Decomposing Changes in Male Employment in Australia: A Propensity Score Re-Weighting Approach

David Black, Yi-Ping Tseng and Roger Wilkins

Melbourne Institute of Applied Economic and Social Research

University of Melbourne

12 September 2007

Abstract

Using Population Census sample files spanning the period 1981 to 2001 we document changes in the employment-population rate of working-age males in Australia and apply the propensity score re-weighting decomposition approach of DiNardo, Fortin and Lemieux (DFL) (1996) to investigate the sources of changes. Specifically, we investigate the extent to which employment rate changes can be attributed to changes in the socio-demographic characteristics of males. We find that changes in characteristics account for little of the large decline in male employment over the period. However, changes in characteristics are found to be important for population sub-groups. Also notable is that a very large decline in the employment rate of 55-64 year olds with university qualifications is found between 1981 and 1991. We also illustrate that, analogous to matching methods in quasi-experimental evaluation, validity of inferences depends on 'correct' implementation of the re-weighting procedure employed by the DFL method.

Keywords: employment-population rate, decomposition

JEL Code: 21

1. Introduction

Male employment – especially full-time employment – has declined substantially in recent decades in Australia. In 1966, 79% of males aged over 15 years were employed full-time. By 2006, this had fallen to 58% (Figure 1). This decline is only partially accounted for by growth in part-time employment, since the aggregate employment-population rate of males over 15 years of age fell from 83% to 68% over the same period. This is a pattern to some extent experienced by all OECD countries in the post-1970 period, although there is considerable variation in the extent and precise timing of the decline. For example, in the UK, the 15-64 year old male employment-population rate fell from 93% in 1973 to 72% in 1993, whereas in the US the decline was concentrated in the pre-1983 period, since when the male employment rate has been reasonably stable (OECD 1996, 2006).

Figure 1: Male employment rates 1966 to 2006

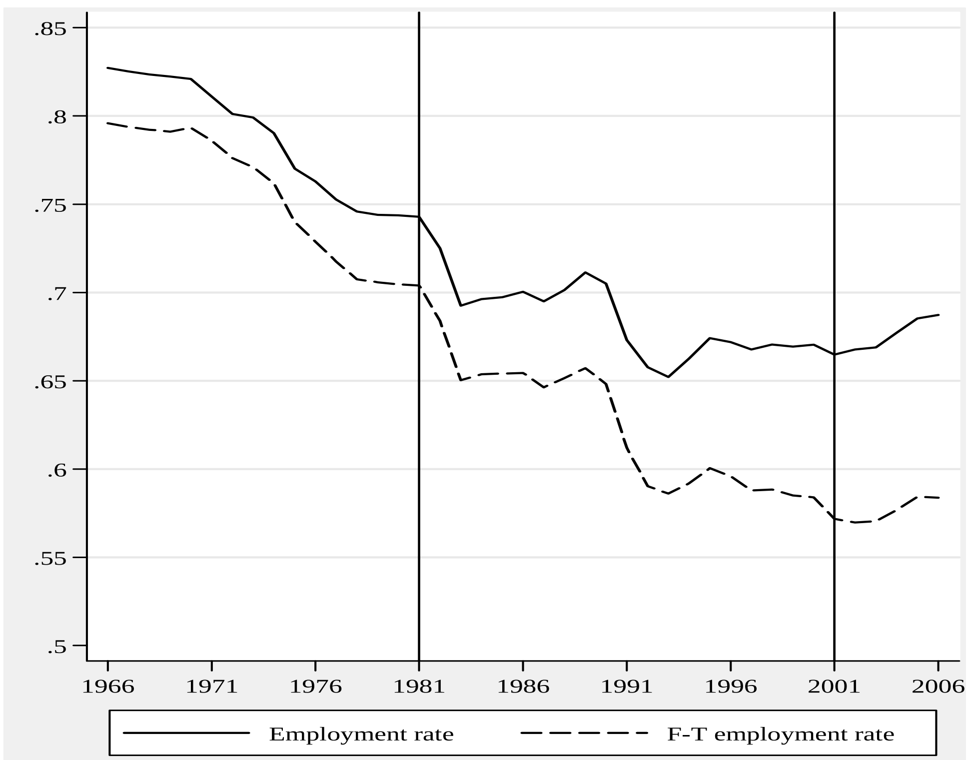

Source: ABS Labour Force Survey.

A range of demand, supply and institutional factors potentially underlie the decline in male employment. Changes in the age and education structure of the male population, changes in male attitudes towards paid work and increased incomes are some of the potential supply factors; while changes in industry structure and production technologies are some of the potential demand factors. Also potentially important is that female employment has increased substantially in recent decades, which could impact on male labour demand via substitution towards female labour, and which could also impact on (and be impacted by) male labour supply via family-level labour supply decision-making.

Given available data, it is extremely difficult to disentangle the effects of the many factors that have potentially driven the decline in male employment. In this paper, we focus on the more limited goal of investigating the nature of the decline in male employment by considering the changes in the observed characteristics of males between 1981 and 2001. Specifically, we decompose changes into those potentially attributable to changes in male personal and family characteristics and those due to changes in employment rates of males of given characteristics.

The 1981-2001 period is dictated by availability of suitable data. While much of the decline in the employment rate pre-dates 1981, it is nonetheless the case that there was a significant decline between 1981 and 2001. There have also been substantial changes in the characteristics of males over this period, suggesting significant potential for much of the decline to be attributable to these characteristics changes. This includes population ageing, increased educational attainment, decreased incidence of partnering and dependent children, and increased educational attainment and employment of partners for those males who are partnered (see Appendix Table A2). In addition to examining aggregate employment rates, we also investigate the role of characteristics changes in producing changes in employment rates for male population subgroups defined by the key socio-demographic characteristics of age, education and partner status.

A number of studies have considered the reasons for the decline in male employment in Australia since the early 1970s, including Stricker and Sheehan 1981, Merrilees 1982 1983, Moir 1982, Miller 1983, Hughes 1984, Borland 1995, Kenyon and Wooden 1996, Connolly and Kirk 1996, Borland 1997, Kennedy and Hedland 2003, and ABS 2003. While a variety of data sources and approaches have been employed, none of the studies takes a decomposition approach. Furthermore, although Borland (1995) and ABS (2003) use micro data to investigate changes in males employment rates, they consider only age, education and, in the case of ABS (2003), cohort effects. Neither study considers the effects of other characteristics, nor allows for interaction effects between these characteristics that may affect the employment rate. In this paper, by contrast, we estimate the male employment rates that might have prevailed after 1981 if the characteristics distribution of the Australian population had remained unchanged from its 1981 distribution, taking a semi-parametric approach that allows for non-linearities and interaction effects between characteristics. We show that it is indeed important to inferences to allow for such interaction effects, something missing from existing studies, which use either cross-tabulations or parametric regression analysis.

The intuition for this finding is that characteristics and their effects can be ‘moving in different directions’ for different population sub-groups. For example, partnering with an employed female may be decreasing over the period for low-educated males, but increasing for highly-educated males. Furthermore, effects associated with partnering with an employed female may differ for these two groups of males. Failing to allow for such interaction effects will, naturally enough, lead to failure to fully identify the effects of the socio-demographic change.

The decomposition technique we adopt is the propensity score re-weighting approach attributable to DiNardo, Fortin and Lemieux 1996 (DFL). This has become a popular decomposition technique in recent years, particularly for decomposing changes over time in earnings or income (e.g., Butcher and DiNardo 2002, Biewen 2001, Daly and Valletta 2006, Hyslop and Mare 2005, Lehmann and Wadsworth 2001). The approach is a generalisation of the Oaxaca (1973) and Blinder (1973) decomposition method (e.g., see DiNardo 2002 for the relationship between Oaxaca/Blinder decompositions and propensity score re-weighting) that is very flexible and has several advantages over alternative decomposition methods, including the ability to decompose measures that are not themselves decomposable.

An issue that we confront, which does not appear to have received any attention in the literature applying this decomposition technique, is the validity of the weights obtained from the re-weighting process. The analogy we draw is with the quasi-experimental program evaluation literature, in which it is standard practice when implementing propensity score matching methods to undertake so-called ‘balancing’ tests, whereby the validity of the propensity score as a measure of comparability of individuals in the treatment and control groups is ascertained. Of course, the end-purpose of the propensity score is quite different for matching methods, where it is used to identify the mean effect of a treatment, in essence by comparing an outcome across persons deemed comparable by their propensity scores. In decomposition analysis, by contrast, the propensity score is used to re-weight observations in one sample such that the distribution of observed characteristics is the same as in another sample. Nonetheless, for both approaches it is important that the propensity score ‘mean the same thing’ for the two samples being examined – that is, correspond to the same expected values of the characteristics variables. In the context of re-weighting, the key requirement is that re-weighting achieves its goal of replicating the distribution of characteristics in the original sample.

We therefore undertake tests of the re-weighting, which we illustrate, via our application to male employment rates, have the potential to be important to inferences. It is our judgement that testing the re-weighting scheme should be standard practice for decomposition analyses based on propensity score re-weighting.

2. Data

The data we use to investigate sources of changes in the male employment rate comprise the publicly released unit record files for the one per cent samples of the Australian population censuses over the period 1981 to 2001. The samples contain individual-level information on a variety of personal and household characteristics, including family structure, country of birth, educational attainment, labour force status at the time of the census and personal and household income.

Several issues in relation to these datasets warrant specific mention. The first issue concerns the timing of the censuses. The censuses to 1986 were conducted on 30th June, while the last three were conducted in early August (see Table A1). There may therefore be seasonal effects for the first two censuses compared with the last three. There may furthermore be 'day of the week effects' present, although all were conducted on a weekday, and generally mid-week. Second, the labour force status information available from the census data does not precisely match the labour force survey. The census is a snapshot of Australia at the date of the census, with labour force status and hours worked defined with respect to the week preceding the census date. An implication is that some employed persons will have reduced or zero hours due to sickness or holidays. We do not attempt to adjust for this.[1]

Table 1 presents employment-population and full-time employment-population rates. Panel A compares the census data with estimates derived from the labour force surveys. The censuses give very similar estimates of the total employment rate, but give consistently lower estimates of the full-time employment rate, in the order of 4 to 5 percentage points lower. This is expected given the approach taken to defining the employment status and hours worked when using the census data. The relative consistency of the differential means that changes over the

[1] Between 3 and 5 per cent of employed persons in each census have hours recorded as 'not stated'. These workers are assumed to have the same distribution of hours worked as workers who reported working hours.

period are quite similar for the two data sources. Panel B presents estimates for the 15-64 years age-range on which we focus. Changes over the period are little-affected, suggesting population ageing explains little of the decline in the employment rate among all males (over 15).

Table 1: Males employment rates, 1981 to 2001

	1981	1986	1991	1996	2001	Change 1981 to 2001
A. Comparisons between labour force survey and census estimates						
Employment rates – Males aged 15+ years						
Labour force surveys	74.40	70.55	66.58	66.93	66.78	-7.62
Census	74.19	70.03	66.07	65.98	67.00	-7.19
Full-time employment rates – Males aged 15+ years						
Labour force surveys	70.34	65.83	60.41	59.01	56.96	-13.38
Census	65.74	61.09	55.90	54.07	53.15	-12.59
B. Effect of age restriction						
Census employment rate estimates – Males aged 15-64 years						
Employment rate	81.30	77.50	73.31	74.08	75.03	-6.27
Full-time emp. rate	72.46	67.93	62.30	61.15	60.02	-12.44

Notes: The estimates from the labour force surveys are not seasonally adjusted and are June estimates in 1981 and 1986 and August estimates in 1991, 1996 and 2001.

3. Methods

3.1 The decomposition method

The employment rate at a point in time, $E_t\left(e\right)$, is a mean of an individual employment rate variable and, following the decomposition approach of DFL, may be expressed as the integral of the mean conditional on a set of characteristics *x* and on a date t_x, $E\left(e \mid x, t_e\right)$, over the distribution of individual characteristics $F\left(x \mid t_x\right)$ at date t_e:

$$\begin{aligned} E_t(e) &= \int_{x \in \Omega_x} E\left(e \mid x, t_e = t\right) dF\left(x \mid t_x = t\right) \\ &\equiv E\left(e; t_e = t, t_x = t\right) \end{aligned} \tag{1}$$

where Ω_x is the domain of definition of the individual characteristics.

The notation in the second line of Equation (1) allows us to express equations for counterfactual employment rates, with t_e denoting the date from which the set of employment rates for each ‘characteristics bundle’ is drawn, and t_x denoting the date from which the distribution of characteristics is drawn. For example, while $E(e; t_e = 2001, t_x = 2001)$ represents the actual employment rate in 2001, $E(e; t_e = 2001, t_x = 1981)$ represents the employment rate that would have resulted in 2001 had the distribution of individual characteristics remained as it was in 1981. This hypothetical employment rate is identified as follows:

$$\begin{aligned} E\left(e;t_e = 2001, t_x = 1981\right) &= \int E\left(e \mid x, t_e = 2001\right) dF\left(x \mid t_x = 1981\right) \\ &= \int E\left(e \mid x, t_e = 2001\right) \psi_x(x) dF\left(x \mid t_x = 2001\right) \end{aligned} \tag{2}$$

where $\psi_x(x)$ is a “re-weighting” function:

$$\psi_x(x) \equiv \frac{dF\left(x \mid t_x = 1981\right)}{dF\left(x \mid t_x = 2001\right)} \tag{3}$$

The equation for the counterfactual employment rate is identical to the equation for the 2001 employment rate except for the function $\psi_x(x)$, so that once an estimate of this function, $\hat{\psi}_x(x)$, is obtained, the counterfactual employment rate can be estimated as:

$$\hat{E}\left(e;t_e = 2001, t_x = 1981\right) = \sum_{i=1}^{n_{2001}} \frac{\hat{\psi}_x(x_i)e_i}{n_{2001}} \tag{4}$$

where n_{2001} is the 2001 sample size and the summation is over observations in the 2001 sample. Applying Bayes’ rule to the ratio $\frac{dF\left(x|t_x=1981\right)}{dF\left(x|t_x=2001\right)}$ gives the following equation for the re-weighting function:

$$\psi_x(x) = \frac{\Pr(t_x = 2001)}{\Pr(t_x = 1981)} . \frac{\Pr(t_x = 1981 \mid x)}{\Pr(t_x = 2001 \mid x)} \tag{5}$$

The probability of being in period t given characteristics x $\left(\Pr(t_x = t \mid x)\right)$ can be estimated non-parametrically, by identifying the proportion of individuals with each characteristic combination at each date, or by a discrete choice model like the probit, with the x’s entered in a reasonably flexible way. For the case where x is a set of dummy variables, DFL is identical to Oaxaca/Blinder with the x variables used in an equation for the employment rate (DiNardo, 2002).

An issue that has not received attention to date concerns the appropriateness of the specification (functional form) of the model used to re-weight observations. The potential exists for the re-weighting function to fail to produce a re-weighted sample with the same distribution of attributes as the original sample, in turn potentially leading to incorrect inferences. In principle, the problem can be avoided by adopting a completely non-parametric specification of the Probit; but in practice the 'curse of dimensionality' means this is usually not viable. The question therefore essentially reduces to "what are the deviations from a non-parametric specification that don't cause any significant differences in distribution of attributes between the re-weighted sample and the original sample?" There is no guarantee that such a parametric specification exists, but of course this does not alter the fact that it is important that a specification achieve its intended function if inferences are to be valid. Inability to find a specification that can be estimated with the available data and correctly-re-weights observations simply means the data available do not support the making of valid inferences.

Identifying an appropriate model specification is an iterative process that contains an element of arbitrariness, and the approach taken will depend on the specific circumstances. Nonetheless, in all cases, at a minimum we should test that the mean value of each element of x is the same in the two samples. As in propensity score matching, t-tests of individual variable means, or Hotelling tests of all characteristics simultaneously, can be undertaken. If there are variables that differ in mean values, the researcher then needs to experiment with modifications to functional form and/or interactions between variables. The process also potentially involves testing whether specific combinations of characteristics – the mean values of specific interactions between variables – are the same in the two samples, if the researcher believes these important to the outcome under consideration.[2] It is, in general, a labour-intensive process.

3.2 Application of the decomposition method to the census data

The unit record data available for the five censuses conducted from 1981 to 2001 allow consideration of the roles played by changes in the characteristics composition of males in terms of age, educational attainment, partner status, the presence of dependent children, immigrant status, student status, English proficiency and partner characteristics (including employment situation).

As has been mentioned earlier, we allow for interactions, which are potentially important to inferences. This could apply to any group of characteristics, but possibly most important in this regard is partner's employment activity. It is likely that growth in female labour force participation has had implications for both male labour supply (particularly if labour supply

[2] The researcher could, of course, include such interactions in the original probit or logit specification, but this is likely to lead to the same identification problems that preclude a completely non-parametric specification.

decisions are made at the family level) and demand for male labour. We furthermore hypothesise that changes in family status have not been uniform across age-education groups, and that the extent to which such family status changes can account for changes in employment also differ by age-by-education group. For example, marital status may have different influences on the labour supply decision for old and young men.

Figures 2 and 3 provide support for these hypotheses. They also illustrate the more general point of the potential importance of testing the re-weighting function and allowing for interactions. Figure 2 shows changes in the marital status composition of the male population are quite variable across education and age groups – the trend decline is not uniform across education-age groups (non-monotonic). This does not of itself have implications for the re-weighting function, but it does in the context of Figure 3, which shows quite different changes in the employment rate of married males across education-age groups.

Figure 2: Proportion of individuals who are married: differences between 1981 and 2001 (1981 minus 2001)

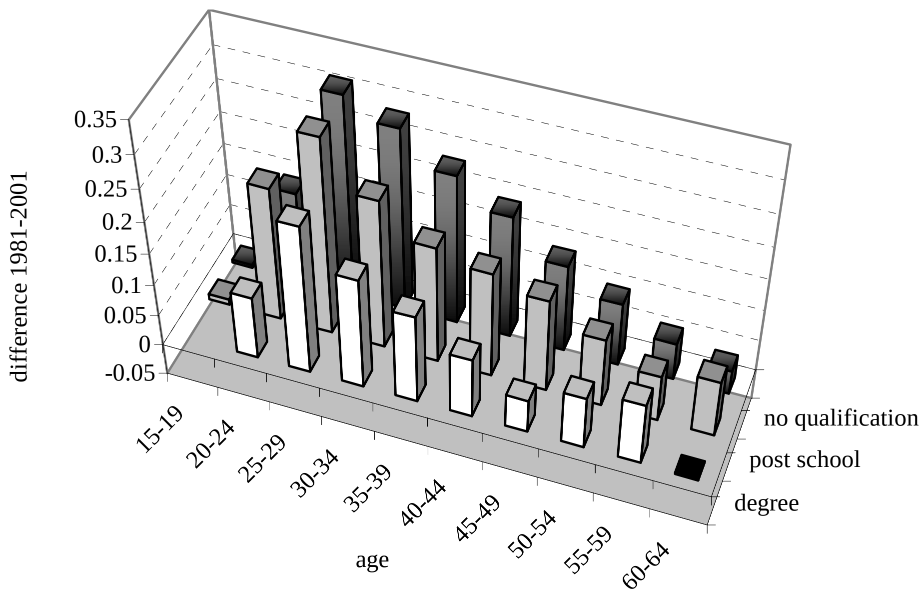

Figure 3: Proportion of married men who are employed: differences between 1981 and 2001.

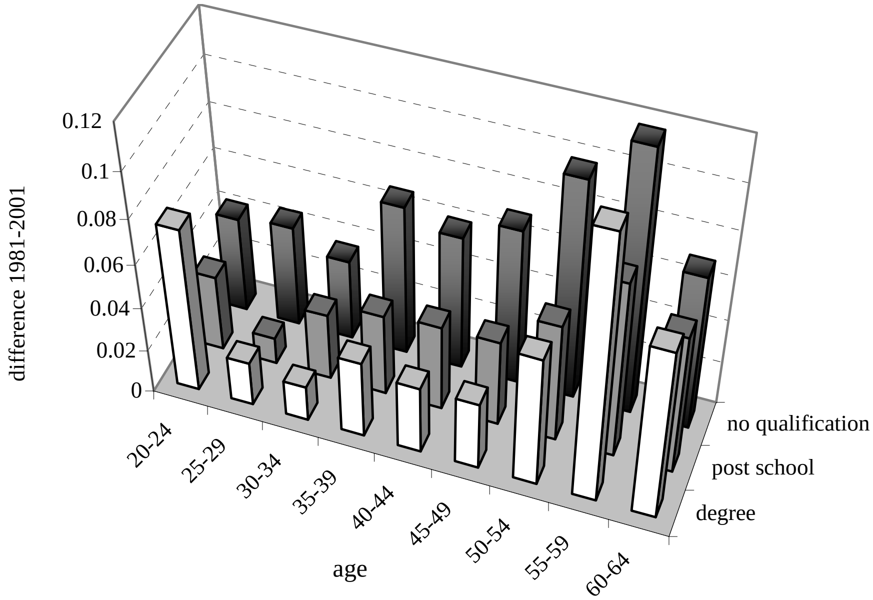

Note: The age group 15-19 years is excluded because of the very small number of observations.

3.3 Choosing the re-weighting function

While it is understood that a 'reasonably flexible' specification of the model used to produce the re-weighting function is required, little attention has been paid in existing research to evaluating the re-weighting procedure. Researchers appear to simply assume that the re-weighting function adopted achieves appropriate re-weighting of samples. However, we argue that it should be a matter of course in decomposition analyses to test the re-weighting function, since there is no guarantee any given semi-parametric specification will achieve correct re-weighting. In principle, a completely nonparametric approach solves this problem, but in practice this is not viable because of the multitude of covariate combinations possible and the occurrence of cells with zero observations.

The approach we take is as follows. We estimate logit models that treat 1981 as the base year, so that each sample is re-weighted to contain the same distribution of attributes as 1981. We first estimate the logit models to obtain the weights using a set of elementary demographic characteristics: age, educational attainment, immigrant status, student status, and the presence of dependent children. In addition, we include a set of interactions of partner status, partner's employment status, and the difference in education level between person and their partner, which produces indicator variables such as 'has full-time employed partner who is more educated' and 'has non-employed partner who is less educated'.

We apply the weights and conduct t-tests for the equality of the mean values for each of these characteristics across the samples. Statistically significant inequalities form the basis for inclusion of additional interaction terms, with decisions based on both the t-test results and (largely intuitive) reasoning on the potential drivers of the inequalities. For example, upon observing that education levels are not equal across the samples we can speculate that this derives from changing levels of educational attainment over time, and so a suitable remedy might be to interact education levels with age. The logit model is then re-estimated with the additional interaction terms to generate new weights. The t-tests are re-run and further interaction terms are added if significant inequalities in means are again observed. This iterative process is repeated until no significant inequalities are obtained.

Once accurate re-weighting has been achieved in the overall samples, we then proceed to consider the distribution of characteristics within groups defined by the key socio-demographic characteristics of age, educational attainment, and partner status. That is, we examine the mean values of characteristics variables within each age group, within each education group, and within each partner status group. When inequalities in the means of characteristics are encountered within these disaggregated samples, an often beneficial strategy is to interact the characteristics in question with each of the categories of the disaggregation. For example, inequalities in the immigrant status variables within each of the age categories may be resolved by interacting these immigrant status variables with each age category. As before, the logit model is then re-estimated (on the entire sample) with the additional interaction terms and the resulting re-weighting function tested. This process may involve numerous iterations, which each consisting of the inclusion extra interaction variables. Once all of these t-tests are satisfied, we produce sets of weights for each of the samples of data for 1986, 1991, 1996, and 2001 that enable us to re-weight these samples such that they exhibit the same distributions of characteristics as the 1981 sample.[3]

Figure 4 illustrates the point that correct re-weighting matters to inferences. It presents, for each education-by-age cell, the difference between the 2001 employment rate holding characteristics constant at 1981 levels using the incorrect (naïve) re-weighting function and the 2001 employment rate change holding characteristics constant at 1981 levels using the correct re-weighting function. That is, for each age-by-education cell, the bar represents:

---

[3] We do not report the interactions resulting from this process because they are too numerous, and differ across years. Details are, however, available from the authors on request. Appendix Table A3 presents the means of the uninteracted variables for the re-weighted samples in 1991 and 2001.

$$\hat{E}(e;t_e = 2001, t_x^n = 1981) - \hat{E}(e;t_e = 2001, t_x^c = 1981)$$

A positive value indicates that the incorrect weighting scheme positively over-states the change. Figure 4 demonstrates that, firstly, correct re-weighting does indeed generally matter, since non-zero differences are evident. Furthermore, the manner in which it matters differs substantially across education-by-age groups.

Figure 4: Employment rate changes 1981-2001, by age-education cell

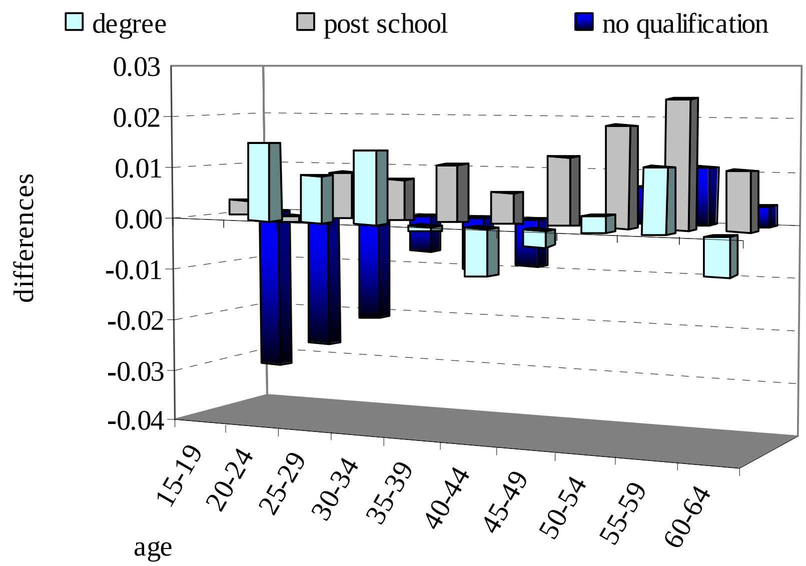

4. Results

Table 2 presents counterfactual employment rate changes when the 1981 distribution of characteristics prevails. Somewhat surprisingly, changes in characteristics appear to account for none of the aggregate decline in male employment. Table 2 also presents counterfactual changes when the re-weighting function is obtained from a 'naïve' Probit specification that contains all the characteristics variables but no interaction terms. Although re-weighting does not achieve the 1981 distribution of characteristics, inferences are in this case little-affected. Nonetheless, it is notable that in moving from the naïve re-weighting scheme to the correct re-weighting scheme, the (employment-increasing) effects attributable to changes in characteristics become slightly smaller.

Table 2: Effects of characteristics changes on aggregate employment rate changes, 1981-2001

	1981	1986	1991	1996	2001
Employment rate					
Raw	0.810	0.768	0.729	0.734	0.746
Change from 1981	-	-0.042	-0.081	-0.076	-0.064
Change from 1981, keeping observed characteristics at 1981 level	-	-0.038	-0.081	-0.086	-0.067
Change from 1981, keeping observed characteristics at 1981 level (basic)	-	-0.041	-0.086	-0.091	-0.073
FT Employment rate					
Raw	0.702	0.657	0.590	0.593	0.579
Changes from 1981	-	-0.045	-0.112	-0.109	-0.123
Changes from 1981, keeping observed characteristics at 1981 level	-	-0.040	-0.105	-0.116	-0.125
Changes from 1981, keeping observed characteristics at 1981 level (basic)	-	-0.043	-0.110	-0.121	-0.132

Note: 'Basic' means re-weighting the sample using inverse predicted probabilities obtained from the logit model with the basic or 'naive' functional form (i.e., the functional form without interactions). This sets of weights does not pass 'balancing' tests.

In Table 3, we consider 'within-cell' changes in employment rates for cells defined by key demographic characteristics. Specifically, for the 1981-1991 and 1981-2001 periods, we present changes in employment rates for individual age groups, education groups and partner status groups, and identify the changes attributable to characteristics changes within these groups. The 'explained' changes are those attributable to changes in characteristics (actual employment rate minus the counterfactual employment rate for the sample). A negative value implies that, keeping characteristics at the 1981 distribution, the employment rate would have been higher than that actually realised – that is, changes in characteristics have acted to lower the employment rate. It turns out that correct re-weighting most matters within groups, rather than for aggregate changes. Sizeable negative effects of characteristics changes are evident for males under 35 years of age and males without post-school qualifications, while sizeable positive effects of characteristics changes are evident for those over 45 years of age and within both partner-status cells.

Table 3: Male employment rates by characteristics – Total change and change due to changes in observed characteristics

	Level in 1981	Change 1981 to 1991		Change 1981 to 2001	
		Explained	Total	Explained	Total
Employment rate					
By Age					
15-19	0.521	-0.066	-0.139	-0.070	-0.108
20-24	0.821	-0.026	-0.120	-0.042	-0.082
25-29	0.885	-0.016	-0.084	-0.023	-0.062
30-34	0.919	-0.003	-0.071	-0.012	-0.056
35-39	0.925	0.006	-0.075	-0.004	-0.078
40-44	0.914	0.009	-0.060	0.004	-0.058
45-49	0.906	0.013	-0.055	0.015	-0.069
50-54	0.872	0.017	-0.076	0.039	-0.070
55-59	0.773	0.025	-0.099	0.056	-0.098
60-64	0.525	0.024	-0.062	0.056	-0.055
By education					
Degree+	0.935	0.001	-0.041	-0.012	-0.044
Post school qualification	0.910	0.006	-0.065	-0.012	-0.073
No PS qualification	0.755	-0.017	-0.105	-0.029	-0.102
By marital status					
Not married	0.693	0.001	-0.075	0.018	-0.033
Married	0.890	0.018	-0.066	0.020	-0.057
Full-time employment rate					
By Age					
15-19	0.411	-0.067	-0.196	-0.069	-0.251
20-24	0.704	-0.040	-0.164	-0.067	-0.191
25-29	0.771	-0.022	-0.105	-0.035	-0.107
30-34	0.817	-0.008	-0.098	-0.012	-0.109
35-39	0.819	-0.002	-0.096	-0.009	-0.109
40-44	0.811	0.003	-0.084	0.008	-0.099
45-49	0.811	0.005	-0.098	0.016	-0.120
50-54	0.754	0.008	-0.096	0.040	-0.102
55-59	0.666	0.014	-0.120	0.048	-0.141
60-64	0.420	0.011	-0.085	0.037	-0.092
By education					
Degree+	0.808	-0.003	-0.048	-0.014	-0.061
Post school qualification	0.806	-0.001	-0.093	-0.017	-0.114
No PS qualification	0.647	-0.024	-0.140	-0.035	-0.180
By marital status					
Not married	0.574	0.000	-0.106	0.028	-0.107
Married	0.789	0.009	-0.096	0.020	-0.098

Note: Change is (end year) minus (start year)

Figure 5 presents similar information to Table 3, but for each of 30 age-by-education cells. It displays the proportion of the total employment rate change between 1981 and 2001 'explained' by changes in characteristics within the cell – i.e., using our earlier notation,

$$p = \frac{\hat{E}(e;t_e = 2001, t_x = 2001) - \hat{E}(e;t_e = 2001, t_x = 1981)}{\hat{E}(e;t_e = 2001, t_x = 2001) - \hat{E}(e;t_e = 1981, t_x = 1981)} \quad (6)$$

where p is evaluated for each cell (rather than in the aggregate).

$p < 0$ implies effects of characteristics changes have operated in the opposite direction to the actual change, while $p > 1$ implies the change due to characteristics changes is in the same direction as the actual change, but is larger. The graphs present a striking picture of the proportion of the employment rate change explained by characteristics changes decreasing in age, particularly for the total employment-population rate.

Figure 5: Proportion of employment rate change within age-by-education cells explained by changes in within-cell characteristics

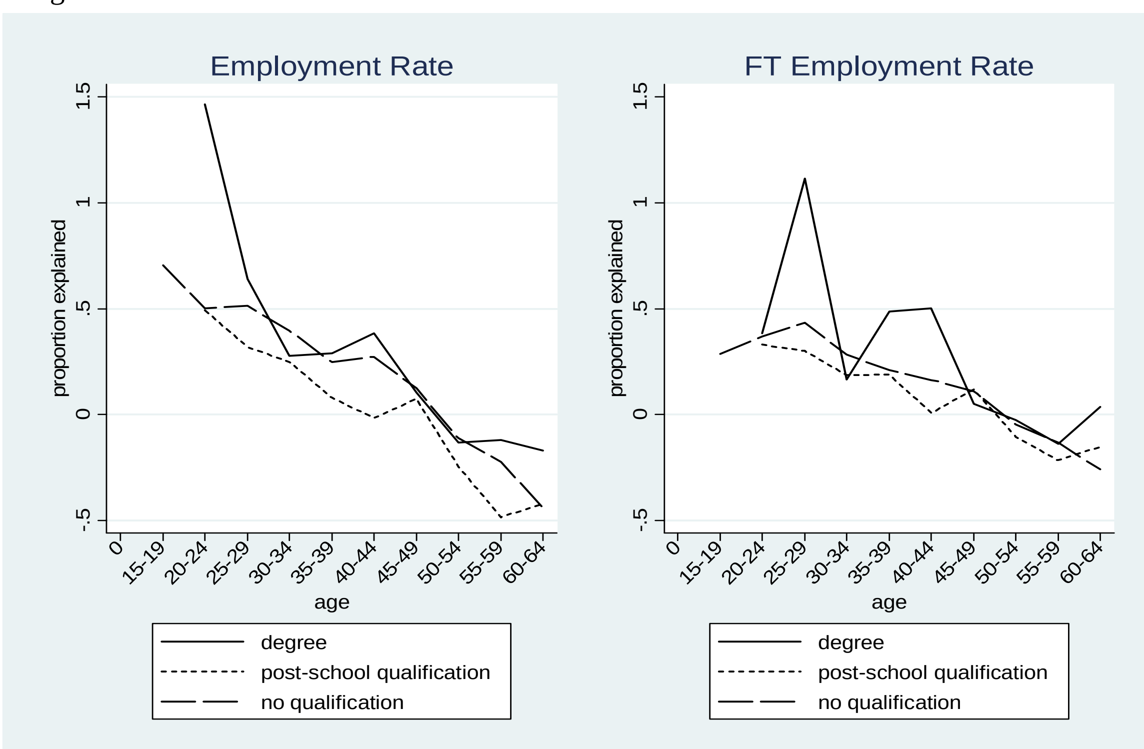

Figure 6 shows changes in the employment rates for each age-education cell when characteristics are held constant at 1981 levels, i.e., $\hat{E}(e;t_e = 2001, t_x = 1981) - \hat{E}(e;t_e = 1981, t_x = 1981)$. We have labelled these as 'unexplained' changes, but they can equally be considered 'real' changes in employment rates, in the sense that they are changes in employment rates that are not simply artefacts of changes in the characteristics composition of the cell.

The most notable feature of the graphs is the big drop for 55-59 year old degree-holders, which reverses the ordering of employment rate changes by educational attainment evident for younger age ranges. That is, once we control for changes in other characteristics of this age-education group, we find a large decline in the employment rate. This is highly suggestive of a labour-supply induced decline in employment among older males, deriving from an increase in early retirement.

Figure 6: Unexplained changes in employment rates

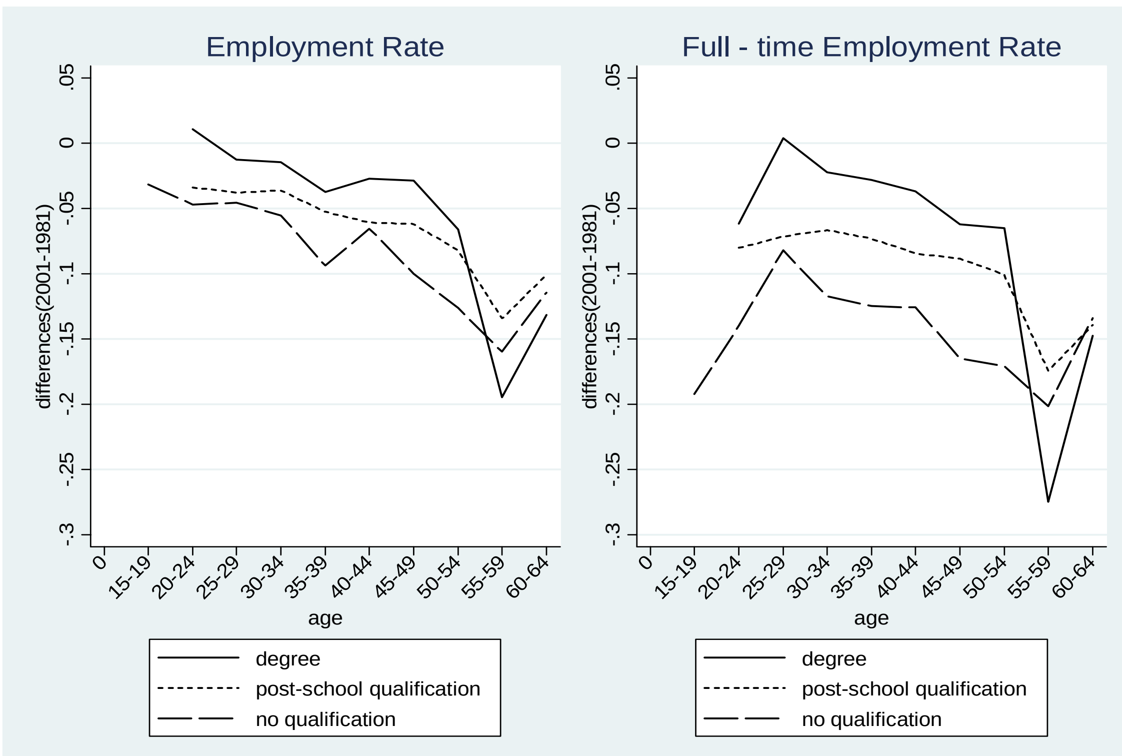

Based on the preceding results, in Figure 7 we break down the full-time employment rate changes by sub-period to pinpoint the timing of the change. It is clear that it occurred in the 1980s, with the pattern not at all evident for the 1991-2001 sub-period.

Figure 7: Unexplained changes in the full-time employment rate, by sub-period

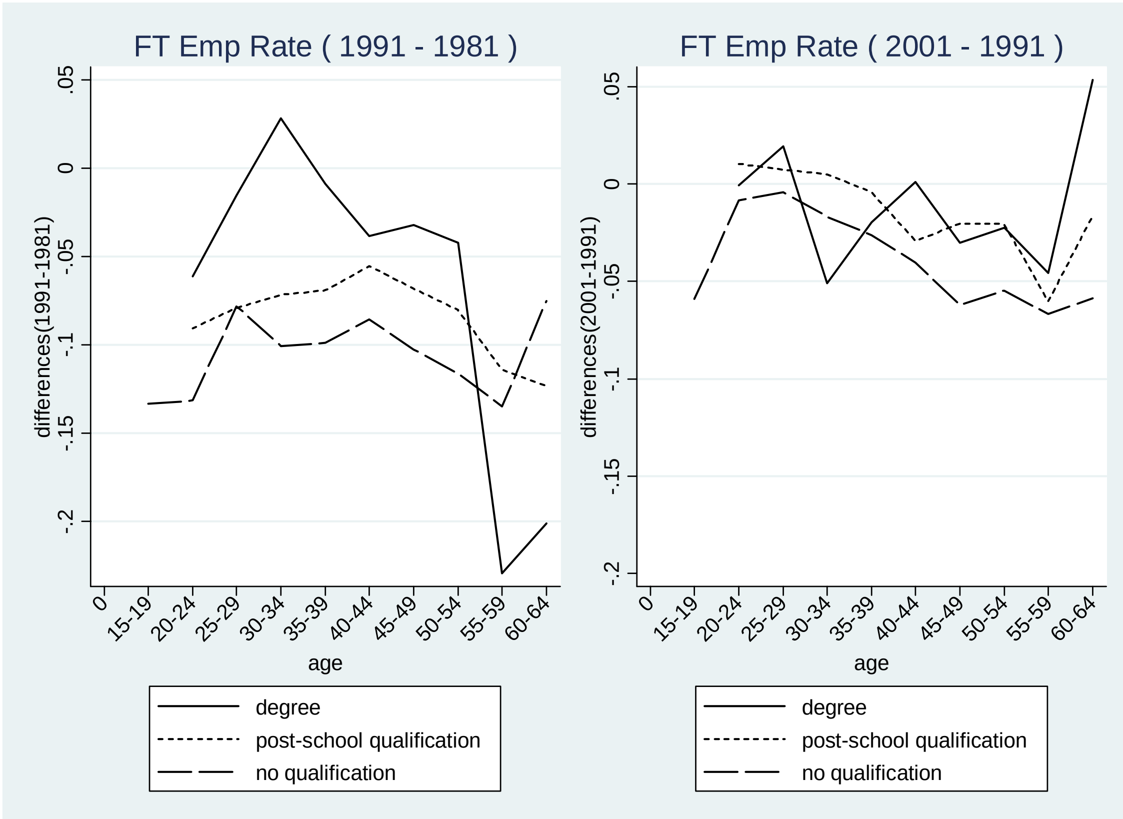

Figure 8 takes a birth cohort perspective on the evolution of employment rates over time once effects of characteristics changes have been accounted for. For each of six cohorts, each comprising males born in a five-year window, the full-time employment rate is plotted against age, holding constant other characteristics at their 1981 distribution using the re-weighting function. The six cohorts are males born in 1936-40, 1946-50, 1956-60, 1961-65, 1966-70 and 1976-80. There at most five data points for each cohort, corresponding to the five censuses conducted between 1981 and 2001.

The figure shows that earlier cohorts tend to have higher employment rates. Interestingly, the decline in the full-time employment rate for the 1936 birth cohort is somewhat similar across the three education groups. Thus, our interpretation the finding presented in Figure 8 is that the negative effect associated with being aged 55-64 years for degree-holders has tended to converge to that evident for other education groups, having been substantially less pronounced in 1981.

Figure 8: Full-time employment rate age profiles by cohort, holding constant other characteristics

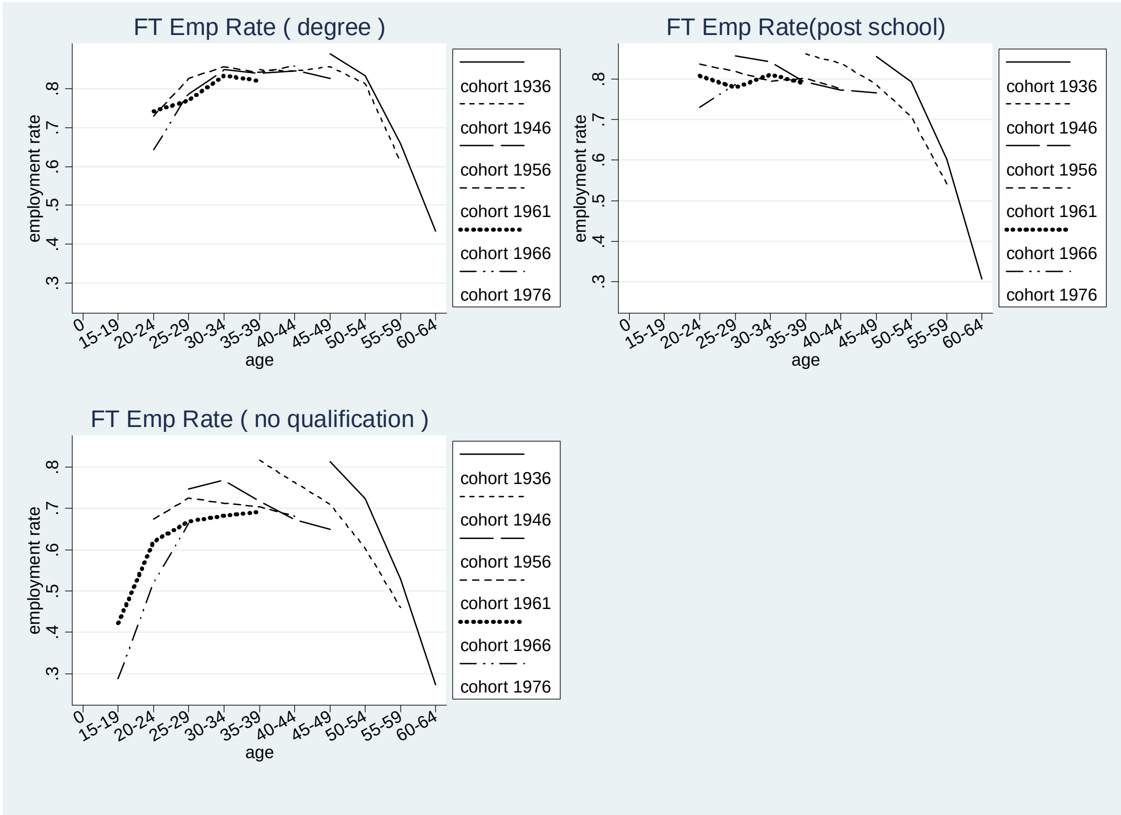

5. Conclusion

The male employment-population rate, and more particularly the full-time employment-population rate, fell substantially over the 1981-2001 period examined in this study. We find that, at the aggregate level, this is not attributable to changes in the socio-demographic characteristics of males. However, within cells defined by age group, educational attainment or partner status, we find that characteristics changes do in fact potentially account for much of the employment rate changes. Our ability to identify these effects to a significant extent derives from our careful attempt to correctly re-weight the samples when using the DFL decomposition method, which results in us accounting for a number of interaction effects. Indeed, a significant element of this source of employment rate change is the changes to partner status and partner educational attainment and employment status, which differ markedly in both their changes and their employment rate effects across males of different ages and educational attainment.

Perhaps the most striking finding of our analysis is the large decline in the employment rate of 55-64 year-old men with bachelor's degrees between 1981 and 1991. At in excess of 20 percentage points, this is an important development, particularly since there is no evidence of a recovery in the employment rate of this group after 1991. Further investigation of this change would seem to be a valuable line of inquiry.

6. References

Australian Bureau of Statistics (ABS) (2003) *Unemployment and Participation Rates in Australia: A Cohort Analysis*, Catalogue No. 6105.0.

Borland, J. (1995) “Male Labour Market Participation in Australia”, *Journal of Industrial Relations*, 37, pp. 587-609.

Borland, J. (1997) “Changes in the Supply of Labour in Australia”, in Industry Commission, *Changing Labour Markets: Prospects for Productivity Growth*, Melbourne, pp. 19-44.

Connolly, G. and B. Kirk (1996) “Determinants of the Full-Time Participation Rate for Males Aged 45-54 Years” in A. Blunn (ed.) *Social Security Early Retirement Seminar*, Australian Government Publishing Service, Canberra, pp. 97-109.

Daly. M.C. and Valletta, R.G. (2006) ‘Inequality and Poverty in United States: The Effects of Rising Dispersion of Men’s Earnings and Changing Family Behaviour,’ *Economica* 73(289) 75-98.

DiNardo, J., Fortin, N. and Lemiex, T. (1996) ‘Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,’ *Econometrica*, 64(5), 1001-44.

Hughes, B. (1984) “Labour Force Participation: What are the Issues?”, in A. Kaspura (ed.) *Labour Force Participation in Australia* (Canberra, AGPS), pp. 1-11.

Kennedy, S. and D. Hedley (2003) “Educational Attainment and Labour Force Participation in Australia” Australian Government Treasury Working Paper.

Kenyon, P. and M. Wooden (1996) “Changes in the Supply of Labour in Australia”, in K. Norris and M. Wooden (eds.), *The Changing Australian Labour Market*, Economic Planning Advisory Commission Paper No 11, AGPS, Canberra, pp. 15-38.

Merrilees, W. (1982) “The Mass Exodus of Older Men from the Labour Force: An Exploratory Analysis”, *Australian Bulletin of Labour*, 8, pp. 81-94.

Merrilees, W. (1983) “Pension Benefits and the Decline in Elderly Male Labour Force Participation”, *Economic Record*, 59, pp. 260-270.

Miller, P. (1983) “On Explanations of Declining Labour Force Participation Among Older Males”, Discussion Paper no. 74, Centre for Economic Policy Research, Australian National University, Canberra.

Moir, H. (1982) “Age Structure of Industries and the Position of Older Men in the Labour Market”, *Australian Bulletin of Labour*, 8(3), pp. 156-76.

OECD (1996) *Employment Outlook 1996*, OECD, Paris.

OECD (2006) *Employment Outlook 2006*, OECD, Paris.

Stricker, P. and P. Sheehan (1981) Hidden Unemployment: The Australian Experience (Melbourne Institute of Applied Economic and Social Research, University of Melbourne).

7. Appendix

Table A1: Census dates

Year	Date of census	Day of the week
1971	30 June	Wednesday
1976	30 June	Wednesday
1981	30 June	Tuesday
1986	30 June	Monday
1991	6 August	Tuesday
1996	6 August	Tuesday
2001	7 August	Tuesday

Table A2: Male Characteristics by Census Year (%)

	1981	1986	1991	1996	2001
Age group					
15-19	13.5	12.9	12.7	11.2	11.1
20-24	13.2	12.5	12.2	11.4	9.9
25-29	12.4	12.5	11.8	11.4	10.4
30-34	13.0	11.8	12.7	11.6	11.2
35-39	10.5	12.1	11.2	11.8	11.3
40-44	8.9	9.6	11.3	11.2	11.3
45-49	7.6	8.1	9.0	10.7	10.7
50-54	7.9	6.9	7.1	8.6	10.0
55-59	7.5	7.3	6.1	6.7	7.9
60-64	5.7	6.3	6.0	5.4	6.2
Place of birth					
Australian born	73.3	73.1	73.3	73.4	73.8
ESB	11.7	10.9	11.2	10.9	10.5
NESB	15.0	16.0	15.5	15.7	15.7
Educational attainment					
Bachelor Degree or higher	6.2	7.3	9.8	13.0	14.7
Other post-school qualification	28.4	29.4	28.5	29.4	31.4
No post-school qualification	65.4	63.4	61.7	57.6	53.9
Family with Dependent Child/ren	49.0	47.8	47.0	44.3	43.1
Partner statsu					
Partnered	59.6	57.0	54.2	51.9	49.8
Single	40.4	43.0	45.8	48.1	50.2
Partner employment and education status					
Partnered, Partner More Educated and Not Employed	19.6	17.7	13.2	12.0	10.7
Partnered, Partner More Educated and Full-Time Employed	10.8	10.2	9.7	10.0	10.1
Partnered, Partner More Educated and Part-Time Employed	7.2	7.6	7.9	8.5	9.1
Partnered, Partner Less Educated and Not Employed	8.5	7.2	5.9	5.7	5.1
Partnered, Partner Less Educated and Full-Time Employed	3.6	3.6	3.7	3.6	3.4
Partnered, Partner Less Educated and Part-Time Employed	3.2	3.6	4.0	4.4	4.4
Partnered, Education or Employment Status ‘missing’	6.7	7.0	9.8	7.7	7.1
Currently studying	13.3	13.7	17.2	16.5	16.5
Sample (N)	44,988	48,003	51,310	53,066	56,253

Table A3: Male characteristics re-weighted to 1981 distribution

Characteristics	1981	1991		2001	
		Basic	Accurate	Basic	Accurate
15-19	13.5	13.4	13.5	13.3	13.5
20-24	13.2	13.2	13.3	13.1	13.3
25-29	12.4	12.4	12.4	12.2	12.4
30-34	13.0	12.9	12.9	13.0	12.9
35-39	10.5	10.5	10.5	10.5	10.5
40-44	8.9	8.8	8.9	8.9	8.9
45-49	7.6	7.5	7.5	7.5	7.6
50-54	7.9	7.9	7.8	7.9	7.9
55-59	7.5	7.6	7.5	7.6	7.5
60-64	5.7	5.8	5.7	5.9	5.7
Australian born	73.3	73.2	73.4	73.4	73.2
ESB	11.7	11.7	11.6	11.5	11.8
NESB	15.0	15.1	15.0	15.0	15.1
Bachelor Degree or higher	6.2	6.2	6.2	6.2	6.2
Other post-school qualification	28.4	28.4	28.4	28.4	28.4
No post-school qualification	65.4	65.4	65.4	65.4	65.4
Family with Dependent Child/ren	49.0	48.5	48.8	48.3	48.9
Partnered	59.6	59.5	59.5	58.7	59.5
Single	40.4	40.5	40.5	41.3	40.5
Partnered, Partner More Educated and Not Employed	19.6	19.6	19.5	19.3	19.6
Partnered, Partner More Educated and Full-Time Employed	10.8	10.7	10.8	10.4	10.8
Partnered, Partner More Educated and Part-Time Employed	7.2	7.2	7.2	7.1	7.2
Partnered, Partner Less Educated and Not Employed	8.5	8.5	8.5	8.4	8.5
Partnered, Partner Less Educated and Full-Time Employed	3.6	3.6	3.6	3.5	3.5
Partnered, Partner Less Educated and Part-Time Employed	3.2	3.3	3.2	3.3	3.2
Partnered, Education or Employment Status ‘missing’	6.7	6.7	6.7	6.7	6.7
Currently studying	13.3	13.5	13.4	13.6	13.3
Sample (N)	44,988	51,310		56,253	