

Job polarisation and earnings inequality in Australia^{*}

Michael Coelli^a and Jeff Borland^b

Department of Economics, University of Melbourne

24 February, 2015

Abstract

We investigate changes in the occupation structure of employment in Australia between 1966 and 2011, and the effect of these changes on the earnings distribution. There has been substantial growth in the employment share of high skill jobs throughout this period. In the 1980s and 1990s the share of middle skill jobs declined, and the share of low skill jobs rose – consistent with what has become known as job polarisation. In the 1970s and 2000s, however, employment shares of both middle and low skill jobs decreased. Changes in the structure of employment by occupation between 1966 and 2011 are consistent with the loss of jobs that were high in routine task intensity. We find that the changes in the occupational composition of employment, and associated changes in average earnings by occupation, have contributed significantly to growth in overall earnings inequality in Australia from the mid-1980s to the mid-2000s.

Keywords: polarisation, technology, inequality, occupation, earnings.

JEL codes: J24, J31

* We thank Mike Keating, Alan Manning and seminar participants at the University of Melbourne and the Economic Society of Australia meeting in Hobart 2014 for useful comments and suggestions. Jing Wen Wu, Mashiyat Khan and Jia Sheen Nah provided painstakingly detailed research assistance. This research was funded in part by the Australian Research Council (ARC) Discovery Project DP140102491 and the University of Melbourne (UoM). The ARC and UoM, however, had no role in the research design or interpretation of results. The views expressed are solely those of the authors, as are all errors and omissions.

a Corresponding author.
Department of Economics,
4th floor, FBE Building,
111 Barry Street,
The University of Melbourne,
Parkville, Victoria, 3010,
AUSTRALIA
mcoelli@unimelb.edu.au

b Department of Economics,
4th floor, FBE Building,
111 Barry Street,
The University of Melbourne,
Parkville, Victoria, 3010,
AUSTRALIA
jib@unimelb.edu.au

I. Introduction

Recent research has identified considerable “job polarisation” over the past three decades in many developed countries: the United Kingdom (UK), the United States (US), Germany, Canada and other European countries (see for example, Autor, 2014; and Goos et al., 2014). Job polarisation refers to a pattern of changes in the composition of employment whereby there is an increase in the share of employment in high skill jobs, a decrease in the share in middle skill jobs, and an increase in the share in low skill jobs.¹

The leading explanation of these changes in the occupational composition of employment is the routinisation hypothesis of Autor, Levy and Murnane (2003). Falling computer prices have caused rapid adoption of computer technology. The new technology has replaced routine cognitive and manual tasks previously undertaken by middle skill workers. At the same time, computer technology has been complementary to the non-routine cognitive and interactive tasks undertaken by high skill workers, thus raising their productivity and in turn the demand for these workers. Non-routine manual tasks undertaken by low skill workers, at present, are not easily undertaken by computers and hence have not been affected by routinisation.² The loss of middle skill jobs, predominantly in the areas of manufacturing production and clerical work, may also have been intensified by the greater scope for offshoring of such tasks.

The first objective of this paper is to ascertain whether job polarisation has occurred in Australia. If the Australian economy is affected by technological change and globalisation in the same way as other developed countries, we might expect to observe this outcome. We measure changes in the skill composition of employment using common measures of skill classified by the occupation category of jobs. Our findings suggest that job polarisation has occurred in Australia, but concentrated in the 1980s and 1990s, rather than happening throughout the period since the early 1970s. Nevertheless, comparisons against the US and Europe, using the same methods of classifying jobs by skill level as used in major

¹ Note that polarisation, in the way the term has been used in this literature, does not imply equal-sized increases in the employment shares of high and low skill workers. Instead, in international studies of job polarisation, the pattern of changes in employment shares by skill level, ordered from low skill to high skill, has generally had the appearance of a ‘J curve’.

² Strictly speaking the routinisation hypothesis implies an increase in high skill jobs and a decrease in middle skill jobs. However, where the decrease in middle skill jobs is greater than the increase in high skill jobs, it follows that the share of low skill jobs will increase.

international studies, suggest that the extent of job polarisation has been similar in Australia to those other regions. Consistent with the routinisation hypothesis, we establish that there have been large declines in the employment shares of occupations that were initially high in routine task intensity.

Our second objective is to investigate whether job polarisation - or occupation changes more generally - have contributed to growing earnings inequality in Australia. We investigate the contributions of both changes in occupational composition of employment and changes in average earnings by occupation to the observed growth in earnings inequality in Australia. We find that these changes can explain a large fraction of the increase in earnings inequality that occurred in Australia between the mid-1980s and mid-2000s.

The remainder of the paper is organised as follows. The related international and Australian literature is discussed in Section II. Occupational changes and the extent of job polarisation in Australia from 1966 to 2011 are investigated in Sections III and IV. In Section V we examine whether the routinisation hypothesis can explain changes in the skill composition of employment in Australia. Trends in earnings inequality in Australia are documented in Section VI. Sections VII and VIII examine how earnings inequality has been affected by changes in the occupational composition of employment and changes in average earnings by occupation, respectively. Section IX concludes.

II. Related Literature

II.A International Literature

Two main approaches have sought to explain how growing use of information technology has affected labour market outcomes in industrialised countries. One approach, developed initially to explain the increase in earnings inequality that occurred in several developed countries in the 1980s and 1990s, is known as the Skill Biased Technical Change (SBTC) hypothesis (Johnson, 1997). Subsequently, evidence of job polarisation prompted the refinement of the SBTC hypothesis into a second approach, based on the routinisation hypothesis of Autor, Levy and Murnane (2003); hereafter, ALM.

The SBTC hypothesis suggests a monotonic positive relation between changes in the relative demand for labour and a worker's skill level. By substituting for tasks undertaken by low skilled workers and raising the productivity of high skill workers, it is argued that increased use of computers has lowered the relative demand for low-skill workers and raised the

relative demand for high-skill workers. Hence the SBTC hypothesis implies that increased use of computers will cause the change in the relative employment share of an occupation to be positively correlated with its required skill level. It follows that the employment share of the highest skill jobs will increase and of the lowest skill jobs will decline.

According to the ALM routinisation hypothesis, computers are able to implement commands that can be coded into routines and thus replace labour in performing routine tasks. More specifically, computer-based technologies are substitutable for workers performing tasks that are cognitive and routine (such as basic clerical jobs) or non-cognitive and routine (such as operation of basic machinery); but not for tasks that are cognitive and non-routine (such as management and medicine), or non-cognitive and non-routine (such as aged care). ALM show that increased use of computers in the US has been associated with a decline in routine tasks undertaken by labour in all industries, occupations and education groups. As well, industries that invested more in computer technology were those where production involved more routine tasks. The ALM routinisation hypothesis has implications for the share of employment by skill level. Where routine clerical and manual jobs are concentrated in the middle of the skill distribution, it is predicted that job polarisation will occur.

Empirical analysis of job polarisation began with a study of the UK by Goos and Manning (2007) and studies of the US by Autor, Katz and Kearney (2006, 2008; hereafter AKK). Goos and Manning (2007) investigate the extent of job polarisation in the UK between 1979 and 1999 using mean wage by occupation as a measure of the skill level of a job. They find significant job polarisation in the UK, and that occupation changes can explain a large fraction of the widening of the earnings distribution in the UK over that period. AKK document how job polarisation occurred in the US in the 1980s and 1990s using the average years of schooling undertaken by workers in each occupation as a measure of skill. Acemoglu and Autor (2011) extend the analysis of job polarisation in the US into the 2000s. They show that the loss of jobs in office/administration, production and operators/labourers occurred in every decade from 1980s, 1990s and 2000s. Autor (2014) further develops that analysis, as well as providing a review of evidence on job polarisation in the US.

Later studies have examined job polarisation in other industrialised countries. Goos, Manning and Salomons (2009 and 2014; hereafter GMS) examine the extent of job polarisation in 16 European countries. They find polarisation to have been quite pervasive. GMS also attempt to disentangle the potential sources of job polarisation by considering the evidence for several

possible explanations: (1) routinisation, (2) offshoring and (3) increasing wage inequality leading to increased demand for low skill service jobs by the highly paid. There is much stronger evidence for the hypothesis that computers have replaced labour in undertaking routine tasks than for offshoring or growth in demand for low-skill jobs in affecting the skill composition of employment. They obtain similar findings using alternative measures of the routine task content of jobs (derived using O*NET and Dictionary of Occupational Titles).

Spitz-Oener (2006) and Dustmann, Ludsteck and Schonberg (2009; hereafter DLS) examine job polarisation in Germany. DLS find significant job polarisation in Germany in the 1980s and 1990s; and that these changes in the occupational composition of employment added to earnings inequality, particularly at the top of the distribution. Spitz-Oener (2006) also finds evidence of job polarisation in Germany using both a direct measure of skill and average earnings by occupation as a proxy for skill. Using a unique dataset that allows her to look at changes in skills used within occupations over time, she finds significant changes in tasks performed by workers within occupations (increases in non-routine cognitive and interactive tasks, and declines in routine cognitive and manual tasks). Green and Sand (2011) examine the extent of job polarisation in Canada. They find a marked pattern of job polarisation. Wage changes by occupation were by contrast monotone, increasing most in high income occupations and least in low income occupations.

In summary, there is widespread evidence of job polarisation in Europe and North America from the 1980s to 2000s. There is some support for routinisation having caused this observed job polarisation. There is only limited evidence of polarisation in earnings by skill level. Nevertheless, for some countries it appears that changes in the composition of employment by skill level has had quite large effects on inequality in the distribution of earnings.

II.B Australian Literature

Several studies have examined changes in the occupational composition of employment in Australia over short periods between the 1970s and the 1990s (Aungles et al, 1993; Gregory, 1993; Dunlop and Sheehan, 1998; Cully, 1999; Wooden, 2001; Pappas, 2001). These studies find mixed evidence regarding whether changes to the demand for labour by skill level have been monotonic (SBTC) or have exhibited a J-shaped pattern (job polarisation). Wooden (2001) suggests that the diversity of findings is explained by the different studies being

undertaken over different stages of the business cycle and using different measures of employment - persons employed or total hours of work.

Studies by Esposto (2011) and Wilkins and Wooden (2014) examine changes in the occupational composition of employment for more recent time periods. Esposto (2011) investigates employment changes in detailed occupation categories using Census data from 1971 to 2006, classifying occupations using the five skill categories constructed by the Australian Bureau of Statistics (ABS) (2006). He finds support for job polarisation rather than for monotonic up-skilling, especially among full-time workers. This finding holds over 1989 to 2009 using either total hours of work or the number of workers as the measure of employment. Wilkins and Wooden (2014, pp.423-24) examine changes in the occupational composition of employment between 1993 and 2013 with data from the ABS Labour Force Survey. They use both the ABS measure and average earnings in an occupation to classify the skill levels of jobs. They claim there is little evidence of job polarisation or that computer use has affected the occupational composition of employment in Australia.

An important contribution was made by Keating (2003) who examines changes in earnings and employment by detailed occupation from 1986 to 2000. He finds that growth in average earnings did not vary across occupations during this period. There was, however, substantial dispersion in changes in employment by occupation. Keating also examines the effect of shifts in the occupational composition of employment on earnings inequality between 1989 and 2000, and shows that these compositional changes explain about one-quarter of the increase in the 90/10 earnings ratio for males and about one-third of the change for females.

We expand on these studies in several main ways. First, we investigate changes in employment by occupation over a longer period (45 years) than the previous Australian studies, and apply alternative measures of skill to study changes in employment by the skill level of jobs. Second, we investigate the routinisation hypothesis by examining the relation between changes in the occupational composition of employment and the routine task intensity of occupations. Third, we investigate the effect on earnings inequality of both changes in employment and earnings by occupation. Fourth, we undertake our analysis separately by gender. Gender differences have received little prior attention in the international literature on job polarisation.

III. Occupation Composition Changes in Australia

III.A – Measurement

Our analysis of the extent of job polarisation in Australia is based primarily on employment data by detailed 4-digit occupation constructed from the five-yearly Australian Censuses from 1966 to 2011. Over this time period, the ABS has changed the way it categorises workers into occupations on several occasions. Occupation data were categorised as follows.

1. 1966 to 1981 - the Classification and Classified List of Occupations (CCLO).
2. 1986 and 1991 – the Australian Standard Classification of Occupations first edition (ASCO1).
3. 1996 and 2001 – the Australian Standard Classification of Occupations second edition (ASCO2).
4. 2006 and 2011 – the Australian and New Zealand Standard Classification of Occupations first edition (ANZSCO).

To create a consistent classification of occupations, we employ occupation link files provided by the ABS. These link files are constructed by taking a subsample of observations in a particular Census, with each individual coded into occupations using both the new and old occupation categories. For example, a link file was constructed using a sample of 1986 Census observations where each worker was classified according to the CCLO and ASCO1 occupation structures. We use these link files, with occupations defined at the 4-digit level, to create occupation concordances over time. These concordances are not one-to-one matches for many occupations, even at the 4-digit level. We use the proportions of workers in a particular ASCO1 occupation allocated to different CCLO occupations in 1986 to re-weight the number of workers in each CCLO occupation in years 1966 to 1981 to yield the number of workers by ASCO1 occupation. These proportions were constructed separately by gender. This proportioning method is essentially the same as that used by Lefter and Sand (2011) for the US, by Green and Sand (2011) for Canada, and by Esposto (2011) for Australia.

III.B Benchmarking against the US and Europe

We begin by benchmarking changes in employment by skill level in Australia against US and European experience. To undertake the benchmarking against the US we use the same skill classifications used by Acemoglu and Autor (2011); and for the comparison with Europe the skill classification from GSM (2014).

For the US Acemoglu and Autor (2011) classify workers into 10 occupation groups. Figure 1 depicts decadal employment growth rates for each of these ten groups for Australia. Overall we find a similar pattern of changes in the occupational composition of employment in Australia to what is shown for the US by Acemoglu and Autor (2011, Figure 12). High growth rates are observed in high skill occupations: managers and professionals in particular, technicians to a lesser extent. Low growth rates are observed for middle skill occupations of office/administration (clerical), production and operators/labourers, although for office/administration, the slower growth did not start until the 1980s. There has been high growth in low skill service jobs of sales, food and cleaning, and in particular, personal service. Protective service jobs also have grown considerably, yet these occupations are probably more correctly considered middle rather than low skill in Australia (see Table 1).³

In interpreting Figure 1 it is important to keep in mind that it is displaying growth rates in employment. Hence, where there is an occupation group with a small initial share of total employment, even a relatively small absolute increase in employment will translate into a high rate of growth. Table 1 (column 3) therefore shows the change in the share of employment for each of the 10 occupation groups between 1966 and 2011. It can be seen that looking at the changes in employment in this way does make a difference. For example, the changes for managers and personal care workers become less pronounced. Nevertheless it is still the case that there is fairly strong evidence of job polarisation. It also is interesting to note that the polarisation is not a recent phenomenon, with middle skill job losses occurring as far back as the 1980s.

For Europe GSM (2014) classify occupations into three categories: low-paying; middle-paying; and high-paying. Table 2 shows changes in employment shares for each of these categories for Australia over a comparable time period to the GSM (2014) study, and also shows the average changes in employment shares for the European countries in their study. In both Australia and Europe there is a common pattern of job polarisation; and the size of changes in employment shares by skill level are quite close.

III.C Australian rankings of skill level by occupation

In this section we undertake a more thorough analysis of the evidence regarding job polarisation in Australia. We do this by developing and applying skill rankings of

³ The occupations experiencing the largest increases and decreases in employment shares over the 1966 to 2011 period by gender are listed in Appendix Table A1.

occupations that are appropriate for Australia. The recent international literature has most often used median or mean wage levels by occupation in some base period as a measure of the skill content of occupations (see for example, Goos and Manning, 2007; GMS, 2009; Green and Sand, 2011; and Acemoglu and Autor, 2011). A concern with this approach is that earnings may be a biased measure of skill: such as where earnings reflect the influence of factors such as union bargaining, discrimination / segregation, and compensating wage differentials. As an alternative AKK (2006, 2008) use average years of schooling by occupation as their skill measure. As already mentioned, some Australian studies have also used an ABS system which classifies jobs into five skill levels based on a judgment regarding the skill level required to complete tasks required in each job.

The main approaches we apply are to measure the skill content of an occupation by: (i) average earnings; and (ii) the skill levels attributed to occupations by the ABS. To rank an occupation's skill level using average earnings, we use average weekly earnings for full-time adults in 1986. This is the earliest year for which we have data on average earnings by occupation at the 4-digit level for all occupations.⁴ Using earnings to measure skill follows the method used by the majority of the international literature on job polarisation. It also assists in the subsequent analysis where we investigate the effect of occupation changes on overall earnings inequality. The ABS has created a measure of the skill content of occupations (ABS, 2006) by allocating all ANZSCO-defined occupations to one of five skill groups based on what skills are required to competently perform the set of tasks required for that occupation. This encompasses: (a) the level or amount of formal education and training, (b) the amount of previous experience in a related occupation, and (c) the amount of on-the-job training. The ABS sought advice from employers, industry training bodies, professional organisations and others when attributing these skill levels to occupations. Highest skill occupations are allocated to skill group 1, while the lowest are allocated to skill group 5. Using the ABS skill measure, production occupations are ranked in the middle of the skill distribution, while operators are ranked close to the bottom.

We do not use years of education as a method for ranking an occupation's skill content. Defining average "years of education" is complicated in Australia due to the need to attribute a "years of education" to different post-secondary qualifications. Rankings of the skill level

⁴ These earnings data were provided directly by the ABS based on their (currently biannual) Survey of Employers. See ABS release *Employee Earnings and Hours*, catalogue no. 6306.0.

of occupations can be quite sensitive to how that is done. For example, Certificate-level study in Australia is often accompanied by apprenticeship training in production trades. Allocating a standard “years of education” of, say, one year for a certificate III/IV results in production and operator occupations (where significant job losses have occurred) being ranked towards the bottom of the skill distribution. However, these same occupations are ranked in the middle of the distribution when skill level is proxied using average earnings.

Table 1 shows each occupation group’s rank by skill level using the two alternative approaches to measuring skill as well as the Acemoglu and Autor (2011) ranking. Most occupations maintain approximately the same rank irrespective of skill measure employed. A notable exception is the operators and labourers group: it is ranked as low skill according to the ABS skill categories, but middle skill using 1986 average earnings. The office/administration group is ranked as low skill using 1986 average earnings, but middle skill using the ABS skill categories. This may reflect the increasing skill level of office/administration jobs over time due to computerisation. Using average income in 2006 (Census income for all full-time employees) or average earnings in 2011 (Employer Survey using all employees) also place this particular occupation group towards the middle of the skill distribution.

III.D Changes in Employment by Skill Level - Australia

This section presents results on changes in employment shares by skill level for the alternative approaches to ranking skill levels of occupations: using average weekly earnings; and using the ABS classification of skill level.

We begin with the approach that uses average weekly earnings as the measure of skill. Figure 2 depicts changes in employment shares by skill level for each decade between 1971 and 2011. In this figure, 1986 average weekly earnings of full-time adult employees by 4-digit ASCO1 occupations are used to measure an occupation’s skill ranking. Workers are grouped into skill quintiles using employment in 4-digit occupation categories as weights (derived from 1986 Census total employment data). We then examine changes over time in the share of employment accounted for by each of the 1986 quintiles. For example, suppose that in 1986 the top quintile of employment consisted of workers in the managerial and professional occupations. Then the top quintile in Figure 2 would show the change in the share of employment accounted for by those occupations between 1971 and 2011.

Employment changes were constructed using Census data for each decade. For example, the 1970s employment changes are constructed based on 1971 and 1981 data.

Strong growth is observed in the share of employment in the top quintile in each of the past four decades. The employment shares of the second lowest and middle quintiles have decreased in all decades. For occupations in the lowest quintile, the employment share fell in the 1970s and 2000s, was essentially unchanged in the 1990s, but rose considerably in the 1980s.

Figure 3 depicts changes in employment shares by the skill levels attributed to each 4-digit occupation by the ABS. Strong employment share increases are observed in every decade among occupations in the highest skill category. Smaller increases are also observed in the second highest skill category. Employment shares of the lowest three skill categories have declined in the vast majority of cases. In all decades, declines have been largest among the middle skill group, where most production occupations are located. Thus there is some evidence of job polarisation when occupations are defined in this manner, with the strongest evidence in the 1990s, and again very little evidence in the 2000s.⁵

Figure 4 depicts changes in employment by skill category over the entire period from 1966 to 2011, using initial earnings to define skill (left panel) and using the ABS skill categories (right panel). In both cases, employment growth is strongest in the highest skill category. Employment share losses are greatest in the second lowest skill category when skill is measured using initial earnings, and in the middle category when skill is measured using the ABS categories. While there are employment share declines in the lowest skill category irrespective of skill definition, these losses are more moderate than in middle skill categories. Thus employment changes in Australia can be described as having a J-curve with respect to

⁵ Wilkins and Wooden (2014) examine changes in employment shares from 1993-2013 using the ABS skill categories and data on employment from the ABS Labour Force Survey (LFS). In Appendix Table A2 we compare Wilkins and Wooden's findings with analysis for 1991 to 2011 using Census data applied to the ABS skill categories. Our analysis using Census data obtains similar findings – except for the bottom quintile where Wilkins and Wooden find a decreased share of 3.4 percent compared to a decrease of 1.8 percent using Census data. The difference appears to be due to the LFS data showing much larger declines in employment of clerical and office support and factory workers than the Census data in the 1990s. In the Census data these occupations experience the largest declines in employment in the 1980s. Hence a comparison with Wilkins and Wooden leaves us comfortable with our argument that job polarisation occurred in Australia in the 1980s and 1990s.

skill, rather than monotonic up-skilling. The gains at the top have outweighed the losses in the middle, thus the employment share at the bottom has declined moderately.

Summary

Our analysis has examined changes in the skill composition of employment in Australia from 1966 to 2011. For the whole of this period we find that the employment share of high skill jobs has increased substantially, there has been a large decline in the share of middle skill jobs, and a small decline in the share of low skill jobs. Looking within sub-periods we find evidence of job polarisation in the 1980s and 1990s; whereas the pattern of changes in employment shares in the 1970s and 2000s is monotonic and hence more consistent with the SBTC hypothesis. The findings are somewhat sensitive to the measure of skill that is used.

IV. Occupation Composition Changes – The Story by Gender

In this section we investigate changes in the occupational composition of employment separately for males and females. The same alternative approaches are used for measuring the skill level of an occupation.

We commence by calculating employment changes by skill level using average earnings by occupation as the skill measure that are disaggregated by gender. The results are presented in Figure 5. Note that this figure uses a measure of earnings by occupation constructed by taking employment-weighted averages of male and female average earnings; that is, it does not incorporate separate wage measures by gender.⁶ Figure 5 suggests that job polarisation has been a feature of occupation changes for males only, and polarisation has occurred in essentially all four decades from the 1970s to the 2000s. Changes in employment shares for females, however, increase monotonically with the skill level of jobs.

In Figure 6 we present changes in employment shares by occupations grouped according to the ABS skill categories. Using this measure there appears to be job polarisation for males in the 1990s with weaker evidence in other decades, and more monotone relationships in all decades for females. Differences by gender are, however, much less apparent in Figure 6 than in Figure 5, where the skill level of an occupation was ranked according to initial earnings.

⁶ The decision to use average earnings over both genders was in part due to necessity. Average earnings by 4-digit occupation were not provided for both genders in a number of cases. If the number of workers of a particular gender in an occupation was small, no average earnings measure was provided by the ABS.

Our analysis by gender therefore shows that job polarisation has been predominantly a male phenomenon. Using the earnings measure of skill we find job polarisation for males in all decades; although with the ABS skill measure there is strong evidence only for the 1990s. For females none of the measures of skill suggest job polarisation at any time.

Why has there been a difference in the experiences of males and females? To answer this question it is not possible to simply interpret the findings for males and females as the effect of different changes in the demand for labour by skill by gender. Changes to the skill composition of employment within each gender group can be caused by changes in the occupational composition of employment, but also by changes in the gender composition of employment. The former factor has already been described. The latter factor can occur where males and females have different initial skill compositions of employment, and there is a change in the gender composition of employment within occupations. This describes the situation in Australia from the 1970s onwards. In 1971 females were concentrated in low skill occupations, whereas males were more evenly spread across low, medium and high skill occupations. As well, over the period from 1971 to 2011 the female share of employment increased from 32.8% to 47.4%. Suppose that the overall occupational composition of employment remained unchanged while the growth in the female share of employment occurred. Even where, for example, the growth in female employment was spread evenly across the skill distribution, given that females began the period heavily over-represented in the lowest skill occupations, there would be a reduction in the share of female employment in low skill jobs and an increase in high skill jobs.

We undertake a simple decomposition analysis to establish the contributions of each factor to occupation changes by gender (details are provided in the Appendix). We construct the decompositions using the following thought experiments. First, suppose that the overall composition of employment by occupation is changed according to what is observed in the data, but the proportion of females within each occupation remained the same (constant gender mix). We interpret the predicted changes in employment shares by skill category for each gender group due to the change in the overall occupational composition as the effect of changes in the “occupation structure”. Second, suppose that the overall occupational composition of employment is held fixed, but the proportion of females within each occupation is changed according to what is observed in the data. We interpret the predicted

changes in employment by skill category within each gender group due to changes in the proportion of females in each occupation as the effect of changes in “gender composition”.

Results from this analysis (see Appendix Figure A1) reveal differing stories for males and females. Changes in employment shares by skill level for males have been primarily driven by changes in the overall skill composition of employment. For females, however, both changes in the overall skill composition of employment and changes in the gender composition of employment are about equally important in explaining changes to the skill composition of jobs. Female employment was heavily concentrated in occupations at the bottom of the skill distribution in the 1960s and 1970s. Some of these occupations experienced large job losses in the following years (effect of change to overall composition of employment); and at the same time, as females have come to account for a larger share of total employment, this growth has been spread throughout the skill distribution (effect of change to the gender composition). Both factors therefore appear to have driven the decreasing share of female employment in low skill categories. Male employment, by contrast, was more evenly distributed by skill level in the 1960s and 1970s, and with occupations that have experienced job losses concentrated in the middle of the skill distribution, changes in the occupational composition of male employment exhibit a greater degree of job polarisation.

The explanation for the differing findings on job polarisation for males between the alternative skill measures appears to be due to differences in the classification of specific occupations between the measures, specifically operators and labourers. Operators and labourers are classified as middle skill by the (1986) earnings-based measure of skill, yet are low skill when the ABS measure is used to define the skill level of an occupation. Operators and labourers account for a large fraction of male employment, and there was a large decline in employment of operators and labourers in the 1970s and 1980s. It follows that male employment will show a decline in middle-skill jobs using the (1986) earnings measure, compared to a decline in low-skill jobs using the ABS measure.

V. Task Content of Occupations

The routinisation hypothesis is the idea that computer technology provides a cheap substitute for labour in completing routine tasks, raises the productivity of workers in high skill jobs requiring the ability to undertake non-routine abstract tasks, and is not yet able to undertake

non-routine manual tasks at low cost. Thus occupations that initially involved routine tasks would be predicted to have decreasing shares of employment, occupations requiring the completion of abstract tasks to have rising employment shares, and occupations undertaking non-routine manual tasks to have had stable or increasing employment shares.

ALM (2003) test their routinisation hypothesis by comparing changes in the relative demand for workers with the capacity to complete different tasks – especially routine versus non-routine tasks. Measures of task requirements for workers in each occupation were constructed using information from the US Dictionary of Occupation Titles (DOT). The DOT is a comprehensive categorisation and description of occupations at a detailed level in the US, and includes information on the skills and attributes required in these occupations. These requirements are interpreted as revealing the extent of demand for labour with the capacity to complete specific types of tasks in that occupation. For example, the level of mathematical aptitude required in a job is interpreted as revealing the level of demand in that job for workers who can complete abstract cognitive tasks. ALM (2003) used information from the fourth edition (1977) and fourth edition revised (1991) versions of the DOT, which include information on over 12,000 individual occupations.

There is no equivalent of the DOT for the Australian labour market. Therefore, we used the US DOT as the basis for our analysis of the effect of computer technology and routinisation on the skill composition of employment in Australia. This follows the recommendation of Autor (2013) and GSM (2014) to use an off-the-shelf measure of the task content of occupations.

Our first step is to link the US DOT to the Australian occupation structure. We use DOT measures of the requirements of each occupation that were constructed by the US National Academy of Sciences for a sample of workers taken from the 1971 US Current Population Survey.⁷ The measures are provided for over 3,800 individual DOT occupations. Each of these DOT occupations was matched into the Australian ASCO1 occupation structure at the 4-digit level (282 occupations) using job descriptions from the DOT and from the ABS. The measures of task requirements within each of the 282 ASCO1 occupations are calculated as the average of each of the DOT occupations allocated to an ASCO1 occupation. This

⁷ This information is provided by the Inter-University Consortium for Political and Social Research based at the University of Michigan (ICPSR 7845). The measures were often constructed using text recognition procedures based on the detailed descriptions for each occupation provided in the DOT Fourth Edition (1977).

averaging was done separately by gender.⁸ Table 3 lists the five DOT categories for which we have collected data on job requirements by occupation, along with the type of task for which each category is interpreted to reveal the level of demand. Our choice of categories and interpretation of the categories follows ALM (2003); and we follow AKK by combining the first two measures into one measure of abstract tasks, and combining the second two measures into one measure of routine tasks.⁹ This results in separate measures of abstract, routine and manual tasks.

Our second step is to create measures of the average task content of jobs in Australia. Our approach follows AKK. For each task measure we assign a percentile score to each 4-digit occupation according to its rank-level on that measure. We do this using 1966 employment by occupation and gender. By construction, for each task measure the average percentile score in 1966 is 50. We also construct the Routine Task Intensity (RTI) index, taken from Autor and Dorn (2013):

$$RTI\ Index = \ln(routine) - \ln(manual) - \ln(abstract) ;$$

and we transform this RTI Index into percentiles in the same way using 1966 employment weights.

We then calculate average percentile scores for each task measure and for the RTI Index for each Census year up to 2011, again using employment by occupation and gender in those years as weights. Hence the average percentile measure (APM) for each task in year t is:

$$APM_t = \sum_{i=1}^n p_i \cdot w_{it}$$

where the p_i are occupation by gender percentile scores in 1966 and the w_{it} are the employment shares for each occupation / gender group i in year t .

Where an average percentile measure (APM) for a particular task increases over time, this shows that changes to the occupational composition of employment are tending to increase the relative demand for labour to perform that task; that is, a relatively greater share of total

⁸ Males and females were often employed in different DOT occupations within the ASCO1 categories.

⁹ We combined the two measures by first putting both measures on a zero to ten scale, then simply averaging them.

employment is in occupations that rank above average on the task measure. By contrast, a decrease in APM reflects a decrease in relative demand for labour to perform that task.

Figure 7 shows the evolution of the change in demand for labour to complete each type of task using the APM. A strong increase in demand for labour to complete abstract tasks is clearly observed, as is a strong decline in demand for routine tasks. A decline in demand for labour to complete manual tasks is also observed up to 2001, which is consistent with the findings of ALM (2003). Many production and operator jobs had high manual task measures, and underwent large decreases in employment over the period. The manual task measure has remained relatively stable over the 2001 to 2011 period. Finally, the RTI Index falls considerably over the period.

In Figure 8 we show changes in the task content of employment separately for females (left panel) and males (right panel). Trends in the task content of employment for each gender have been similar over the sample period. However, there are gender differences in the levels of task content. First, routine task content was initially higher among females, while abstract and manual task content are higher among males. As a consequence, the RTI index is much higher for females than for males. These findings for Australia are consistent with the ALM (2003) results for the US. Second, the decline in routine tasks and the increase in abstract tasks are both greater among females. Third, the reduction over time in manual task content is more muted within each gender group than the overall decline depicted in Figure 7. The larger overall decline is due mainly to the growth of female employment relative to male employment from 1966 to 2011 (from 30.5% to 47.4% of total employment), and that female manual task content is at a considerably lower level than for males throughout the period.

In Figure 9 we show the relationships between our measures of task content and a measure of the skill content of occupations (using average earnings to measure skill). The percentiles of the task measures (using the initial 1966 employment weights) for each occupation have been graphed by order of the skill level of each occupation, and smoothed. As expected, the level of requirement to complete abstract tasks increases with the skill level of occupations. The level of requirement to complete routine tasks falls slightly with skill level at the bottom of the distribution, rises to a peak at the 60th percentile of the skill distribution then falls again at high skill levels. The level of requirement to complete manual tasks rises with skill level up to around the 70th percentile, and then falls. This again is a function of production occupations being reasonably well-paid in Australia, particularly up until the 1980s. The RTI

Index is relatively stable with occupation skill level until the 40th percentile, after which it falls considerably with skill.

Have the changes in the skill composition of employment been related to changes in the relative demand for labour to perform the different tasks? Descriptive analysis suggests that a relation does exist. Figure 10 displays changes in employment by occupation across the distribution of the RTI index. In the left hand panel, changes in employment shares over the entire period from 1966 to 2011 are plotted as a function of the initial RTI Index distribution percentiles. Growth is strongest among occupations at the very bottom of the RTI Index distribution, and is weakest among occupations at the very top of the distribution. In the right hand panel, employment changes are broken down by decade. While employment growth is negatively related to RTI in all decades, the relationship is most negative from the mid 1980s to the early 2000s. This is consistent with our finding that job polarisation was most evident in the 1980s and 1990s in Australia.

To further investigate this issue we estimate the following simple model of the relation between occupation-level employment and the RTI Index in Australia:

$$\ln E_{it} = \beta \cdot RTI_i \cdot t + \alpha_t + \gamma_i$$

Where E_{it} is employment in occupation / gender group i in year t , RTI_i is the RTI Index for each occupation / gender group i , and t is a simple linear time trend. We also include time fixed effects, α_t and occupation / gender fixed effects, γ_i .

The results of these regressions are presented in Table 4. Column (1) reports results from a pooled model for the whole of our sample period from 1966 to 2011. Column (2) tests whether there is a significant difference between the effect of the RTI Index on male and female employment. Columns (3) and (4) seek to get closer to the analysis undertaken for Europe in GMS (2014, Table 3) by pooling genders and restricting the sample period to 1991 to 2011. In each model the interaction of the RTI Index and the time trend is found to be significantly negatively related to the level of employment in an occupation. We find no statistically significant difference in this effect between males and females (see column (2)). Interestingly, when we seek to get as close as possible to the approach of GMS (2014), we obtain almost exactly the same estimated effect of the RTI Index variable on log employment. Note that, to assist comparison with the estimates in GMS (2014), the coefficient estimates reported in Table 4 were all multiplied by 100, and the RTI Index was

first transformed into a mean zero, standard deviation one variable. To interpret the size of these estimates, an occupation that is one standard deviation higher in the RTI Index distribution has annual employment growth that is on average 0.9 to 1.1 per cent lower over the period.

Overall, our findings are quite consistent with routinisation being an important influence on changes in the composition of employment in Australia. First, we have shown evidence of an association between growth in the level of employment in an occupation and the level of routinisation in tasks performed by workers in that occupation. Second, it appears that the timing of the most substantial job polarisation in Australia matches with the period where there are the largest changes in demand for workers to undertake routine tasks.

VI. Earnings Inequality

The objective of the following two sections (VII and VIII) is to understand how changes in the occupational composition of employment and changes in average earnings by occupation have affected earnings inequality in Australia. In this section, we set the scene by documenting the trends in earnings inequality that we wish to understand. In Australia, there are three main sources of information on earnings inequality:

1. Labour Force Survey (LFS) measures provided by ABS - individual worker reports.
2. Employer Survey (ES) measures also provided by ABS - employer reports.
3. Income Distribution Survey (IDS) micro-data - individual worker reports.

The first two sources provide aggregate-level information on the distribution of earnings via ABS publications.¹⁰ The third source provides data on individual earnings from which inequality measures can be constructed.

To begin, we employ the third data source to document changes in the distribution of earnings for male and female full-time employees using information on weekly earnings.¹¹ Figure 11 depicts real earnings growth at each percentile of the earnings distribution over

¹⁰ Labour force Survey information is provided in *Employee Earnings, Benefits and Trade Union Membership*, ABS catalogue number 6310.0. Employer Survey information is provided in *Employee Earnings and Hours*, ABS catalogue number 6306.0.

¹¹ It was not possible to generate hourly earnings information for all employees with this data over time, as hours of work in several years is only provided in broad categories.

three periods: 1982 to 1990,¹² the 1990s and the 2000s. Note that the top and bottom three percentiles are not presented due to excessive noise at the very ends of the earnings distribution.

For the 1990s and 2000s, real earnings growth is essentially monotonically increasing in earnings percentile; that is, earnings growth is strongest at the top of the earnings distribution and weakest at the bottom. This is most evident for males: earnings grew by approximately 30 per cent over each of these decades at the top of the distribution, while at the bottom of the distribution there was essentially no growth in the 1990s and growth of only 10 percent in the 2000s. For females we see growth of approximately 15 per cent in earnings at all percentiles below the median, with increasing growth above the median. At the top of the distribution, growth reached approximately 25 per cent in the 1990s and 35 per cent in the 2000s.

In the period from 1982 to 1990, workers below the 25th percentile had positive but still moderate real earnings growth, whereas for other workers there was little change in real earnings. This was the period of the Wages and Incomes Accord in Australian wage-setting, where the objective was to achieve overall wage moderation while allowing some earnings growth for those at the bottom of the earnings distribution (Borland, 1991).

The international earnings inequality literature, and the recent job polarisation literature in particular, tend to examine specific measures of earnings inequality: the 90/50 percentile log wage gap (upper tail inequality), and the 50/10 percentile log wage gap (lower tail inequality). Figures 12 and 13 depict these two measures using all three sources of information on weekly earnings inequality listed above, for full-time and all workers respectively.

Among full-time workers (Figure 12), upper tail inequality has risen steadily for both genders over the post-Accord period, from the late 1980s to 2011. Upper tail inequality was relatively stable during the 1980s, with some evidence of rising inequality among females prior to the Accord. Over the whole period from 1975 to 2011, the magnitude of increase in upper tail inequality is similar for males and females, with the 90/50 percentile log wage gap increasing by 0.2 log points. Note that upper tail inequality is larger for males than females.

¹² Note that earnings data from the Income Distribution Survey micro-data file for 1982 is on current weekly earnings rather than usual weekly earnings, as it is in all other years. This difference is only of small consequence among full-time workers.

Evidence on lower tail inequality among full-time workers is more mixed, with the alternative sources of earnings information giving somewhat different results – although all show widening inequality since the mid-1990s. For full-time females the LFS and IDS reveal a narrowing in the lower tail from the early 1980s to the mid-1990s, then a constant or slowly widening lower tail since then; whereas the ES reveals a steadily widening lower tail from the mid-1980s onwards.¹³ For full-time males, the ES also reveals a steady widening in the lower tail over the period; and by contrast, the LFS and IDS reveal relatively stable lower tail inequality from the beginning of the 1980s to the mid-1990s, then slowly rising inequality after that time.

Figure 13 provides information on earnings inequality (weekly earnings) among all workers, i.e. part-time workers are also included. This affects the lower tail inequality measures in particular. The upper tail inequality measures reveal steady increases for both males and females over the period, with more moderate increases in the 1980s. Lower tail inequality increased sharply from 1990 to 2000, then flattened or fell to the end of the period.

To summarise, earnings inequality increased markedly in Australia since the mid-1970s. For males, both upper and lower tail inequality increased, especially since around 1990. For females, upper tail inequality increased since 1990, but lower tail inequality remained more stable. Therefore, although we have documented signs of job polarisation in Australia in the 1980s and 1990s, there is no strong evidence of earnings polarisation over any of these periods. This is similar to what has been observed in many other developing countries for which this analysis has been undertaken (the UK, Germany and Canada), but differs from the US experience of the 1990s, where earnings grew faster at the bottom of the earnings distribution than at the middle (the 50/10 earnings gap narrowed, implying earnings polarisation).

In the next two sections we turn to investigating whether the changes in the occupational composition of employment described in Sections III and IV have contributed to earnings inequality changes in Australia. The sources of any changes in earnings inequality can in general be decomposed into four components:

- (a) Effect of changes in dispersion in average earnings between occupations;

¹³ The ABS no longer provides earnings distribution information for full-time employees separately by gender in their standard release (6306.0). Figures since 2006 were constructed using ABS CURF micro-data.

- (b) Effect of changes in the extent of dispersion in earnings within occupations;
- (c) Effect of changes in the shares of employment between occupations with different average earnings (for example, shifts towards occupations with average earnings at the top and bottom of that ordering); and
- (d) Effect of changes in the shares of employment between occupations with different degrees of within-occupation earnings dispersion (for example, shifts towards occupations with higher intra-occupation earnings dispersion).

In section VII we investigate the effect of the third component (c) on earnings dispersion, and in section VIII the effect of the first component (a). Data limitations - specifically the absence of individual-level data at detailed occupation levels over the period - mean we are not able to examine the second component (b) or the fourth component (d). While it would be valuable to be able to undertake an analysis of the complete set of potential influences, we find that the first and third components are able to explain a large fraction of the change in earnings inequality in Australia.

VII. Employment Changes by Occupation and Earnings Inequality

In this section we examine the effect of changes in employment shares by occupation on earnings inequality. We do this by constructing pseudo-distributions of earnings that show the effect of changes to the occupation composition of employment, while holding constant workers' earnings. To construct these measures we use total employment by occupation and gender at the 4-digit level taken from each Census from 1966 to 2011. For each Census year and gender, we construct a pseudo-distribution of earnings using that year's Census employment weights and earnings by occupation for 1986 from the ABS Employer Survey. In essence, we are attributing the average earnings in the occupation to each worker employed in that occupation. We then construct the 90/50 and 50/10 percentile log earnings gaps from our pseudo-distributions. These gaps are presented in Figure 14. Changes over time in these constructed log earnings gaps provide a measure of the contribution of changes in employment shares by occupation to increases in earnings inequality.

Note that the log earnings gaps for our pseudo-distributions in Figure 14 are smaller than the actual gaps presented in Figures 12 and 13. This is expected given that our pseudo-distributions do not include any dispersion of individual earnings within occupations. In Figure 14 we can see that for males the 90/50 gap increases markedly from 1966 to 1991

(approximately 0.08 log points), but has been relatively stable since. The 50/10 gap for males also increased over the same period (0.04 log points), but again has been relatively stable since. For females, the 90/50 gap increased (a little haphazardly) over the whole period (0.09 log points). The 50/10 gap increased from 1986 to 2001, but declined slightly both pre and post this period. Given the lumpiness of data on employment by occupation, it was not unexpected to see uneven movements in these pseudo-wage gaps.

Overall, changes in occupational composition for males have potentially contributed 10% to the increase in upper tail earnings inequality over the period from 1976 to 2011 (0.02 of the approximate 0.2 log point increase observed in Figure 14). It has also potentially contributed one quarter to the increase in lower tail inequality over the same period (0.04 of the approximate 0.15 log point increase). For females, the occupation changes observed above have also added to upper tail earnings dispersion over the period, contributing nearly a quarter (0.05 of approximately 0.2 log points) to the 90/50 gap. Note also that the general upskilling observed among females has potentially contributed to the higher overall wage growth observed for females relative to males over this period.

VIII. Earnings Changes by Occupation

Estimates of the effect of employment changes by occupation on earnings inequality derived in the previous section are based on the assumption that earnings by occupation do not also change at the same time. The full effect of increased demand for high skill workers on earnings inequality may be higher than that depicted in Figure 14 if earnings of high skill workers increased in response to demand pressures, limiting the increase in employment. We consider the effect of changes in earnings by occupation on overall earnings inequality in this section.

It is not clear *a priori* whether employment and earnings by occupation should co-vary positively or not. The sign of co-variation will depend on a number of factors, including the following:

1. If changes in the demand for workers induced the observed changes in employment, it suggests positive co-variation. For example, increased demand for highly skilled workers may bid up wages in high skill occupations.
2. If supply changes induced the employment changes, then a negative relationship is expected. For example, if those workers who are pushed out of declining middle skill

occupations do not have the skills required to enter the growing high skill occupations, they may increase the pool of workers looking for low skill work, bidding down wages in low skill occupations.

3. Worker composition changes may imply a negative relationship. For example, in declining middle skill production occupations, the most educated and experienced workers may keep their jobs, and such workers are paid more than those who are forced out, raising average wage levels.
4. Changes in skill requirements within occupations where computers replace specific tasks may induce up-skilling and higher wages. For example, administrative work has changed considerably, with routine tasks replaced with more non-routine and interactive tasks. The education levels of workers who obtain administrative positions may now be higher, and wages also higher to attract skilled workers, yet there are fewer administrative jobs.

As noted above, the ABS changes the way it categorises workers into occupations each decade or so. In Sections II and III, we constructed a consistent occupation structure for obtaining long-run estimates of the share of employment by occupation by using 4-digit link files. This involved applying the proportions of each occupation in one categorisation that were allocated to each occupation in the other categorisation. This re-weighting seems defensible as long as those proportions remain reasonably stable over time. With earnings, however, re-weighting occupations in this manner is less defensible. For example, suppose one occupation in the old categorisation is split into two occupations under the new categorisation. Re-weighting would essentially be assuming that both the new occupations earn on average the same as the one old occupation. This may not be correct.

Therefore, in this section, we restrict our attention to investigating separate periods within which the occupation categorisation remains constant. We investigate earnings by occupation movements over two periods:

1. 1986 to 1995, using earnings by occupation at the 4-digit level using ASCO1 definitions.¹⁴

¹⁴ These average weekly earnings by 4-digit occupation for full-time adult employees were provided directly by the ABS in unpublished tables.

2. 1996 to 2006, using earnings by non-managerial occupation at the 3-digit level using ASCO2 definitions.¹⁵

We begin in the left panel of Table 5 by providing measures of changes over time in average earnings by occupation separately by the initial earnings quintile of each occupation. To avoid a mechanical negative relationship between earnings changes and initial earnings, we constructed earnings changes over the period from 1986 to 1995, and used earnings in 1987 to identify initial earnings quintiles. Similarly, we use earnings in 1998 to identify the initial earnings quintiles when measuring earnings changes over the 1996 to 2006 period.¹⁶ Note also that when constructing average earnings for an occupation, we use a weighted average of average earnings of males and females in that occupation. These weights by gender are held constant over each period (at initial year levels) in order to avoid gender composition changes within occupations affecting these earnings by occupation measures.

The figures in the left panel of Table 5 reveal that earnings growth was lower in initially low paying occupations over both periods (the lowest two quintiles). This should add to lower tail inequality. Earnings growth is quite similar over the top three initial earnings quintiles; thus earnings differences by occupation should not impact upper tail inequality to any real extent.

The figures in the right panel of Table 5 are measures of changes in occupational earnings where occupations are split into five groups based on employment growth over the relevant time period. These figures do not reveal any consistent pattern between earnings growth and employment growth over either period. Simple linear regressions of log earnings growth by occupation on log employment growth by occupation separately for the two time periods yielded coefficients on employment growth that were both economically small and statistically insignificant. These coefficients are provided in the top row of Table 6.

We then split occupations into two groups based on initial earnings level: the top 40% and the bottom 40%. We ran regressions of log earnings growth on log employment growth by occupation separately among these two groups. The results are in the bottom two rows of Table 6. We find a positive relationship for initially high paying occupations, which is

¹⁵ These average weekly earnings by 3-digit occupation for full-time non-managerial adult employees were collected from annual publications of the ABS's *Employee Earnings and Hours*, cat. no. 6306.0.

¹⁶ The ABS has published these earnings numbers by detailed occupation on a bi-annual basis since 1996, while they were published on an annual basis prior to 1996.

significant at the 10% level over the 1986-1995 period. We find a negative or zero relationship for initially low paying occupations over both periods. It suggests that high paying (high skill) jobs have grown due to demand effects (skill-biased technological change), whereas low paying (low skill) jobs may have grown due to supply effects. DLS (2009) and Autor (2014) arrive at similar conclusions for Germany and the US, respectively.

We now provide measures of the effect of these changes in average earnings by occupation on overall earnings inequality. We start by measuring the effect of earnings changes in isolation, holding the composition of employment by occupation constant at initial levels in each period. We then provide a measure of the combined effect of earnings and employment changes by occupation on the earnings distribution. These measures are provided in Table 7. The measures are again based on the construction of pseudo-distributions using earnings by occupation from the ABS Employer Surveys and employment weights from the Censuses.

In the left panel of Table 7, we provide estimates of the changes in the log 90/50 and log 50/10 earnings differentials for men and women using our pseudo-distributions and allowing earnings by occupation to change according to observed movements in the earnings data, but keeping the employment weights constant at initial levels. We find that between 1986 and 1995 changes in earnings by occupation have only small negative effects on upper tail inequality, and moderate positive effects on lower tail inequality. This is consistent with the earnings changes by initial earnings level observed in the left panel of Table 5. Over the 1996-2006 period, earnings changes added to both upper and lower tail inequality for males, with a larger effect on the upper tail. For females, earnings changes only affected lower tail inequality positively over this second period.

In the right panel of Table 7, we provide our measures of the combined effect of both employment and earnings changes by occupation on overall earnings inequality. We now observe positive contributions for both genders to upper and lower tail inequality over both time periods. The combined effect of the change in average earnings by occupation and employment shares by occupation over the 1986 to 2006 period appear to explain essentially all of the overall observed increase in lower tail inequality for females. These same changes also explain a small proportion of upper tail inequality growth among females. For males, these combined changes potentially explain approximately half of the overall increase in inequality in both tails. Thus, changes in average earnings and the composition of

employment by occupation have been extremely important in understanding earnings inequality growth in this country.

IX. Conclusions

There has been considerable interest internationally in how technological change has affected labour markets. A major question is whether technological change has been purely skill-biased, raising demand for skilled versus unskilled workers, or has changed the relative demand for workers according to their capacity to perform routine tasks, causing job polarisation.

Many countries have exhibited occupation changes consistent with job polarisation (UK, US, Canada, Germany and many other European countries). We add to this growing evidence to show that Australia has also experienced job polarisation. This occurred primarily in the 1980s and 1990s, whereas in the 1970s and 2000s changes in the composition of employment were more consistent with general up-skilling. We also show that job polarisation in Australia has been mainly a male phenomenon. Consistent with the routinisation hypothesis, we establish that there were large declines in the employment shares of occupations that were initially high in routine task intensity.


We further find that changes in the occupation structure of employment can explain large proportions of the overall increases in earnings inequality that have occurred in Australia. Changes in average earnings by occupation have also contributed to earnings inequality changes. Pressure on the earnings distribution appears to be a common consequence of the technological change that has happened here and internationally over the past several decades.

References

- Acemoglu, Daron and David Autor, 2011. Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics* Vol. 4B, O. Ashenfelter and D. Card (eds), 1043-1171: Elsevier (North-Holland).
- Aungles, Phil, Lorraine Dearden, Tom Karmel and Chris Ryan, 1993. Through a rear-view mirror darkly: Occupational change 1971-86. *Australian Bulletin of Labour* 19(2): 97-113.
- Australian Bureau of Statistics, 2006. *ANZSCO - Australian and New Zealand Standard Classification of Occupations, First Edition*. ABS Catalogue Number 1220.0, Canberra.
- Autor, David, 2013. The 'task approach' to labor markets. *Journal of Labour Market Research* 46(3): 185-99.
- Autor, David, 2014. Polanyi's paradox and the shape of employment growth. National Bureau of Economic Research, Working Paper 20485.
- Autor, David and David Dorn, 2013. The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review* 103(5): 1553-1597.
- Autor, David, Lawrence Katz and Melissa Kearney, 2006. The polarization of the U.S. labor market. *American Economic Review, Papers and Proceedings* 96(2): 189-194.
- Autor, David, Lawrence Katz and Melissa Kearney, 2008. Trends in US wage inequality: revising the revisionists. *The Review of Economics and Statistics* 90(2): 300-323.
- Autor, David, Frank Levy and Richard Murnane, 2003. The skill content of recent technological change: an empirical investigation. *Quarterly Journal of Economics* 118(4): 1279-1333.
- Borland, Jeff, 1991. Incomes policies in Australia. *Australian Economic Review* 24(2): 45-50.
- Cully, Mark, 1999. A more or less skilled workforce? Changes in the occupational composition of employment, 1993 to 1999. *Australian Bulletin of Labour* 25(2): 98-104.
- Dunlop, Yvonne and Peter Sheehan, 1998. Technology, skills and the changing nature of work. *Working for the Future: Technology and Employment in the Global Knowledge Economy*, Peter Sheehan and Greg Tegar (eds), Victoria University Press: Melbourne.
- Dustmann, Christian, Johannes Ludsteck and Uta Schonberg, 2009. Revisiting the German wage structure. *Quarterly Journal of Economics* 124(2): 843-881.
- Esposito, Alexis, 2011. Upskilling and polarisation in the Australian labour market: A simple analysis. *Australian Bulletin of Labour* 37(2): 191-216.

- Goos, Maarten and Alan Manning, 2007. Lousy and lovely jobs: the rising polarisation of work in Britain. *Review of Economics and Statistics* 89(1): 118–133.
- Goos, Maarten, Alan Manning and Anna Salomons, 2009. Job polarisation in Europe. *American Economic Review, Papers and Proceedings* 99(2): 58-63.
- Goos, Maarten, Alan Manning and Anna Salomons, 2014. Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review* 104(8): 2509-26.
- Green, David and Benjamin Sand, 2014. Has the Canadian labor market polarized? Canadian Labour Market and Skills Researcher Network (CLSRN) working paper no. 133.
- Gregory, Robert, 1993. Aspects of Australian and US living standards: The disappointing decades 1970-1990. *Economic Record* 69(1):61-76.
- Johnson, George, 1997. Changes in earnings inequality: The role of demand shifts. *Journal of Economic Perspectives* 11(2): 41-54.
- Keating, Michael, 2003. The labour market and inequality. *Australian Economic Review* 36(4): 374-90.
- Lefter, Alexandru and Benjamin Sand, 2011. Job polarization in the U.S.: A reassessment of the evidence from the 1980s and 1990s. University of St. Gallen, Department of Economics Discussion Paper no. 2011-03.
- Pappas, Nick, 2001. Earnings inequality and skill. *Work Rich, Work Poor*, Jeff Borland, Robert Gregory and Peter Sheehan (eds), Victoria University Press: Melbourne (196-218).
- Spitz-Oener, Alexandra, 2006. Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics* 24(2): 235-270.
- Wilkins, Roger and Mark Wooden, 2014. Two decades of change: The Australian labour market 1993-2013. *Australian Economic Review* 47(4): 417-31.
- Wooden, Mark, 2000. The changing skill composition of labour demand. *Australian Bulletin of Labour* 26(3): 191-98.

Figure 1: Per cent Change in Employment by Occupation, 1971-2011, Australia



Sources: Australian Censuses, 1971 to 2011, occupations defined at 4-digit level prior to grouping, all employed individuals, excluding agricultural and military occupations.

Table 1: Occupation Groups and Average Skill Levels, Alternative Measures

Occupation Group	Per cent share of total employment			1986 Earnings		ABS Skill Category (av.)	
	1966	2011	change	Level	Rank	Level	Rank
Managers	9.0	11.8	2.8	651	1	1.37	2
Professionals	11.3	22.4	11.1	561	2	1.01	1
Technicians	2.3	2.5	0.2	481	4	2.08	3
Sales	7.5	9.8	2.3	380	9	4.65	10
Office/administration	16.3	15.1	-1.2	380	8	3.62	7
Production	17.9	10.1	-7.8	394	6	3.00	4
Operators	24.8	12.4	-12.4	421	5	4.37	9
Protective	1.1	1.3	0.2	485	3	3.17	5
Food/cleaning	7.0	7.5	0.5	342	10	4.25	8
Personal care	2.8	7.1	4.3	385	7	3.50	6


Sources: Employment – Australian Censuses, 1966 to 2011; Earnings – ABS Employer Survey; ABS Skill Categories – ANZSCO definitions.

Table 2: Changes in Employment Shares by Occupational Skill Level, Australia and Europe

	Lowest pay occupations	Middle pay occupations	Highest pay occupations
Australia 1966-2011	+2.2	-19.2	+17.0
Australia 1991-2011	+1.5	-8.5	+7.0
Europe average 1993-2010	+2.7	-9.9	+7.2


Sources: Australia: Authors' calculations using occupation classification from GSM (2014, Table 1); and Europe: GSM (2014, Table 2).

Figure 2: Employment Changes by Skill Quintile, 1986 Earnings



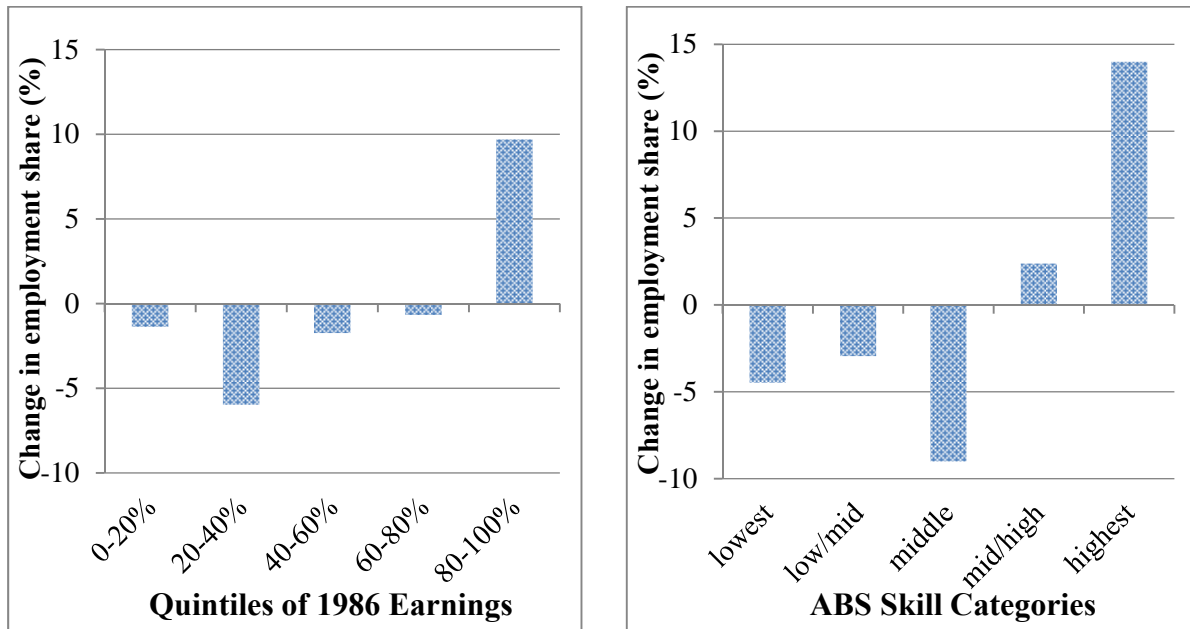
Sources: Employment– Australian Censuses, 1971 to 2011; Earnings – ABS Employer Survey. Occupations grouped into quintiles using employment weights for each occupation.

Figure 3: Employment Changes by ABS Occupation Skill Categories



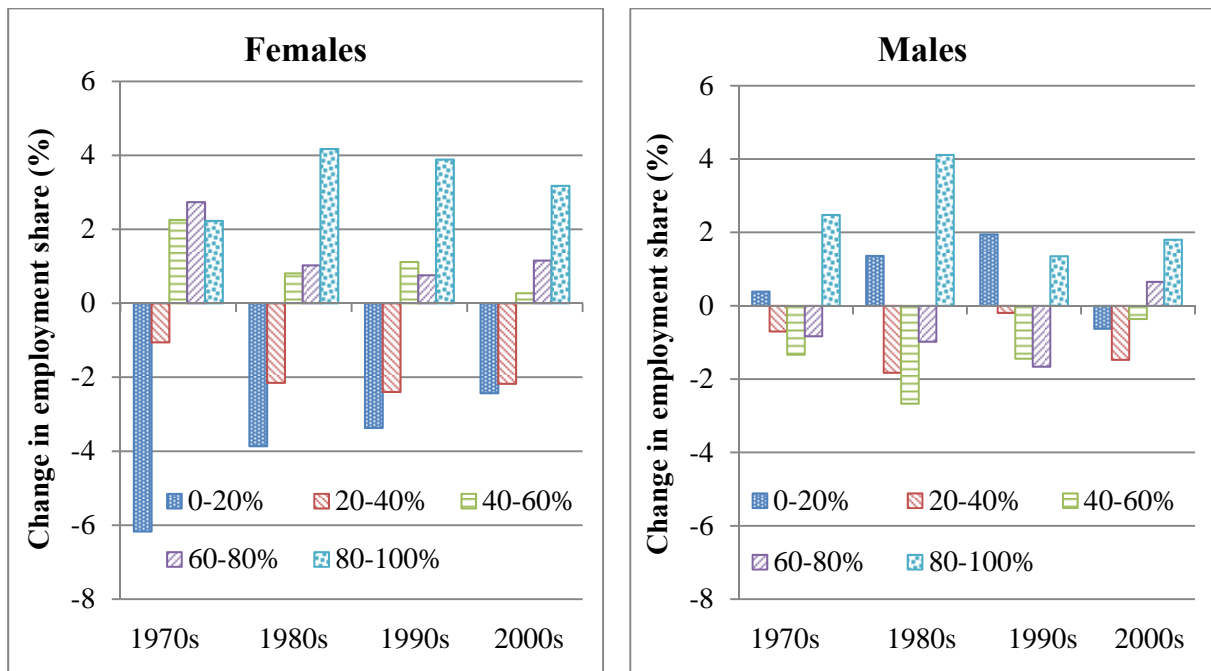
Sources: Employment– Australian Censuses, 1971 to 2011; Skill Categories – ANZSCO definitions.

Figure 4: Employment Changes by Skill Categories over Full Period, 1966-2011



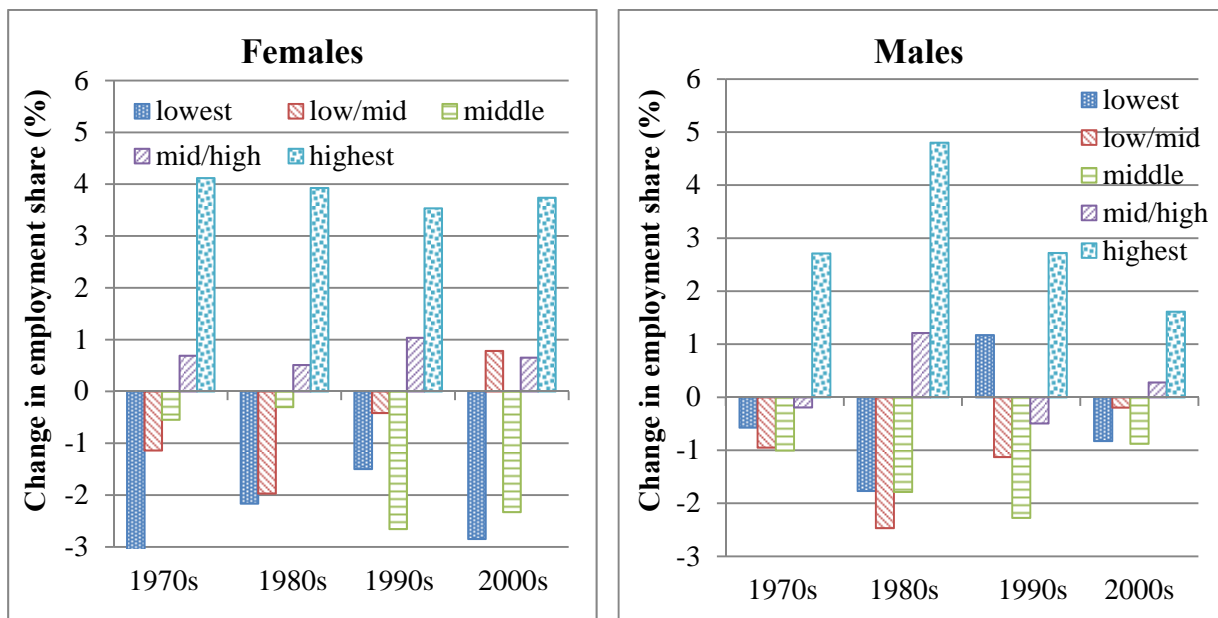
Sources: See Figures 2 and 3.

Figure 5: Employment Changes by Skill Quintile (1986 Earnings) and Gender



Sources: Employment – Australian Censuses, 1971 to 2011; Earnings – ABS Employer Survey. Occupations grouped into quintiles using employment weights for each occupation.

Figure 6: Employment Changes by ABS Skill Categories and Gender




Sources: Employment – Australian Censuses, 1971 to 2011; Skill Categories – ANZSCO definitions.

Table 3: Categories of DOT Measures and Interpretation

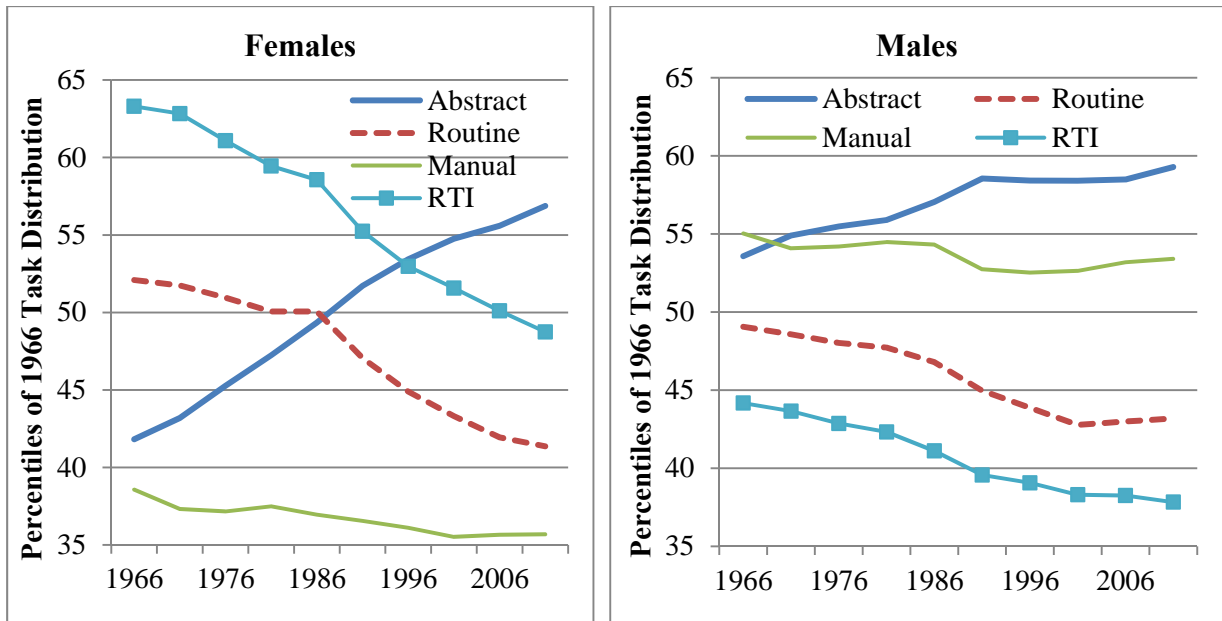
DOT category	Interpretation of revealed demand for task completion
General Educational Development Mathematical (GED-MATH)	Non-routine cognitive tasks – analytical reasoning
Adaptability to accepting responsibility for the direction, control and planning of activities (DCP)	Non-routine cognitive tasks – interactive, communication and managerial
Finger dexterity (FINGDEX)	Routine manual tasks
Adaptability to situations requiring the precise attainment of set limits, tolerances or standards (STS)	Routine cognitive tasks
Eye-hand-foot coordination (EYEHAND)	Non-routine manual tasks

Figure 7: Average of Percentiles of Task Measures over Time




Notes: Task measures constructed by linking ASCO1 occupations to the US Dictionary of Occupational Titles measures for 1977, and using the measures of Autor and Dorn (2013), where RTI refers to their Routine Task Intensity Index. The percentiles were constructed using 1966 employment weights over occupations and genders.

Figure 8: Average of Percentiles of Task Measures over Time by Gender



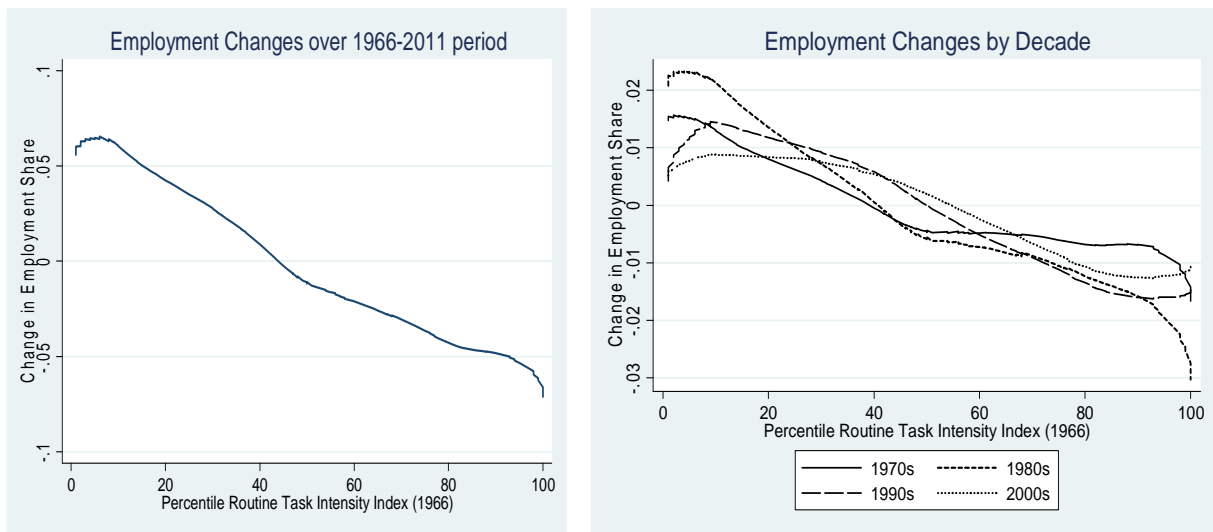
Notes: As per Figure 7.

Figure 9: Average of Task Measures by Occupational Skill (Initial Earnings)



Notes: Task measures constructed by linking ASCO1 occupations to the US Dictionary of Occupational Titles measures for 1977, and using the task and RTI measures of Autor and Dorn (2013). Average percentile of task measure smoothed using locally weighted regressions and a bandwidth of 0.6.

Figure 10: Employment Share Changes by Occupation Routine Task Intensity



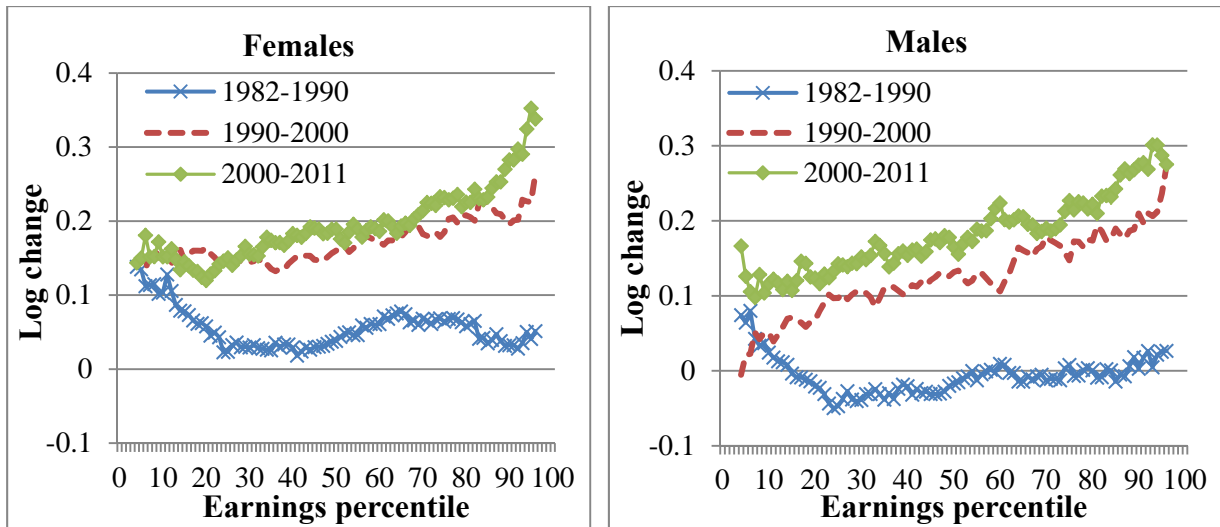
Notes: Routine Task Intensity (RTI) Index constructed by linking ASCO1 occupations to the US Dictionary of Occupational Titles measures for 1977, and using the RTI measure of Autor and Dorn (2013). Changes smoothed using locally weighted regressions and a bandwidth of 0.8.

Table 4: Regression Estimates—Employment and the RTI Index in Australia

Dependent variable (logged)	(1) Employment by occupation by gender	(2) Employment by occupation by gender	(3) Employment by occupation	(4) Employment by occupation
Sample period	1966 to 2011	1966 to 2011	1966 to 2011	1991 to 2011
RTI × Linear time trend	-1.101*** (0.135)	-0.810*** (0.063)	-1.107*** (0.151)	-0.912*** (0.170)
RTI × Linear time trend × Female		-0.357 (0.264)		
R-squared	0.0617	0.0623	0.0499	0.0542
Observations	5,396	5,396	2,710	1,355

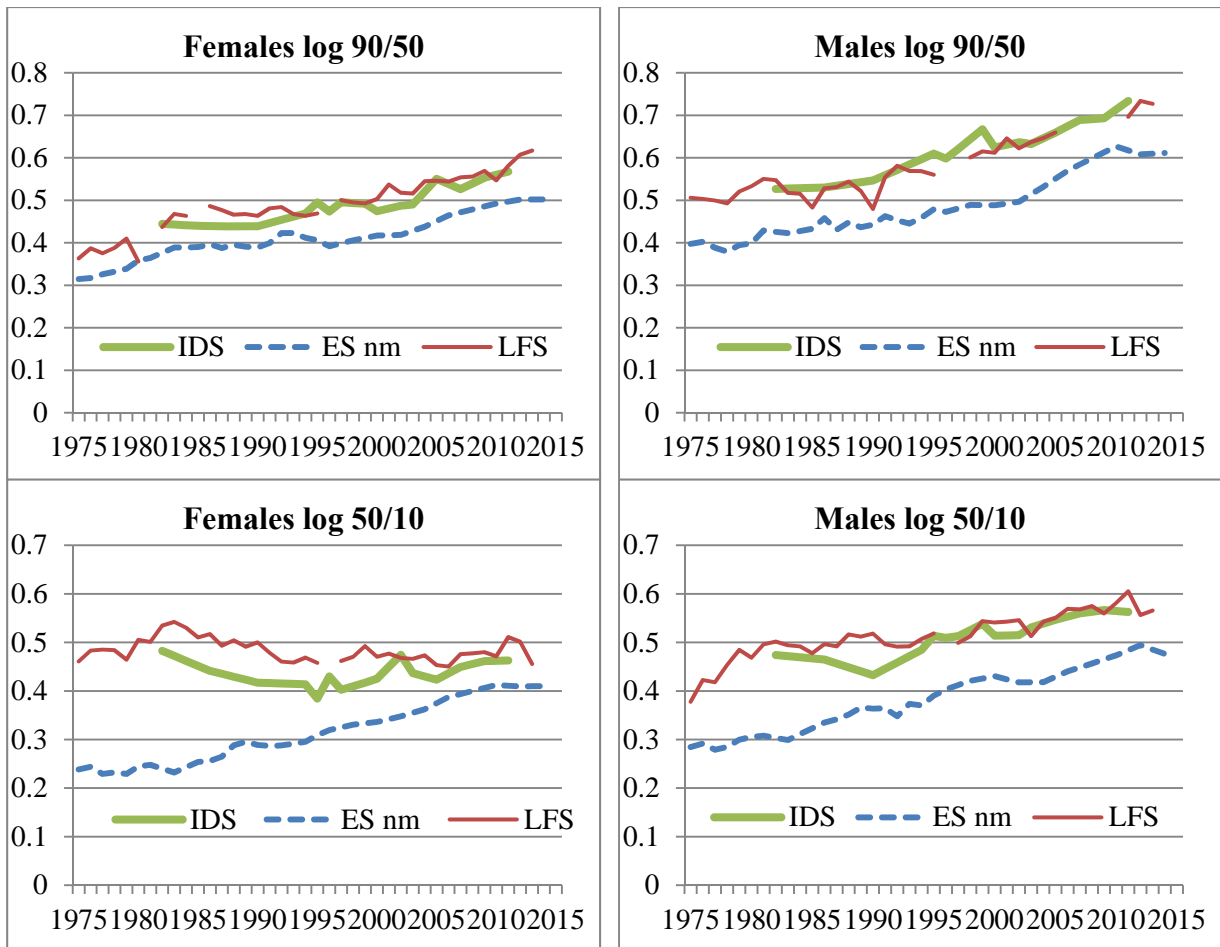
Notes: Regression estimates (multiplied by 100) using Australian Census employment over time and the RTI Index constructed using 1966 employment weights. Occupation by gender and time fixed effects are also included in columns (1) and (2), while occupation and time fixed effects are included in columns (3) and (4). Standard errors (provided in parentheses) were constructed allowing for clustering at the occupation by gender level in columns (1) and (2), and at the occupation level in columns (3) and (4). *** denotes significant at the 1% level.

Figure 11: Real weekly wage growth by percentile, Full-Time Employees aged 15-64



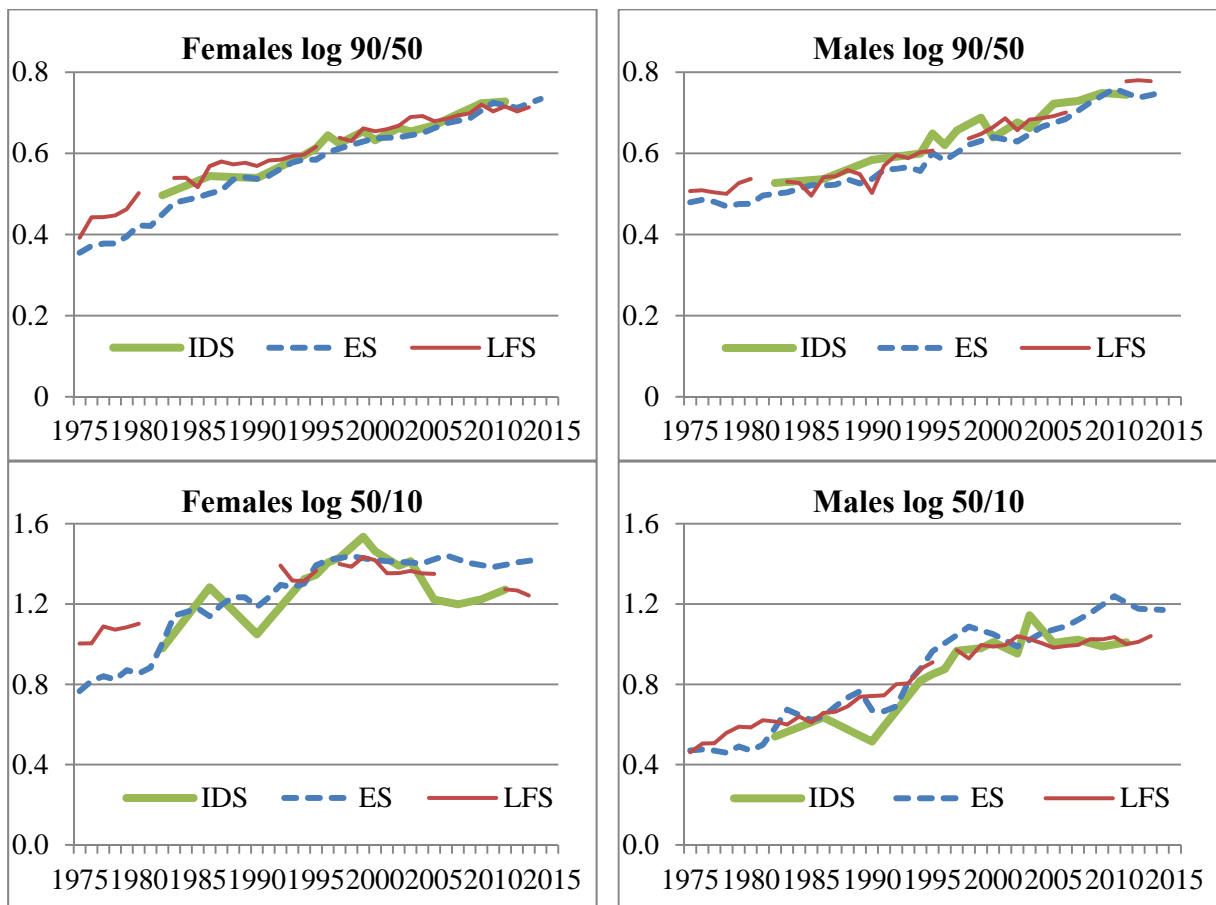
Sources: IDS – Income Distribution Survey, CURF Micro-data. Weekly wages were first deflated using the Australian Consumer Price Index constructed by the ABS.

Figure 12: Earnings Inequality over Time, Full-Time Employees aged 15-64




Sources: IDS – Income Distribution Survey, CURF Micro-data; ES – Employer Survey, ABS cat. No. 6306.0; LFS – Labour Force Survey, ABS cat.No. 6310.0.

Figure 13: Earnings Inequality over Time, All Employees



Sources: IDS – Income Distribution Survey, CURF Micro-data; ES – Employer Survey, ABS cat. No. 6306.0; LFS – Labour Force Survey, ABS cat.No. 6310.0.

Figure 14: Effect of Employment Changes by Occupation on Earnings Inequality



Sources: Employment weights – Australian Censuses, 1971 to 2011; Earnings – ABS Employer Survey for 1986. Log earnings gaps calculated based on pseudo-distributions of earnings where each worker is attributed the average wage in 1986 for their occupation.

Table 5: Real Earnings Growth by Initial Earnings Level and Employment Growth

“Initial” Earnings Quintile	1986 to 1995	1996 to 2006	Employment Growth Quintile	1986 to 1995	1996 to 2006
	All FT 4-digit	FT non-m 3-digit		All FT 4-digit	FT non-m 3-digit
1 – lowest	-1.43	8.06	1 – lowest	1.43	13.16
2 – low-mid	0.65	13.25	2 – low-mid	1.80	18.27
3 – middle	2.51	15.86	3 – middle	1.48	13.33
4 – mid-high	2.31	14.04	4 – mid-high	-0.33	8.74
5 – highest	2.65	18.45	5 – highest	2.23	18.04

Notes: All measures are earnings growth constructed at the occupation level using initial employment weights from the Census. In the left panel, the “initial” earnings quintiles were constructed using 1987 earnings and 1986 employment weights for the 1986 to 1995 earnings growth figures, and using 1998 earnings and 1996 employment weights for the 1996 to 2006 earnings growth figures. In the right panel, employment growth quintiles were constructed using Census data at the occupation level, from 1986 to 1996 and from 1996 to 2006 respectively. “non-m” refers to non-managerial workers only.

Table 6: Log Earnings Growth Regressions on Log Employment Growth by Occupation

Occupations	1986-1995, All FT, 4-digit		1996- 2006, FT non-m, 3-digit	
	Coefficient (standard error)	Observations	Coefficient (standard error)	Observations
All occupations	0.0064 (0.014)	269	0.0259 (0.032)	75
Top 40% by initial earnings	0.0409* (0.022)	152	0.0517 (0.056)	46
Bottom 40% by initial earnings	-0.0242 (0.019)	76	0.0026 (0.072)	17

Notes: * denotes significant at the 10% level. Data sources as per Table 5.

Table 7: **Changes in Earnings Inequality Attributable to Earnings and Employment Changes by Occupation**

		Earnings Changes Only		Earnings and Employment Changes	
		1986-1995	1996-2006	1986-1995	1996-2006
Change in log 90-50	Males	-0.013	0.092	0.038	0.077
	Females	-0.009	-0.044	0.017	0.003
Change in log 50-10	Males	0.039	0.048	0.039	0.063
	Females	0.028	0.146	0.035	0.146

Notes: Authors' calculations.

Appendix - Decomposition Details

The decomposition was constructed as follows. Female employment in occupation i in period t (F_i^t) can be calculated as total employment in that occupation and time period (T_i^t) multiplied by the proportion of workers in the occupation and time period that are female (pf_i^t). Consider the change in female employment within some occupation i over a particular period, from period 0 to period 1.

$$F_i^1 - F_i^0 = pf_i^1 \times T_i^1 - pf_i^0 \times T_i^0 = (pf_i^1 - pf_i^0)T_i^0 + (T_i^1 - T_i^0) pf_i^1$$

The first component on the right hand side of this equation $(pf_i^1 - pf_i^0)T_i^0$ is what we are interpreting as the gender composition effect on the change in female employment over time. Given a constant level of employment in occupation i (no change in the overall structure of employment), the number of females working in the occupation may still change due to changes in the gender mix. The second component $(T_i^1 - T_i^0) pf_i^1$ is interpreted as the effect of overall changes in the structure of employment on female employment, holding fixed the gender mix.

Note that this decomposition is not unique. We can also construct it as follows:

$$F_i^1 - F_i^0 = pf_i^1 \times T_i^1 - pf_i^0 \times T_i^0 = (pf_i^1 - pf_i^0)T_i^1 + (T_i^1 - T_i^0) pf_i^0$$

In figure A1, we plot the simple average of the estimates constructed using both of these two potential decompositions.

Table A1: Largest Changes in Occupation Employment Share by Gender

	Occupation	Group	Earnings	Δ Share %
Female losers				
1	Typists and typist-clerks	Clerical	1	-7.9
2	Office secretaries and stenographers	Clerical	2	-4.8
3	Cleaners	Food/Clean	1	-2.4
4	Textile sewing machinists	Operators	1	-2.2
5	Other trade assistants and factory hands	Operators	2	-2.1
6	Electronic data processing machine operators	Clerical	3	-1.7
7	Hand packers	Operators	1	-1.7
8	Telephonists	Clerical	2	-1.5
9	Assemblers	Operators	1	-1.3
10	Enrolled nurses	Personal	2	-1.2
Female gainers				
1	Child care, refuge and related workers	Personal	2	3.4
2	Accounting clerks	Clerical	3	3.0
3	Other para-professionals	Technicians	3	1.8
4	Accountants	Professionals	5	1.4
5	Receptionists and information clerks	Clerical	1	1.4
6	Welfare para-professionals	Personal	4	1.0
7	Teachers' aides	Personal	1	0.9
8	Personnel specialists	Professionals	5	0.8
9	Sales and marketing managers	Managers	5	0.7
10	Other business professionals	Professionals	5	0.7
Male losers				
1	Accounting clerks	Clerical	3	-2.0
2	Metal fitters and machinists	Production	4	-1.7
3	Other trade assistants and factory hands	Operators	2	-1.7
4	Trades assistants	Operators	2	-0.9
5	Stock and purchasing clerks	Clerical	3	-0.8
6	Truck drivers	Operators	3	-0.8
7	Communications equipment tradespersons	Production	3	-0.7
8	Shop managers	Managers	3	-0.6
9	Meat tradespersons	Food/Clean	2	-0.6
10	Electrical/electronic engineering associates	Technicians	5	-0.5
Male gainers				
1	Computing professionals	Professionals	5	2.1
2	Sales assistants	Sales	1	1.4
3	Cooks	Food/Clean	2	1.0
4	Other specialist managers	Managers	5	0.9
5	Kitchenhands	Food/Clean	1	0.8
6	Data processing managers	Managers	5	0.8
7	Electrical mechanics	Production	4	0.8
8	Other para-professionals	Technicians	3	0.7
9	General managers	Managers	5	0.7
10	Other labourers and related workers	Operators	3	0.6

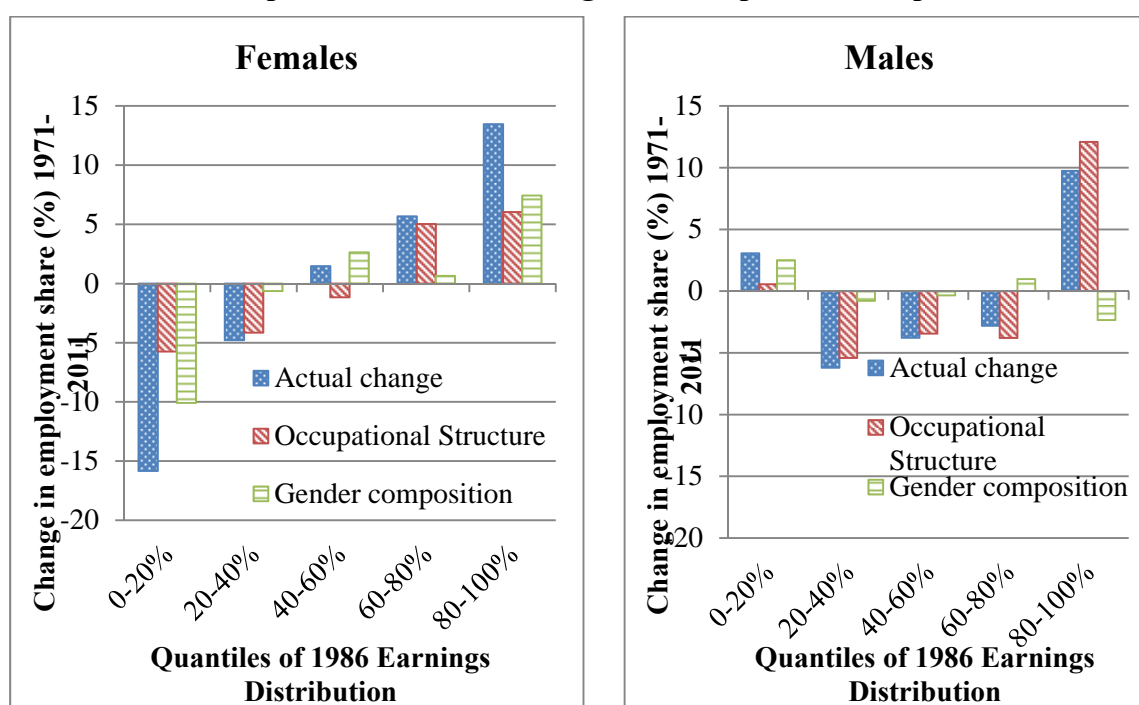
Sources: Employment – Australian Censuses, 1966 to 2011; Earnings – ABS Employer Survey, 1986.
Occupation earnings are quintiles of 1986 distribution. Occupations use ASCO1 definitions.


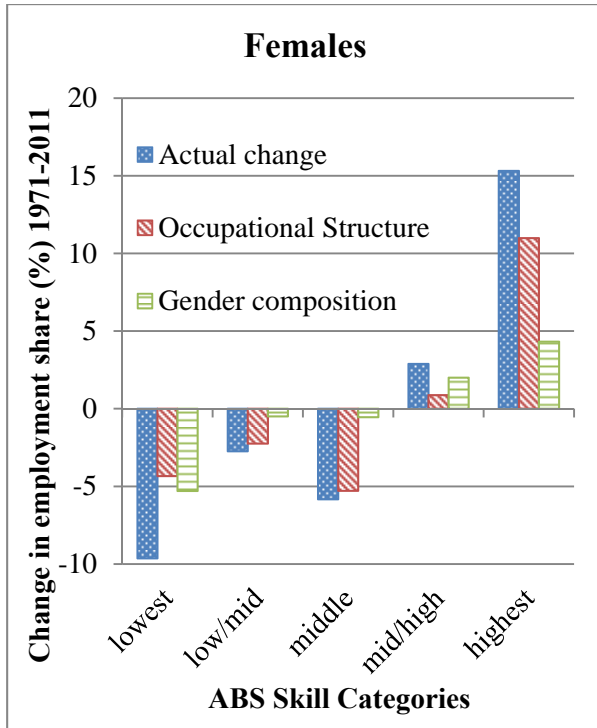
**Table A2: Comparison of results with Wilkins and Wooden (2014):
Changes in occupational employment shares**

Skill quintile (ABS)	Our study 1991-2011 (Census)	Wilkins and Wooden (2014) 1993-2013 (Labour Force Survey)
Top	+5.7	+5.5
Second	+0.3	+0.7
Third	-4.2	-3.8
Fourth	0.0	0.3
Bottom	-1.8	-3.4

Sources: Authors' calculations; and Wilkins and Wooden (2014, Table 4).

Figure A1: Decomposing actual changes in gender employment shares into occupational structure and gender composition components





Sources: Authors' calculations.