** (1) | -9.1101*** (1) | -10.1175*** | -10.0249*** |
| Δpr_nel | -10.1944*** (0) | -10.0815*** (0) | -12.0799*** | -11.9126*** |
| Δpr_nld | -10.5377*** (0) | -10.4592*** (0) | -11.5369*** | -11.4488*** |
| Δpr_otg | -4.3242*** (3) | -4.2824*** (3) | -14.3454*** | -14.1612*** |

| | Augmented Dickey Fuller (ADF) | | Phillips-Perron (PP) | |
|---|---|---|---|---|
| | Without trend | With trend | Without trend | With trend |
| Δpr_sth | -9.6056*** (2) | -9.6209*** (2) | -12.5965*** | -12.3799*** |
| Δpr_tnk | -10.9206*** (0) | -10.8433*** (0) | -12.3216*** | -12.2197*** |
| Δpr_wkt | -10.4076*** (0) | -10.3804*** (0) | -10.6759*** | -10.6438*** |
| Δpr_wlg | -13.3567*** (0) | -13.2338*** (0) | -15.1835*** | -15.1077*** |
| w_akd | 0.61559 (4) | -2.6174 (4) | -1.4972 | -5.9694*** |
| w_bop | -1.1396 (0) | -4.0898** (0) | -0.8184 | -4.1190** |
| w_can | -4.1880*** (0) | -4.1252** (0) | -4.1425*** | -4.0857** |
| w_gis | -0.2229 (4) | -1.9899 (4) | -3.2365** | -4.0861** |
| w_mwt | -2.6652 (0) | -5.6553*** (0) | -2.3245 | -5.7740*** |
| w_nel | -3.2776** (0) | -5.3408*** (0) | -3.4007** | -5.3916*** |
| w_nld | -4.6434*** (0) | -4.9410*** (0) | -4.8230*** | -5.0915*** |
| w_otg | -0.9637 (4) | -1.5114 (4) | -3.7276*** | -5.1660*** |
| w_sth | -2.5224 (4) | -2.8229 (4) | -5.3045*** | -5.6863*** |
| w_tnk | -0.0813 (1) | 0.0368 (1) | -1.4204 | -1.3333 |
| w_wkt | -1.9145 (0) | -5.9262*** (0) | -1.4085 | -6.0007*** |
| w_wlg | -2.0582 (3) | -2.1766 (3) | -3.4545** | -3.1918* |
| Δw_akd | -10.4460*** (2) | -10.4270*** (2) | -11.0407*** | -10.8893*** |
| Δw_bop | -9.0996*** (0) | -8.9987*** (0) | -10.0279*** | -9.9026*** |
| Δw_can | -8.1030*** (1) | -8.0664*** (1) | -9.3708*** | -9.2582*** |
| Δw_gis | -4.4689*** (3) | -4.9406*** (3) | -9.6328*** | -9.4604*** |
| Δw_mwt | -7.5753*** (2) | -7.5763*** (2) | -15.2001*** | -15.0110*** |
| Δw_nel | -6.0554*** (2) | -5.9566*** (2) | -11.6320*** | -11.6729*** |
| Δw_nld | -8.0579*** (1) | -7.9492*** (1) | -13.1597*** | -13.0507*** |
| Δw_otg | -4.5972*** (3) | -4.4395*** (3) | -10.4710*** | -10.5033*** |
| Δw_sth | -9.6028*** (2) | -9.6571*** (2) | -8.8397*** | -8.9035*** |
| Δw_tnk | -9.5711*** (0) | -9.8344*** (0) | -9.2917*** | -9.6513*** |
| Δw_wkt | -7.9245*** (2) | -8.0718*** (2) | -16.1902*** | -16.4146*** |
| Δw_wlg | -8.5727*** (0) | -8.9666*** (0) | -9.0232*** | -10.4191*** |

Notes: Numbers in brackets for the Augmented Dickey-Fuller (ADF) test column indicate the number of lags used in the ADF test. The lag lengths were determined using the Schwarz Bayesian Criterion. For the Phillips Perron (PP) test, the lag truncation used for Bartlett Kernel is four in all cases. The critical values for rejection of a unit root null hypothesis were obtained from MacKinnon (1991).

* = A unit root null hypothesis is rejected at the 10% significance level
** = A unit root null hypothesis is rejected at the 5% significance level
*** = A unit root null hypothesis is rejected at the 1% significance level
Shaded cells represent cases where the unit root null is rejected for at least at the 10% significance level.

# Appendix E: VAR coefficients

The tables below present the detailed coefficients from the pooled VAR models in Section 6 of this paper. It is noted that the standard errors (hence, the t-statistics) shown in the tables below are robust to general heteroskedasticity (see White, 1980).[58] The corresponding tables for individual regions are not reported here, but are available upon request from the authors.

**3-variable VAR (employment share entering in first difference) (adjusted sample: 1987Q1-2001Q2) – White heteroskedasticity-consistent standard errors & covariance**

***Change in employment equation***

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| ΔEM(-1) | -0.172950 | 0.047177 | -3.666024 | 0.0003 |
| ΔEM(-2) | -0.136997 | 0.048929 | -2.799907 | 0.0053 |
| ΔEM(-3) | -0.093997 | 0.048030 | -1.957043 | 0.0508 |
| ΔEM(-4) | 0.013928 | 0.039238 | 0.354960 | 0.7227 |
| ER(-1) | -0.158485 | 0.157136 | -1.008587 | 0.3135 |
| ER(-2) | -0.030644 | 0.164422 | -0.186376 | 0.8522 |
| ER(-3) | -0.187445 | 0.158242 | -1.184550 | 0.2366 |
| ER(-4) | 0.208403 | 0.168213 | 1.238919 | 0.2158 |
| PR(-1) | -0.155922 | 0.088422 | -1.763394 | 0.0783 |
| PR(-2) | -0.107352 | 0.097795 | -1.097720 | 0.2727 |
| PR(-3) | 0.100517 | 0.101007 | 0.995152 | 0.3200 |
| PR(-4) | -0.143497 | 0.088649 | -1.618707 | 0.1060 |
| Fixed Effects | | | | |
| _AKD--C | 0.005094 | | | |
| _BOP--C | -0.009777 | | | |
| _CAN--C | 0.004272 | | | |
| _GIS--C | -0.014561 | | | |
| _MWT--C | -0.016733 | | | |
| _NEL--C | 0.007447 | | | |
| _NLD--C | -0.021167 | | | |
| _OTG--C | -0.015660 | | | |
| _STH--C | 0.003753 | | | |
| _TNK--C | -0.008218 | | | |
| _WKT--C | -0.002063 | | | |
| _WLG--C | 0.012865 | | | |
| R-squared | 0.103452 | | Mean dependent var | -0.001057 |
| Adjusted R-squared | 0.072767 | | S.D. dependent var | 0.037855 |
| S.E. of regression | 0.036452 | | Sum squared resid | 0.892900 |
| Log likelihood | 1329.622 | | F-statistic | 7.049268 |
| Durbin-Watson stat | 1.990184 | | Prob(F-statistic) | 0.000000 |

[58] This form of heteroskedasticity is more general than the cross-section heteroskedasticity, since variances within a cross-section are allowed to differ across *time* (*E-Views* Help function).

### *Employment rate equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| ΔEM | 0.081383 | 0.012867 | 6.324873 | 0.0000 |
| ΔEM(-1) | 0.057209 | 0.018753 | 3.050750 | 0.0024 |
| ΔEM(-2) | 0.003302 | 0.014959 | 0.220710 | 0.8254 |
| ΔEM(-3) | 0.006942 | 0.016596 | 0.418271 | 0.6759 |
| ΔEM(-4) | 0.007648 | 0.012355 | 0.619046 | 0.5361 |
| ER(-1) | 0.445957 | 0.051754 | 8.616813 | 0.0000 |
| ER(-2) | 0.053266 | 0.054570 | 0.976105 | 0.3294 |
| ER(-3) | 0.074150 | 0.054929 | 1.349931 | 0.1775 |
| ER(-4) | 0.121924 | 0.041820 | 2.915447 | 0.0037 |
| PR(-1) | 0.078913 | 0.029404 | 2.683782 | 0.0075 |
| PR(-2) | -0.039515 | 0.030963 | -1.276203 | 0.2023 |
| PR(-3) | -0.008011 | 0.029809 | -0.268738 | 0.7882 |
| PR(-4) | -0.009901 | 0.027840 | -0.355645 | 0.7222 |
| Fixed Effects | | | | |
| _AKD--C | 0.000512 | | | |
| _BOP--C | -0.006330 | | | |
| _CAN--C | 0.001060 | | | |
| _GIS--C | -0.003134 | | | |
| _MWT--C | 0.000649 | | | |
| _NEL--C | 0.002349 | | | |
| _NLD--C | -0.008730 | | | |
| _OTG--C | 0.003441 | | | |
| _STH--C | 0.001963 | | | |
| _TNK--C | -0.002133 | | | |
| _WKT--C | -0.001574 | | | |
| _WLG--C | 0.001988 | | | |
| R-squared | 0.655980 | | Mean dependent var | -0.004084 |
| Adjusted R-squared | 0.643675 | | S.D. dependent var | 0.019058 |
| S.E. of regression | 0.011376 | | Sum squared resid | 0.086840 |
| Log likelihood | 2140.604 | | F-statistic | 106.6224 |
| Durbin-Watson stat | 2.015956 | | Prob(F-statistic) | 0.000000 |

### *Participation rate equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| ΔEM | 0.208449 | 0.023232 | 8.972368 | 0.0000 |
| ΔEM(-1) | 0.068830 | 0.025445 | 2.705062 | 0.0070 |
| ΔEM(-2) | 0.057983 | 0.024054 | 2.410513 | 0.0162 |
| ΔEM(-3) | 0.043520 | 0.023140 | 1.880723 | 0.0604 |
| ΔEM(-4) | -0.007062 | 0.019390 | -0.364182 | 0.7158 |
| ER(-1) | 0.076698 | 0.074058 | 1.035645 | 0.3007 |
| ER(-2) | -0.008629 | 0.078416 | -0.110040 | 0.9124 |
| ER(-3) | 0.075641 | 0.083778 | 0.902884 | 0.3669 |
| ER(-4) | 0.089937 | 0.063767 | 1.410409 | 0.1589 |
| PR(-1) | 0.464892 | 0.045771 | 10.15684 | 0.0000 |
| PR(-2) | 0.119778 | 0.056136 | 2.133702 | 0.0332 |
| PR(-3) | 0.039146 | 0.050461 | 0.775760 | 0.4382 |
| PR(-4) | 0.185604 | 0.046413 | 3.998993 | 0.0001 |
| Fixed Effects | | | | |
| _AKD--C | -0.000893 | | | |
| _BOP--C | 0.002283 | | | |
| _CAN--C | 0.000391 | | | |
| _GIS--C | -0.001214 | | | |
| _MWT--C | -0.006793 | | | |
| _NEL--C | 0.002394 | | | |
| _NLD--C | -0.001240 | | | |
| _OTG--C | -0.007529 | | | |
| _STH--C | 0.005117 | | | |
| _TNK--C | 0.003617 | | | |
| _WKT--C | 0.001884 | | | |
| _WLG--C | 0.003656 | | | |
| R-squared | 0.774768 | | Mean dependent var | -0.007755 |
| Adjusted R-squared | 0.766712 | | S.D. dependent var | 0.035691 |
| S.E. of regression | 0.017239 | | Sum squared resid | 0.199403 |
| Log likelihood | 1851.325 | | F-statistic | 192.3453 |
| Durbin-Watson stat | 1.957759 | | Prob(F-statistic) | 0.000000 |

**3 variable VAR (all variables entering in levels) (adjusted sample: 1986Q4-2001Q2) – White heteroskedasticity-consistent standard errors & covariance**

### *Employment equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| EM(-1) | 0.757036 | 0.047245 | 16.02378 | 0.0000 |
| EM(-2) | 0.026173 | 0.054120 | 0.483621 | 0.6288 |
| EM(-3) | 0.004154 | 0.054711 | 0.075933 | 0.9395 |
| EM(-4) | 0.073489 | 0.047124 | 1.559489 | 0.1193 |
| ER(-1) | -0.092745 | 0.154661 | -0.599669 | 0.5489 |
| ER(-2) | 0.018631 | 0.156935 | 0.118720 | 0.9055 |
| ER(-3) | -0.166005 | 0.155642 | -1.066579 | 0.2865 |
| ER(-4) | 0.275357 | 0.162232 | 1.697302 | 0.0901 |
| PR(-1) | -0.094163 | 0.088237 | -1.067151 | 0.2863 |
| PR(-2) | -0.102940 | 0.094725 | -1.086732 | 0.2775 |
| PR(-3) | 0.135338 | 0.094537 | 1.431584 | 0.1527 |
| PR(-4) | -0.120274 | 0.085500 | -1.406714 | 0.1600 |
| Fixed Effects | | | | |
| _AKD--C | -0.165304 | | | |
| _BOP--C | -0.408875 | | | |
| _CAN--C | -0.275656 | | | |
| _GIS--C | -0.424938 | | | |
| _MWT--C | -0.410416 | | | |
| _NEL--C | -0.433396 | | | |
| _NLD--C | -0.498240 | | | |
| _OTG--C | -0.420674 | | | |
| _STH--C | -0.494287 | | | |
| _TNK--C | -0.504239 | | | |
| _WKT--C | -0.329143 | | | |
| _WLG--C | -0.279130 | | | |
| R-squared | 0.997461 | | Mean dependent var | -2.763553 |
| Adjusted R-squared | 0.997376 | | S.D. dependent var | 0.699640 |
| S.E. of regression | 0.035841 | | Sum squared resid | 0.878656 |
| Log likelihood | 1364.291 | | F-statistic | 24429.21 |
| Durbin-Watson stat | 1.938682 | | Prob(F-statistic) | 0.000000 |

### *Employment rate equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| EM | 0.084022 | 0.012946 | 6.490365 | 0.0000 |
| EM(-1) | -0.020694 | 0.020290 | -1.019923 | 0.3081 |
| EM(-2) | -0.048485 | 0.022417 | -2.162850 | 0.0309 |
| EM(-3) | 0.002251 | 0.018757 | 0.120012 | 0.9045 |
| EM(-4) | -0.003444 | 0.015997 | -0.215282 | 0.8296 |
| ER(-1) | 0.441063 | 0.050102 | 8.803271 | 0.0000 |
| ER(-2) | 0.037867 | 0.053956 | 0.701815 | 0.4830 |
| ER(-3) | 0.086984 | 0.055117 | 1.578166 | 0.1150 |
| ER(-4) | 0.107443 | 0.041520 | 2.587777 | 0.0099 |
| PR(-1) | 0.072551 | 0.028735 | 2.524790 | 0.0118 |
| PR(-2) | -0.035847 | 0.030144 | -1.189201 | 0.2348 |
| PR(-3) | -0.008325 | 0.029695 | -0.280350 | 0.7793 |
| PR(-4) | -0.017314 | 0.027894 | -0.620717 | 0.5350 |
| Fixed Effects | | | | |
| _AKD—C | 0.017347 | | | |
| _BOP—C | 0.032871 | | | |
| _CAN—C | 0.028389 | | | |
| _GIS—C | 0.037150 | | | |
| _MWT--C | 0.039329 | | | |
| _NEL--C | 0.045450 | | | |
| _NLD--C | 0.037895 | | | |
| _OTG--C | 0.043355 | | | |
| _STH--C | 0.050709 | | | |
| _TNK--C | 0.046057 | | | |
| _WKT--C | 0.030379 | | | |
| _WLG--C | 0.030680 | | | |
| R-squared | 0.656845 | | Mean dependent var | -0.004144 |
| Adjusted R-squared | 0.644787 | | S.D. dependent var | 0.019004 |
| S.E. of regression | 0.011327 | | Sum squared resid | 0.087623 |
| Log likelihood | 2180.384 | | F-statistic | 108.9463 |
| Durbin-Watson stat | 2.013674 | | Prob(F-statistic) | 0.000000 |

### *Participation rate equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| EM | 0.211697 | 0.023536 | 8.994575 | 0.0000 |
| EM(-1) | -0.136711 | 0.032378 | -4.222320 | 0.0000 |
| EM(-2) | -0.002163 | 0.032859 | -0.065837 | 0.9475 |
| EM(-3) | -0.010282 | 0.029043 | -0.354043 | 0.7234 |
| EM(-4) | -0.034185 | 0.022874 | -1.494477 | 0.1355 |
| ER(-1) | 0.072286 | 0.073023 | 0.989898 | 0.3226 |
| ER(-2) | -0.041263 | 0.078399 | -0.526318 | 0.5988 |
| ER(-3) | 0.088701 | 0.082870 | 1.070364 | 0.2848 |
| ER(-4) | 0.041701 | 0.068027 | 0.612999 | 0.5401 |
| PR(-1) | 0.445182 | 0.048480 | 9.182753 | 0.0000 |
| PR(-2) | 0.103248 | 0.054385 | 1.898482 | 0.0581 |
| PR(-3) | 0.034676 | 0.048898 | 0.709149 | 0.4785 |
| PR(-4) | 0.190371 | 0.045473 | 4.186431 | 0.0000 |
| Fixed Effects | | | | |
| _AKD--C | 0.033976 | | | |
| _BOP--C | 0.082760 | | | |
| _CAN--C | 0.057183 | | | |
| _GIS--C | 0.081610 | | | |
| _MWT--C | 0.072923 | | | |
| _NEL--C | 0.092499 | | | |
| _NLD--C | 0.094248 | | | |
| _OTG--C | 0.075699 | | | |
| _STH--C | 0.107643 | | | |
| _TNK--C | 0.104271 | | | |
| _WKT--C | 0.068158 | | | |
| _WLG--C | 0.064114 | | | |
| R-squared | 0.770023 | | Mean dependent var | -0.007753 |
| Adjusted R-squared | 0.761942 | | S.D. dependent var | 0.035566 |
| S.E. of regression | 0.017353 | | Sum squared resid | 0.205668 |
| Log likelihood | 1878.345 | | F-statistic | 190.5718 |
| Durbin-Watson stat | 1.924299 | | Prob(F-statistic) | 0.000000 |

### 4-variable VAR (employment share and wages entering in first difference) (adjusted sample: 1990Q2-2001Q1) – White HS-consistent standard errors

#### *Change in employment equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| ΔEM(-1) | -0.106688 | 0.051649 | -2.065639 | 0.0394 |
| ΔEM(-2) | -0.155092 | 0.051586 | -3.006501 | 0.0028 |
| ΔEM(-3) | -0.124742 | 0.056744 | -2.198324 | 0.0284 |
| ΔEM(-4) | -0.002147 | 0.049924 | -0.043012 | 0.9657 |
| ER(-1) | -0.362515 | 0.177357 | -2.043985 | 0.0415 |
| ER(-2) | 0.237710 | 0.184892 | 1.285666 | 0.1992 |
| ER(-3) | -0.380691 | 0.187478 | -2.030595 | 0.0428 |
| ER(-4) | 0.011972 | 0.191138 | 0.062636 | 0.9501 |
| PR(-1) | -0.220165 | 0.102676 | -2.144277 | 0.0325 |
| PR(-2) | 0.024187 | 0.110137 | 0.219605 | 0.8263 |
| PR(-3) | 0.046544 | 0.119721 | 0.388776 | 0.6976 |
| PR(-4) | -0.280980 | 0.107531 | -2.613011 | 0.0092 |
| ΔW(-1) | 0.164830 | 0.199965 | 0.824291 | 0.4102 |
| ΔW(-2) | 0.158875 | 0.251850 | 0.630834 | 0.5284 |
| ΔW(-3) | 0.342886 | 0.245061 | 1.399186 | 0.1624 |
| ΔW(-4) | 0.562992 | 0.230520 | 2.442268 | 0.0149 |
| Fixed Effects | | | | |
| _AKD--C | 0.003335 | | | |
| _BOP--C | -0.016865 | | | |
| _CAN--C | 0.010017 | | | |
| _GIS--C | -0.019809 | | | |
| _MWT--C | -0.022365 | | | |
| _NEL--C | 0.020989 | | | |
| _NLD--C | -0.035074 | | | |
| _OTG--C | -0.009478 | | | |
| _STH--C | 0.009793 | | | |
| _TNK--C | -0.006171 | | | |
| _WKT--C | -0.007855 | | | |
| _WLG--C | 0.017840 | | | |
| R-squared | 0.135350 | | Mean dependent var | -0.000405 |
| Adjusted R-squared | 0.088659 | | S.D. dependent var | 0.037197 |
| S.E. of regression | 0.035509 | | Sum squared resid | 0.630460 |
| Log likelihood | 1027.626 | | F-statistic | 5.217931 |
| Durbin-Watson stat | 1.957261 | | Prob(F-statistic) | 0.000000 |

### *Employment rate equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| ΔEM | 0.074900 | 0.013258 | 5.649489 | 0.0000 |
| ΔEM(-1) | 0.037565 | 0.017432 | 2.155004 | 0.0316 |
| ΔEM(-2) | 0.004048 | 0.016856 | 0.240162 | 0.8103 |
| ΔEM(-3) | 0.027356 | 0.014933 | 1.831888 | 0.0676 |
| ΔEM(-4) | 0.010535 | 0.011903 | 0.885083 | 0.3765 |
| ER(-1) | 0.366479 | 0.055523 | 6.600506 | 0.0000 |
| ER(-2) | 0.048998 | 0.054836 | 0.893538 | 0.3720 |
| ER(-3) | -0.001238 | 0.050574 | -0.024486 | 0.9805 |
| ER(-4) | 0.120640 | 0.041918 | 2.878004 | 0.0042 |
| PR(-1) | 0.092284 | 0.031541 | 2.925870 | 0.0036 |
| PR(-2) | -0.063462 | 0.033044 | -1.920532 | 0.0554 |
| PR(-3) | 0.009969 | 0.032305 | 0.308580 | 0.7578 |
| PR(-4) | -0.032497 | 0.030541 | -1.064026 | 0.2878 |
| ΔW(-1) | -0.039328 | 0.056336 | -0.698104 | 0.4854 |
| ΔW(-2) | -0.129349 | 0.061735 | -2.095240 | 0.0367 |
| ΔW(-3) | -0.021400 | 0.060797 | -0.351988 | 0.7250 |
| ΔW(-4) | 0.014478 | 0.053175 | 0.272263 | 0.7855 |
| Fixed Effects | | | | |
| _AKD--C | 0.000325 | | | |
| _BOP--C | -0.010831 | | | |
| _CAN--C | 0.002177 | | | |
| _GIS--C | -0.004926 | | | |
| _MWT--C | -0.001417 | | | |
| _NEL--C | 0.005894 | | | |
| _NLD--C | -0.017774 | | | |
| _OTG--C | 0.006702 | | | |
| _STH--C | 0.007615 | | | |
| _TNK--C | -0.003767 | | | |
| _WKT--C | -0.002460 | | | |
| _WLG--C | 0.003747 | | | |
| R-squared | 0.729976 | | Mean dependent var | -0.002887 |
| Adjusted R-squared | 0.714825 | | S.D. dependent var | 0.019221 |
| S.E. of regression | 0.010264 | | Sum squared resid | 0.052574 |
| Log likelihood | 1683.463 | | F-STATISTIC | 84.31163 |
| Durbin-Watson stat | 1.979952 | | Prob(F-statistic) | 0.000000 |

### *Participation rate equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| ΔEM | 0.209954 | 0.024961 | 8.411323 | 0.0000 |
| ΔEM(-1) | 0.048269 | 0.026957 | 1.790568 | 0.0740 |
| ΔEM(-2) | 0.073604 | 0.025993 | 2.831646 | 0.0048 |
| ΔEM(-3) | 0.011048 | 0.027383 | 0.403473 | 0.6868 |
| ΔEM(-4) | -0.009480 | 0.022415 | -0.422902 | 0.6725 |
| ER(-1) | 0.074853 | 0.088612 | 0.844723 | 0.3987 |
| ER(-2) | -0.056922 | 0.085619 | -0.664828 | 0.5065 |
| ER(-3) | 0.113670 | 0.092881 | 1.223830 | 0.2216 |
| ER(-4) | 0.016051 | 0.070751 | 0.226868 | 0.8206 |
| PR(-1) | 0.478619 | 0.053451 | 8.954273 | 0.0000 |
| PR(-2) | 0.064444 | 0.058555 | 1.100558 | 0.2716 |
| PR(-3) | 0.086887 | 0.061769 | 1.406661 | 0.1602 |
| PR(-4) | 0.111115 | 0.056667 | 1.960861 | 0.0505 |
| ΔW(-1) | -0.061202 | 0.091294 | -0.670387 | 0.5029 |
| ΔW(-2) | -0.013446 | 0.100711 | -0.133513 | 0.8938 |
| ΔW(-3) | 0.057816 | 0.108352 | 0.533594 | 0.5939 |
| ΔW(-4) | 0.094659 | 0.097159 | 0.974266 | 0.3304 |
| Fixed Effects | | | | |
| _AKD--C | -0.000794 | | | |
| _BOP--C | -0.000808 | | | |
| _CAN--C | 0.001921 | | | |
| _GIS--C | -0.003993 | | | |
| _MWT--C | -0.010603 | | | |
| _NEL--C | 0.007581 | | | |
| _NLD--C | -0.003642 | | | |
| _OTG--C | -0.006225 | | | |
| _STH--C | 0.005645 | | | |
| _TNK--C | 0.001222 | | | |
| _WKT--C | 0.000512 | | | |
| _WLG--C | 0.005740 | | | |
| R-squared | 0.767202 | | Mean dependent var | -0.005987 |
| Adjusted R-squared | 0.754139 | | S.D. dependent var | 0.033309 |
| S.E. of regression | 0.016516 | | Sum squared resid | 0.136121 |
| Log likelihood | 1432.313 | | F-statistic | 102.7804 |
| Durbin-Watson stat | 1.984812 | | Prob(F-statistic) | 0.000000 |

### *Change in wages equation*

| Variable | Coefficient | Std. Error | t-Statistic | Prob. |
|---|---|---|---|---|
| ΔEM | -0.003319 | 0.013234 | -0.250834 | 0.8020 |
| ΔEM(-1) | 0.016517 | 0.013441 | 1.228863 | 0.2197 |
| ΔEM(-2) | -0.027776 | 0.014290 | -1.943712 | 0.0525 |
| ΔEM(-3) | -0.022896 | 0.013782 | -1.661358 | 0.0973 |
| ΔEM(-4) | 0.030040 | 0.010985 | 2.734571 | 0.0065 |
| ER(-1) | -0.015084 | 0.040226 | -0.374976 | 0.7078 |
| ER(-2) | 0.063225 | 0.045313 | 1.395304 | 0.1635 |
| ER(-3) | -0.046933 | 0.036021 | -1.302919 | 0.1932 |
| ER(-4) | 0.019599 | 0.041071 | 0.477202 | 0.6334 |
| PR(-1) | 0.038396 | 0.025374 | 1.513232 | 0.1309 |
| PR(-2) | -0.003906 | 0.028824 | -0.135496 | 0.8923 |
| PR(-3) | -0.035146 | 0.028878 | -1.217048 | 0.2242 |
| PR(-4) | -0.022015 | 0.023403 | -0.940705 | 0.3473 |
| ΔW(-1) | -0.476646 | 0.058779 | -8.109090 | 0.0000 |
| ΔW(-2) | -0.451034 | 0.069508 | -6.488916 | 0.0000 |
| ΔW(-3) | -0.288373 | 0.063198 | -4.563027 | 0.0000 |
| ΔW(-4) | 0.271044 | 0.055681 | 4.867833 | 0.0000 |
| Fixed Effects | | | | |
| _AKD--C | 0.001530 | | | |
| _BOP--C | -0.001890 | | | |
| _CAN--C | -0.000587 | | | |
| _GIS--C | -0.001242 | | | |
| _MWT--C | -0.003148 | | | |
| _NEL--C | -0.000891 | | | |
| _NLD--C | -0.000239 | | | |
| _OTG--C | -0.002411 | | | |
| _STH--C | -0.001103 | | | |
| _TNK--C | -0.001595 | | | |
| _WKT--C | -0.001121 | | | |
| _WLG--C | 0.000521 | | | |
| R-squared | 0.521673 | | Mean dependent var | -0.000479 |
| Adjusted R-squared | 0.494833 | | S.D. dependent var | 0.011605 |
| S.E. of regression | 0.008248 | | Sum squared resid | 0.033947 |
| Log likelihood | 1798.940 | | F-statistic | 34.01376 |
| Durbin-Watson stat | 2.017860 | | Prob(F-statistic) | 0.000000 |

# Appendix F: Impulse response functions for individual regions

### *3-variable VAR model with employment share modelled in levels*

Vertical axis displays: percentage
Horizontal axis displays: quarter

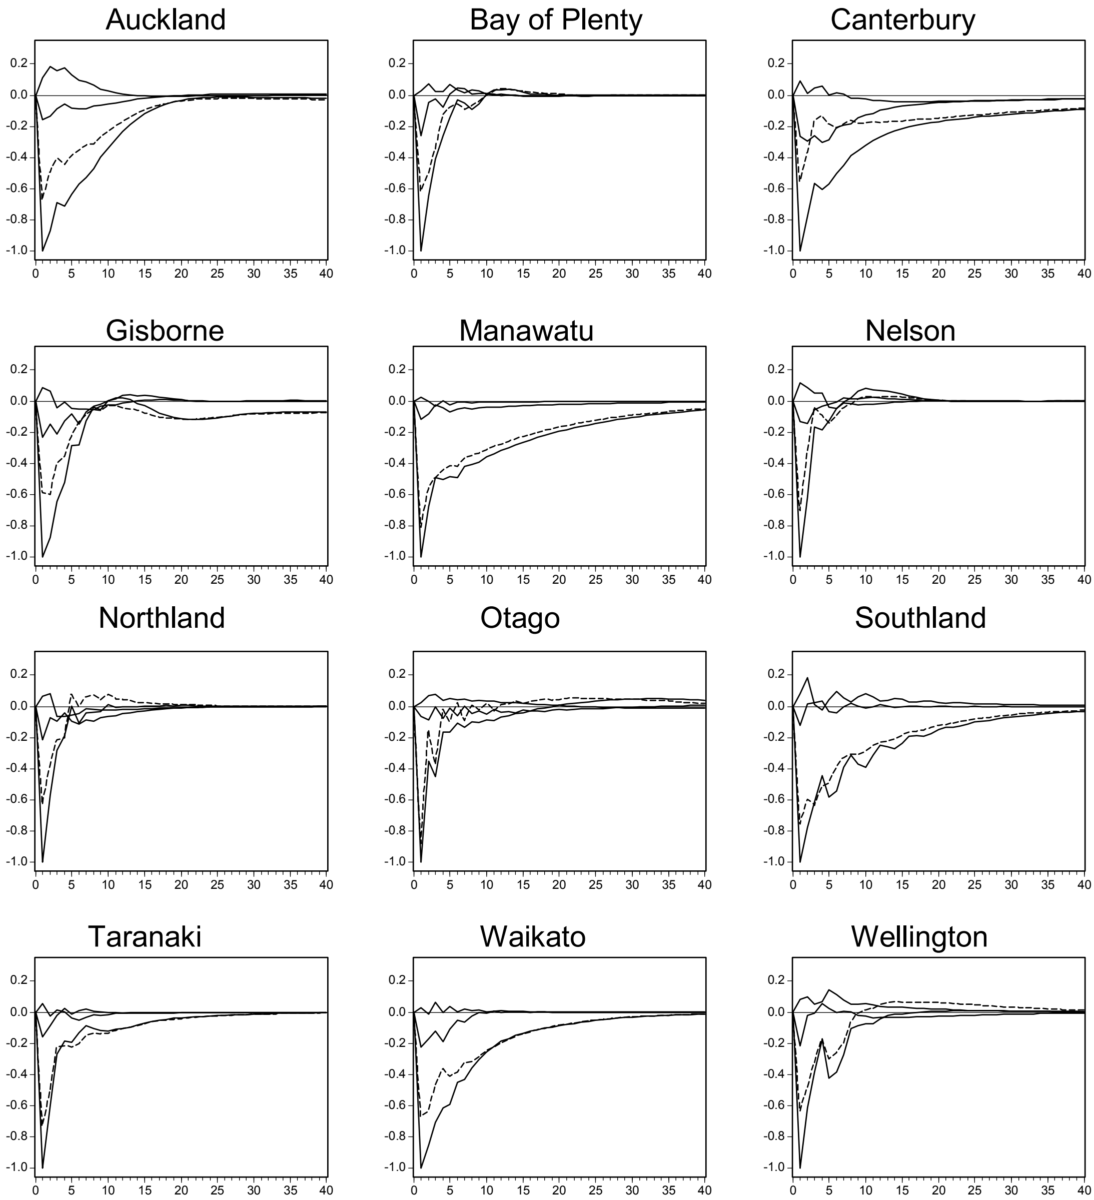


Note: For the three solid lines, from the left of each panel, the upper most line is the unemployment rate response, the middle line is the participation rate response, and the lowest line is the employment response.  The dotted line displays the migration response.