Decomposing changes in South African wage inequality 1994-2011: The role of skills and institutions

Anthea du Toit
Martin Wittenberg
School of Economics, SALDRU and DataFirst
University of Cape Town

August 24, 2015

Abstract

Draft only - please do not quote

1 Introduction

The international literature on earnings inequality has expanded rapidly over the last 20 years, driven by a need to explain the large increases in inequality observed in especially the US and UK since the late 1970s (Machin & Van Reenen, 2007:1). While a fairly standard set of explanations have emerged, mostly linked to changes in skills on the one hand or institutional factors on the other, there is by no means a consensus view on the relative importance of these phenomena.

In South Africa, studies of earnings inequality have been limited to the post-Apartheid era, since detailed household survey data has only been collected since the early 1990s. Indications are that earnings inequality increased during the latter half of the 1990s and has steadily declined since 2000 (Burger & Yu, 2007:13; Leite et al, 2006). During this period, the South African labour market was subject to a range of influences, including the enactment of several important pieces of labour legislation, growing union power, increased openness to trade, rising educational attainment combined with signs of worsening quality of education and the rise of the HIV/AIDS pandemic. In terms of global trends, South Africa has also been exposed to the microcomputer revolution, which has had a dramatic influence on the production technology and labour productivity of companies worldwide.

What are the linkages between these changes and changes in earnings inequality? Theoretical explanations of earnings inequality can be grouped into two categories, namely market-related factors and institutional factors (Machin

& Van Reenen, 2007:3). Market-related factors refer to the supply and demand of different types of labour, with a particular focus on skilled versus unskilled labour. One of the most important demand-side theories that has been proposed is skills-biased technological change (SBTC), which predicts that increasing microcomputer usage will lead to an increase in the demand for skilled labour and rising earnings inequality. Another strand of the literature focuses on the impact of globalisation on labour demand. According to this explanation, competition from developing countries, where unskilled labour is relatively cheaper, will drive down the relative price of unskilled labour in developed countries, which will result in growing earnings inequality.

On the supply side, changes in the level of skills can mechanically lead to changes in earnings inequality. For instance, rising levels of educational attainment will result in higher inequality, since the variation of earnings increases with age and education. Alternatively, a reduction in the quality of education may result in a greater premium being placed on qualifications obtained from better quality institutions, which can also lead to greater inequality (Burger & Woolard, 2005:18). HIV/AIDS is also likely to affect the supply of labour. In particular, the impact of HIV/AIDS on earnings inequality will depend on which segment of the labour force is hardest hit by the pandemic. If younger, less skilled individuals are more likely to drop out of the labour force as a result of HIV/AIDS-related illness, earnings inequality is likely to increase. On the other hand, if the ranks of older, more educated workers are more likely to be depleted by HIV/AIDS, this will reduce earnings inequality.

Institutional factors also have an important role to play in determining earnings inequality. The level of union power and the real minimum wage level are predicted to have an inverse relationship with earnings inequality: higher minimum wages or increased union power will result in lower earnings inequality. As a result, changes in South African labour legislation, which have had a positive impact on both of these variables, would have led to lower earnings inequality. The end of Apartheid also resulted in the removal of discriminatory labour practices and the introduction of legislation aimed at promoting the employment of previously disadvantaged individuals. The expectation is that this will lead to a further reduction in earnings inequality.

In that regard, it is important to be cognisant of the interaction between earnings inequality, unemployment and income inequality. In particular, institutional factors, such as minimum wage provisions, may actually keep earnings inequality at an artificially low level. Therefore, if earnings inequality cannot increase in response to market factors, employment levels may fall instead. As such, measures of income inequality, which in addition to differences in earnings amongst those who are employed also take into account differences between employed individuals and those who are unemployed, may show an increase in inequality at the same time that earnings inequality has fallen.

Overall then, changes in market-related and institutional factors in the South African labour market present different predictions with regard to changes in earnings inequality. Separating out the effects due to individual factors is a difficult (if not impossible) task, which is further complicated by the impact of selection into and out of employment. Keeping these difficulties in mind, one method that has been proposed by Lemieux (2002; 2006a) makes use of a simple human capital framework and focuses on skills-related explanations of changes in earnings inequality. Lemieuxs method provides a way to decompose changes in inequality into two subcomponents, namely changes due to prices and changes due to composition effects, which is similar to the Oaxaca-Blinder decomposition of changes in means. The method presents the advantage of being straightforward to implement and can be used to calculate a variety of inequality measures.

This study uses Lemieux's method to examine changes in South African earnings inequality, in order to distinguish between the relative importance of price effects (e.g. changes due to SBTC or globalisation), compositional effects (e.g. changes due to education or the impact of HIV/AIDS) and other effects (e.g. institutional factors or measurement error). This is the first study to apply Lemieuxs method to South African data. Furthermore, a time series is constructed from eleven cross-sectional datasets with observations for years from 1995 to 2006, which extends the number of data points beyond the usual two or three points in time used in many other South African studies of inequality.

While the results indicate that SBTC or globalisation may have been responsible for the increase in earnings inequality in the top half of the distribution, there is less support for these theories in the bottom half. There is also some evidence that rising compositional effects have led to increases in inequality, although the effects are limited. However, the trend in overall earnings inequality has been driven mainly by declining inequality in the bottom half of the distribution, and changes in the bottom half appear to have been relatively unaffected by skills.

As a result, it is concluded that institutional factors have probably been a more important driver of changes in overall earnings inequality. In that regard, higher union power, higher minimum wages and lower discrimination as a result of legislative changes are all predicted to lead to a reduction in earnings inequality, and are therefore consistent with the observed trends. Furthermore, institutional rigidities may have resulted in wages that cannot fully adjust in reaction to other factors such as SBTC and globalisation, so that the burden of adjustment has been left to employment. Indeed, indications are that unemployment has increased since the end of Apartheid, which supports this argument (Kingdon & Knight, 2007:814; Leite et al, 2006:18).

Section 2 provides a review of the literature on earnings inequality, both from an international and South African perspective. Section 3 goes on to explain Lemieuxs methodology for decomposing changes in earnings inequality. Section 4 discusses data-related issues and the procedures that were followed in creating the dataset for the analysis. Section 5 presents an overview of trends in earnings, education and experience, while Section 6 goes on to present the results from Lemieuxs decomposition methodology. The final section concludes with a discussion of the potential impact of institutional factors on earnings inequality.

2 Literature Review

Studies of the development of South African inequality since 1994 have painted a picture of rising overall inequality, which has been underpinned by declining between-group inequality and rising within-group inequality (Seekings et al, 2004:3; Leibbrandt et al, 2005a:3; Ardington et al, 2006:823). These findings have been fairly consistent, irrespective of the dataset or methodology used to measure changes in inequality. The importance of earnings as a driver of overall inequality has been highlighted in many of these studies, with estimates of the contribution of earnings to overall inequality ranging between 73% and 80%, depending on the dataset used[1]. Earnings as a driver of inequality is important for two reasons. Firstly, labour income makes up the largest part of overall income. Based on the 1993 PSLSD dataset, Leibbrandt et al (2001:76) find that wage income made up 69% of overall income in 1993, with the remainder coming from capital income (13%), self-employment (6%), state transfers (5%), agriculture (4%) and remittances (3%). Secondly, the pervasiveness of unemployment in South Africa means that a significant proportion of the population have to rely on the relatively small pool of non-labour earnings to survive[2]. Therefore, overall inequality is to a large extent driven by differences between the unemployed and employed. This also implies that the difference between estimates of overall inequality and earnings inequality is bound to be non-trivial.

2.1 Changes in South African earnings inequality since 1994

Although the literature has tended to focus more on overall inequality (probably in part due to the non-ignorability of unemployment), some studies have examined earnings inequality in greater detail. The literature on earnings inequality has also been stimulated by the increased availability of detailed labour force surveys since 1993, in the form of the October Household Surveys (OHS) for 1993 to 1999 and the Labour Force Surveys (LFS) since 2000[3].

Woolard & Woolard (2006) examine changes in earnings inequality as measured by earnings gaps for various groups using OHS and LFS data. Overall, they find that the earnings gap between low-skilled and more highly skilled workers widened between 1995 and 2003, which suggests that changes in earnings inequality were driven by changes in between-group inequality. Using data from the OHS and LFS, Burger & Yu (2007:13) find that earnings inequality increased between 1995 and 1998, but decreased thereafter. They find that overall, between-group earnings inequality worsened between 1995 and 2005, although the gender and racial earnings gap did narrow somewhat during the

[1] See, for instance, Leibbrandt et al (2001:76), Seekings et al (2004:9-10) and Leite et al (2006:14).

[2] A recent estimate places unemployment at 28% based on the narrow definition and 42% based on the broad definition (Kingdon & Knight, 2007:814)

[3] As such, most studies of earnings inequality have tended to use these datasets, and all of studies reviewed here have done so.

latter half of the period.

According to Leite et al (2006), earnings inequality increased sharply between 1995 and 1999, and declined somewhat between 2000 and 2004. They focus on unemployment as a cause of earnings inequality[4]. In particular, they rely on the wage curve relationship between earnings and unemployment, which postulates that unemployment and earnings are negatively correlated. In the case of South Africa, they argue that high levels of unemployment have lowered the earnings of unskilled workers by more than those of skilled workers, which has worsened inequality. They go on to claim that the positive correlation between unemployment and the Gini coefficient provides evidence in support of a wage curve relationship in South Africa. However, it is debateable whether correlation implies causation in this case. Indeed, it can just as easily be argued that a third exogenous factor such as SBTC or globalisation has led to a concomitant reduction in inequality and increase in unemployment.

2.2 International literature on earnings inequality

The international literature on earnings inequality has grown rapidly since the early 1990s, when empirical studies based on US data first noted a consistent trend of sharply rising earnings inequality over the preceding decade. Similar evidence of rising inequality has come to light in other countries, including the UK and New Zealand. The growth in overall inequality has been underpinned by rising educational differentials (especially the college-high school gap) as well as rising residual or within-group inequality (Machin & Van Reenen, 2007:1-2). The empirical observations in turn prompted theoretical explorations into the causes of these changes, the most prominent of which are briefly reviewed here.

2.2.1 Skills-Biased Technological Change

Skills-biased technological change (SBTC) has probably enjoyed the most attention as a potential explanation of rising inequality over the last 30 years. The basic premise is that skilled workers are better able to cope with new technology, and in particular, the changes imposed by the "microcomputer revolution" of the 1980s. The rapid uptake of new technology by firms has therefore resulted in a shift in relative demand towards more skilled labour and as a result, the gap between the wages of skilled and unskilled workers has grown over time. In other words, SBTC is expected to result in rising between-group inequality, driven by rising returns to skills or "price effects".

Although there is widespread agreement that SBTC explains an important part of rising inequality, it probably does not tell the whole story. In that regard, Card & DiNardo (2002) review several inconsistencies between the predictions of SBTC and observed inequality since the 1980s. Primarily, they highlight that

[4] This highlights the importance (and difficulty) of establishing causality in studying labour market outcomes. Although unemployment can be seen to drive earnings and inequality, as argued by Leite et al (2006), at the same time, relatively high real wages for some workers may result in unemployment. See Institutional factors below for further details.

inequality stabilised during the 1990s, while computing technology continued to advance exponentially. They also identify declining gender inequality and stable racial inequality as trends that are difficult to reconcile with SBTC, which predicts that discrimation-related inequality will rise in conjunction with overall inequality.

South African studies also make reference to SBTC as a driver of inequality. Bhorat (2004:944-945) mentions SBTC as a driver of rising demand for highly skilled workers and growing unemployment amongst unskilled workers. Using the 1995 to 1999 OHS and 2000 to 2003 LFS, Woolard & Woolard (2006:1) find that the earnings gap between low and high skilled workers widened due to a lowering of earnings for low-skilled workers while earnings for high skilled workers remained roughly constant. Leite et al (2006:17) find that the earnings of low-skilled workers have been more adversely affected by slow growth and rising unemployment than those of high skilled workers.

2.2.2 Globalisation

A second popular explanation of rising international inequality is based on changes in the relative demand for unskilled labour as a result of increased international trade and integration. For developed countries, globalisation brings increased competition from goods produced in emerging economies, where unskilled labour is plentiful and cheap. As a result, the wages of unskilled workers in developed countries are driven downwards, and these "price effects" lead to rising earnings inequality.

In their review, Machin & Van Reenen (2007:7) find little empirical support for this theory. Most notably, the globalisation explanation predicts that as developed countries increasingly specialise in skills-intensive goods, employment should shift from less skill-intensive towards more skill-intensive industries. Within industries, on the other hand, employment should shift towards relatively cheaper unskilled labour. Instead, empirical studies have shown a rise in skilled employment within all industries in developed countries. Furthermore, the international trade explanation predicts that wage differentials in emerging countries will fall as a result of increased competition from skilled labour in developed countries. However, the actual experience of emerging countries has been in line with that of developed countries, with wage differentials growing in these countries as well.

In South Africa, various studies have highlighted the impact of globalisation and increasing openness on the demand for labour. Examining data from the Soufh African manufacturing sector, Bhorat (2000:459-460) finds that tariff liberalisation since the late 1980s has resulted in high job losses for unskilled workers, while skilled workers gained from changes in the relative demand for labour. Edwards & Behar (2006:143) find that tariff liberalisation did not have a significant impact on earnings inequality between 1994 and 2003. They go on to note that South African labour market rigidities may have encouraged firms to shed labour in the face of increased international competition. Therefore, rather than causing an increase in inequality, globalisation may instead have resulted in rising unemployment.

2.2.3 Changes in the composition of the labour force

As pointed out by Lemieux (2006b:197), changes in inequality can also be driven by changes in the composition of the labour force. More specifically, it can be shown that the variance of wages will be larger for more educated and experienced workers, and the impact of rising skills prices will be more pronounced for highly skilled workers. As a result, if the labour force is becoming systematically older and more educated (as has been the case in developed countries), overall inequality will increase. Indeed, Lemieux (2006a:464) finds that a large portion of increases in US inequality after 1980 was actually the result of composition effects.

National survey data indicates that the South African labour force has become more educated since 1994. However, there are potential issues regarding both the quality of educational data and differences in actual quality of education, and as a result, overall levels of education are probably lower than indicated by survey data (Seekings, 2007:18). Rospabe (2002) investigates the evolution of racial wage discrimination using the 1993 PSLSD and 1999 OHS and finds that differences in the skills composition accounts for between 70% and 80% of the gap between White and African wages. She concludes that the slight improvement in African earnings versus White earnings was mainly due to improvements in the skills composition of Africans (Rospabe, 2002:209).

2.2.4 Institutional factors

Changing institutional factors, such as decreasing unionisation or minimum wages, can also result in rising inequality. The decline in the real value of the minimum wage, in particular, has been shown to explain a significant portion of the increase in inequality within the bottom half of the US wage distribution during the 1980s. Similarly, rising inequality occurred at the same time as declining unionisation in the US and UK. The different experience of other developed countries with stricter labour legislation such as Germany, France and Japan, where inequality has not risen as much as in the US or UK, supports the notion that institutional factors are important determinants of overall inequality. Indeed, it may be that due to the relative inflexibility of wages in these countries, SBTC-related pressures have instead resulted in rising unemployment (Machin & Van Reenen, 2007:8-9).

In South Africa, institutional factors have been widely cited as having a restrictive influence on employment. Since 1994, various changes to existing labour legislation have been enacted with the aims of further extending the reach of collective bargaining, increasing minimum wage coverage and promoting employment equity[5]. This has significantly strengthened the union movement and driven a wedge between formal sector wages and competitive market wages.

[5] Relevant legislation includes the Labour Relations Act of 1995, the Basic Conditions of Employment Act of 1997 and the Employment Equity Act of 1998.

Indeed, as noted by Kingdon & Knight (2007:816), high real wages due to labour market rigidities has had a detrimental effect on employment, resulting in a growing chasm between employed "insiders" and those without formal sector employment.

2.2.5 Measurement error

A final (but no less important) issue raised by Lemieux (2006a) is that of "noisy data" or measurement error. It can be shown that if the level of measurement error in earnings data increases over time, earnings inequality will also increase, and vice versa. By comparing two different datasets, he shows that measurement error may have overstated the level and growth of residual inequality observed in US data. Given that South Africa's national statistics agency only started collecting comprehensive household survey data in the mid-1990s, the impact of measurement error is bound to be non-negligible. The changeover from the October Household Survey to the Labour Force Survey in 2000, for instance, may have resulted in lower measurement error since 2000, which in turn would have resulted in lower estimates of inequality.

3 Methodology

This study examines changes in individual earnings inequality between 1995 and 2006 for various subgroups, in order to distinguish between the relative importance of price effects (changes due to SBTC or globalisation), compositional effects (changes due to education or the impact of HIV/AIDS) and other effects (institutional factors or measurement error). Kernel density estimates of earnings are the first distributional measures that are examined, and provide a convenient visual representation of the entire distribution[6]. The analysis also includes other popular summary measures of earnings inequality, namely the variance, 90-10 gap, 90-50 gap and 50-10 gap of earnings[7]. Since the focus is mainly on the impact of human capital variables (specifically education and experience) on earnings, quantile regressions of the 10th, 50th and 90th percentiles of earnings on years of education at the mean age in the given population are used to examine this relationship[8].

Residual or within-group inequality is also analysed, and can be defined as inequality remaining after human capital variables (conventionally, experience and education) have been controlled for, i.e. earnings dispersion among workers with the same education and experience (Lemieux, 2006:461). In a simple human capital model, residual inequality therefore represents the portion of inequality that is due to "unobserved skills", such as ability, motivation and quality of schooling.

[6] All kernel densities were estimated using Stata's default Epanechnikov kernel and bandwidth.

[7] The 90-10 wage gap, for example, is calculated as the difference between the ninetieth and tenth centile of earnings.

[8] This is simlar to the method used in Lemieux (2006b:195-196).

By estimating a regression of the log of earnings on education and experience variables, it is possible to calculate a variety of measures of residual dispersion based on the residuals from this regression[9]. The analysis below focuses specifically on the variance as measure of dispersion. Unlike other measures of dispersion, such as the 90-10 gap, it is possible to decompose the overall variance into an explained and residual component, which facilitates comparisons between the overall and residual variance. As noted by Lemieux (2006:463), the total variance of earnings can be written as $Var(w_{it}) = Var(x_{it}b_t) + Var(\varepsilon_{it})$, where $Var(\varepsilon_{it})$ is the residual variance. As such, the residual variance is the fraction $(1 - R^2)$ of the overall variance of earnings. The same is not true for the 90-10 gap of earnings, where the 90-10 gap of w_{it} is generally not equal to the sum of the 90-10 gap of $x_{it}b_t$ and the 90-10 gap of ε_{it}.

Given the emphasis that skills as a determinant of earnings has enjoyed in the international and South African literature, it is also the focus of this particular study. Lemieux (2002 and 2006a) presents an appealing decomposition framework to investigate the contribution of changes in skills to changes in earnings. His method consists of a "unification" of approaches originally presented by Juhn, Murphy and Pierce (1993) ("JMP") and DiNardo, Fortin and Lemieux (1996) ("DFL"), and Lemieux (2002:648) describes it as "a generalisation of [the] well-known Oaxaca-Blinder decomposition of means to the full distributional case". In short, it provides a way to quantitatively assess the relative contributions of changes in the distribution of skills ("composition effects") and changes in the returns to skills ("price effects") to changes in earnings inequality. As such, it is well suited to test the predictions of the SBTC hypothesis and has the advantage of being relatively straightforward to apply to a variety of distributional measures[10].

3.1 Adjusting earnings for composition effects

The method followed here to adjust earnings for changes in the distribution of covariates was first proposed by DFL (1996) and further elaborated on by Lemieux (2002 and 2006a). Assume that the distribution of earnings $g(w)$ is given by

$$g(w) = \int f(w|x)h(x)dx$$

where $f(w|x)$ is the conditional distribution of earnings and $h(x)$ is the marginal distribution of observed characteristics x. The actual density of earnings in

[9] The earnings regression used in the analysis below contains a quadratic in years of education and a quartic in age, as well as interaction terms between a quadratic in education and a quartic in age.

[10] Lemieux's methodology has not explicitly been applied to studies of inequality using South African data, although the DFL and JMP methodology have been applied separately in at least two studies. Leibbrandt et al (2005b) apply the DFL methodology (which, as shown below, consists of a reweighting to account for composition effects) to data on individual incomes from the 1995 and 2000 Income and Expenditure Survey (IES). Burger & Jafta (2006) apply the JMP methodology (which accounts for changes in the returns to skills) in order to examine racial wage gaps at different percentiles.

period b can therefore be written as

$$g(w|t=b) = \int f_b(w|x)h(x|t=b)dx$$

where $f_b(w|x) = f(w|x, t=b)$. The question of interest is what the distribution of earnings in period b would have been had the characteristics x stayed at the prior period a level. This can be expressed as

$$g_b(w|t=a) = \int f_b(w|x)h(x|t=a)dx \tag{1}$$

In order to solve for the right-hand side integral in equation 1, we use Bayes' Law to derive the following two relationships:

$$h(x) = h(x|t=b)\frac{\Pr(t=b)}{\Pr(t=b|x)} \tag{2}$$

$$h(x) = h(x|t=a)\frac{\Pr(t=a)}{\Pr(t=a|x)} \tag{3}$$

Combining equations 2 and 3, we therefore obtain

$$\begin{aligned}
h(x|t=a) &= \frac{\Pr(t=a|x)}{\Pr(t=a)}h(x) \\
&= \frac{\Pr(t=a|x)}{\Pr(t=a)}\frac{\Pr(t=b)}{\Pr(t=b|x)}h(x|t=b) \\
&= \theta\, h(x|t=b)
\end{aligned}$$

Going back to equation 1, it is now possible to write

$$\begin{aligned}
f_b(w|x)h(x|t=a) &= \theta\, f_b(w|x)h(x|t=b) \\
g_b(w|t=a) &= \int \theta\, f_b(w|x)h(x|t=b)dx
\end{aligned}$$

The counterfactual distribution $g_b(w|t=a)$ can therefore be obtained by kernel density methods using the *actual* period b earnings distribution $f_b(w|x)h(x|t=b)$ but reweighting it with the computed θ variable to ensure that the distribution of covariates resembles that of period a.

The procedure is straightforward to implement by using a standard probability model on the data for years a and b pooled together. The reweighting factor is computed using the estimates from a logit model for the probability of being in year b relative to the base year a. The dependent variable in the logit model is a dummy variable for year b, while the explanatory variables are the age and education variables[11]. The predicted probability that worker i

[11] According to Lemieux (2006a:468), the explanatory variables will be the same age and education variables as in the earnings regression used to estimate the residual variance. The specification used here includes a quadratic in years of education and a quartic in age, as well as interaction terms between a quadratic in education and a quartic in age.

with characteristics x is observed in year b, $P_i = \Pr(t = b|x_i)$, is then used to compute the counterfactual weight as

$$\omega_{ib}^{*} = \left[(1 - P_i) / P_i\right] \omega_{ib}$$

where ω_{ib} is the sample weight[12]. In this way, by simply using the adjusted weights instead of the usual survey weights when calculating inequality measures, it is possible to control for composition effects.

3.2 Adjusting earnings for price effects

The adjustment for changes in the returns to skills presented here is based on Lemieux (2002), which in turn builds on work done by JMP (1993). The starting point is the well-known Oaxaca-Blinder decomposition of changes in means between period a and b, given by

$$\overline{w}_b - \overline{w}_a = \overline{x}_b \left(b_b - b_a\right) + \left(\overline{x}_b - \overline{x}_a\right) b_a \tag{4}$$

where the first right-hand side term reflects differences in the estimated parameters, while the second term shows differences in the average values of the covariates between period a and b. In this case, the term $\overline{x}_b b_a$ can also be interpreted as the average earnings that would have been observed in period b had the returns to skills been the same as in period a. If we let

$$\overline{w}_b^{*} = \overline{x}_b b_a$$

the counterfactual $\overline{w}_b^{*}$ can be used to rewrite equation 4 as

$$\overline{w}_b - \overline{w}_a = \left(\overline{x}_b b_b - \overline{w}_b^{*}\right) + \left(\overline{w}_b^{*} - \overline{x}_a b_a\right) = \left(\overline{w}_b - \overline{w}_b^{*}\right) + \left(\overline{w}_b^{*} - \overline{w}_a\right)$$

Lemieux (2002:652) proceeds to write the individual-specific counterfactual earnings as

$$w_{ib}^{*} = x_{ib} b_a + u_{ib} = w_{ib} + x_{ib} \left(b_a - b_b\right) \tag{5}$$

Therefore, equation 5 shows that by using the coefficients from the base period a but retaining the residuals from the current period b, individual counterfactual earnings can be calculated which can be used to derive any hypothetical measure of earnings dispersion by substituting the individual counterfactual earnings w_{ib}^{*} for actual earnings w_{ib} in the calculation.

In the analysis below, the counterfactual weights and counterfactual earnings were used to calculate hypothetical versions of kernel density estimates, summary inequality measures and measures of residual inequality. The decomposition was performed by first calculating hypothetical measures by controlling

[12] Lemieux (2002:656) notes that the correction factor $P_b/(1 - P_b)$, where P_b is the unconditional probability i.e. $\Pr(t = b)$, included in the original specification of the procedure by DiNardo, Fortin and Lemieux (and in the derivation of θ shown above) can be excluded in practice. The correction factor changes the reweighting only in a proportional way, which falls away in most statistical packages since they automatically normalise the sum of weights when weighted statistics are computed.

for changes in the distribution of covariates, and then additionally controlling for changes in the coefficients. In order to test the robustness of results, the data was reweighted in two ways, by first keeping constant the 1997 distribution of covariates and coefficients, and then the 2006 distribution of covariates and coefficients[13].

4 The Data

The analysis was performed using data from PALMS (Kerr et al 2013), which brings together the national household surveys collected by Statistics South Africa (Stats SA) relevant to labour market analysis, viz. the October Household Surveys (OHS) from 1994 to 1999, the Labour Force Surveys (LFS) biannually from 2000 to 2007 and the Quarterly Labour Force Surveys (QLFSs) for 2010 and 2011. The QLFSs for 2008 and 2009 do not contain earnings information.

4.1 Earnings data

The sample was restricted to working age (15 to 65 years) employed individuals with positive hours worked and positive earnings. Earnings were adjusted for inflation using the South African Reserve Bank's CPI series (KBP7032N), since the main concern of this study is with the purchasing power of workers. In cases where earnings were reported as intervals rather than as point data, a reweighting approach was used to combine this information with the continuous earnings data.

Various authors have reported a fall in real incomes between the mid-1990s and early 2000s, based on data from the OHS and LFS. For instance, Casale et al (2004:990) report a 22% decline in average real earnings between 1995 and March 2003. Burger & Yu (2007:1) investigate this matter further, and find that the decline is mostly the result of data inconsistencies caused by the changeover between the OHS and LFS. By adjusting the data to account for inconsistencies in questionnaire design and the presence of outliers, they find that real formal sector wages have in fact increased since 1994, and that there is little evidence of declining wages in the informal economy.

Initially, results are reported using both hourly and monthly data; however, based on the observation that the trends in monthly and hourly earnings do not differ markedly, most of the key results are presented using monthly data only.

4.2 Independent variables

In order to calculate measures of residual variance, it is necessary to specify an appropriate regression model. In addition, the reweighting approach applied in

[13] As outlined below, data was examined for years between 1995 and 2006, and although it would have been ideal to compare the two outermost years in the available series, 1997 was chosen as the intial year due to problems with the 1995 and 1996 data.

this study also requires the selection of a set of independent variables for the logit model. The chosen specification follows Lemieux (2006a:469) and focuses only on arguably "pure" measures of skills, namely education and experience. A full specification was used, including a quadratic in education, a quartic in experience and interactions between a quartic in experience and a quadratic in education[14].

Education is measured in years, and the data was adjusted for inconsistencies in the coding of the education variable between surveys. Since experience is not directly measured, a suitable proxy needs to be found. Most studies in South Africa have used a measure of "potential experience", calculated as age minus years of education minus six. Keswell & Poswell (2004:836), however, highlight that the effect of potential experience is likely to be overestimated in South African data, due to the prevalence of grade repetition, low educational attainment and job insecurity. Given these difficulties, the analysis below uses age as a proxy for experience. The results from the analysis are presented separately for males and females. In addition, results for African and White males are also analysed.

5 Trends in earnings, education and experience

5.1 Earnings

Trends in earnings inequality differ according to the time period under observation, with notable differences between the OHS and LFS years. Given the concerns highlighted in the previous section, changes over the entire twelve year period should be viewed with circumspection. Although it is to be hoped that at least some of the observed change was real, given the important changes that the South African labour market underwent during this period, changes in the survey instrument may also have played a role in determining the final outcome of the data.

As shown in Figure 1, overall inequality, as measured by the variance and 90-10 gap, increased slightly between 1995 and 1999. Over the same period, the 90-50 gap shows an increasing trend, while the 50-10 gap also shows an increasing trend if 1995 is ignored. The LFS years display a consistently declining trend, if the datapoint for 2002 is ignored. This was driven by a large drop in the 50-10 gap, so that the lower half of the distribution appears to have become less unequal. Although inequality in the top half of the distribution, as represented by the 90-50 gap, increased over the same period, this was of a smaller magnitude than the changes in the bottom half of the distribution. Therefore, overall earnings inequality can be characterised as increasing between 1995 and 1999 (during the OHS years) and declining thereafter (during the LFS years). However, taken together over the entire period (and ignoring 1995, which appears to be an outlier), the net effect of these changes was close to zero, so that

[14] The calculation of the residual variance and the reweighting procedure both incorporate exactly the same set of independent variables.

Figure 1: Trends in overall earnings inequality 1995-2006

inequality in 2006 was very close to levels measured in 1997.

The following two subsections examine the trends for males and females separately, as well as for African and White males. Although the main focus of the analysis is earnings dispersion, it is useful to include a brief overview of trends in mean earnings. As can be seen in Figure 2, the trends in mean hourly and monthly earnings are broadly similar. Apart from 1995, 1997 and 2002, the overall trend in real earnings is upwards. Males and females experienced roughly the same increase in real earnings, so that the gap between male and female wages remained relatively stable. In contrast, Whites experienced a greater rise in earnings than Africans, resulting in a widening of the sizeable gap between these two groups.

Table 1 shows the sample sizes and percentages of males to females and African males to White males used in the analysis[15]. The proportion of males

[15] The difference between the total number of males in the male-female comparison and the

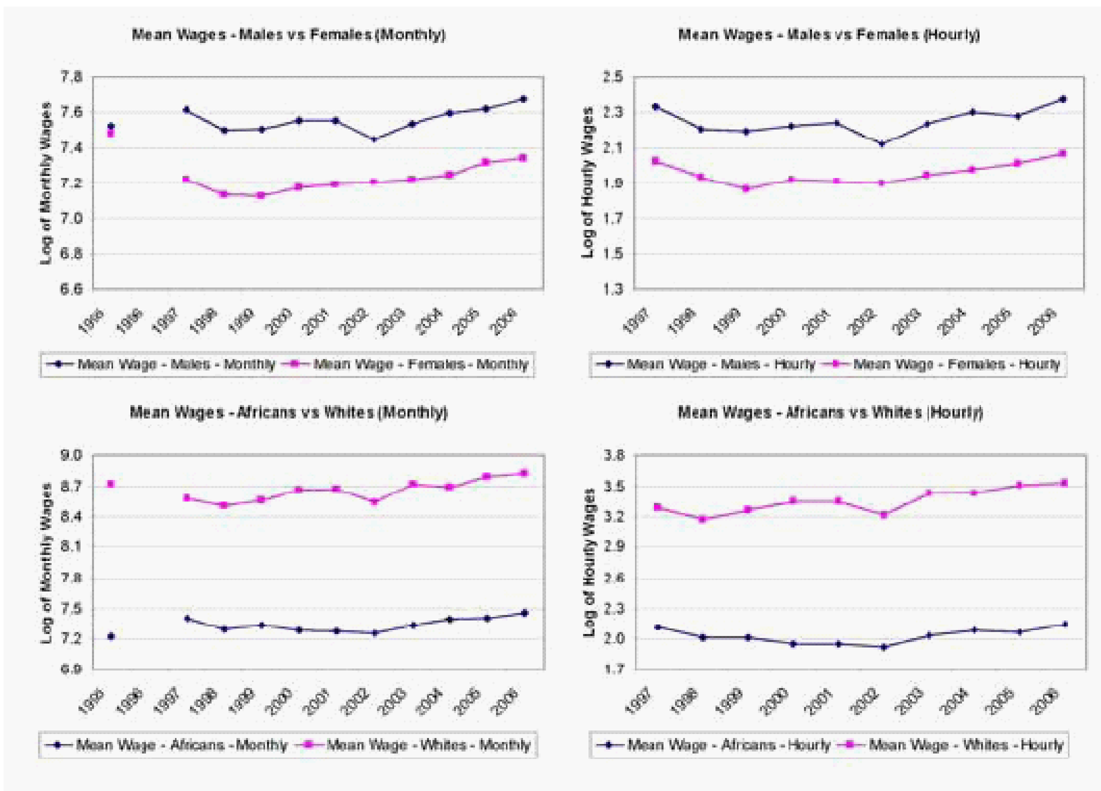

Figure 2: Trends in Mean Earnings 1995-2006: Males versus Females and Africans versus Whites

Sample Sizes											
	1995	**1997**	**1998**	**1999**	**2000**	**2001**	**2002**	**2003**	**2004**	**2005**	**2006**
Male	12,359	7,651	4,261	5,459	8,412	7,394	7,424	6,574	6,663	6,797	7,232
Female	6,209	5,429	2,765	3,883	6,210	5,426	6,817	4,677	5,048	4,954	5,546
African	7,768	5,138	3,183	4,061	5,997	5,394	5,260	4,822	4,701	4,812	5,116
White	1,895	617	332	330	754	645	514	428	423	429	400
Proportions											
	1995	**1997**	**1998**	**1999**	**2000**	**2001**	**2002**	**2003**	**2004**	**2005**	**2006**
Male	67%	58%	61%	58%	58%	58%	52%	58%	57%	58%	57%
Female	33%	42%	39%	42%	42%	42%	48%	42%	43%	42%	43%
African	80%	89%	91%	92%	89%	89%	91%	92%	92%	92%	93%
White	20%	11%	9%	8%	11%	11%	9%	8%	8%	8%	7%

Note: The sample consists of working age individuals in full-time formal sector employment with positive earnings. Proportions are based on unweighted sample sizes. The African-White comparison is based on a subsample of males only.

Table 1: Sample Sizes and Proportions: Males versus Females and Africans versus Whites

is relatively constant at around 58% for most years. The exceptions are 1995, with 67% males, 1998 with 61% males and 2002 with 52% males in the sample. Therefore, 1995 and 1998 contain relatively more males in the sample compared to other years, while 2002 contains relatively fewer males in the sample. The proportion of African males is also fairly similar between 1995 and 2006, remaining in a range of between 89% and 93% of the African-White sample. Once more, 1995 seems to be an outlier, with only 80% Africans in the sample. The sample sizes for the African-White comparison further presents a problem, in that the number of White males in the sample is very small (as low as 330 in the case of 1999). Therefore, it is worth noting that for some years, the relatively small sample size for White males may impact the analysis below.

From the analysis above, the datapoint for 2002 appears to be an outlier, both in terms of overall earnings inequality (which is higher than in other years) and average earnings (which is lower than in other years). Although this feature of the LFS data has not been explicitly noted or explored in the literature, it is nonetheless evident in other studies. Leite et al (2006:16) calculate the Gini index and Theil-L index using OHS and LFS data, which results in markedly higher measurements of inequality in 2002 compared to other years. While Burger & Yu (2007:11) do not analyse earnings inequality for 2002, they do examine trends in mean earnings. In this case, mean earnings are lower in the September 2002 LFS data compared to other years[16]. One possible explanation for the earnings outlier is that more females (i.e. relatively fewer males) were

total number of males in the African-White comparison is due to the omission of the remaining population groups ("Coloured" and "Indian/Asian") from the analysis.

[16] Burger & Yu (2007) use both the March and September datasets for the LFS years. As such, it is necessary to exclude the March datapoints when comparing their analysis to the results presented in this study.

sampled in 2002 compared to other years. Indeed, as shown in the following sections, higher inequality appears to have been driven by the female subgroup (and to a lesser extent, African males). However, at the same time it is puzzling that higher means appear to have been driven by the male subgroup.

5.1.1 Males versus Females

Kernel density estimates of the log of monthly earnings for males and females are shown in Figure 3[17]. As expected, female densities are flatter and slightly to the left of those for males, indicating that female earnings are lower than male earnings and are more widely dispersed[18]. Male and female densities for 1995 appear to coincide the most, with marked differences between male and female densities for the remaining years. The unimodal shape of the densities for males have become less peaked over time. Female densities have undergone a much more dramatic change in shape over time: first flattening out over the OHS years, then assuming a distinct bimodal shape during the earlier LFS years, and finally moving closer again to the male distribution from 2004 onwards.

Figure 4 shows changes in the male-female wage gap over time, which gives an indication of the "between-group" difference between males and females. The wage gap is calculated as the difference between the mean value of the log of monthly or hourly earnings for males and that for females. The observation for 1995 is clearly an outlier when compared the rest of the years. As noted above, relatively more employed males were surveyed in 1995 compared to other years, which may explain the outlier. However, this raises the more fundamental question as to why the proportion of males was higher in 1995. The observation for 2002 also appears to be an outlier for both hourly and monthly earnings. Again, the sample proportion of females versus males was different in 2002 compared to other years (in this case, relatively more females were included in the sample), although it is puzzling that this should lead to an outlier in the same direction as the one for 1995.

Despite the rather volatile trend in the male-female gap based on hourly earnings, on average it seems to have remained at around 0.3 on a log scale. The male-female gap based on monthly earnings shows a slightly declining trend between 1997 and 2006 from around 0.4 to 0.3 on a log scale. Lemieux (2006a:464) notes that increases in measures of "discrimination", such as the male-female or White-African wage gap, can be interpreted as evidence of increasing returns to unobserved skills. As such, the stable or declining trend in the male-female wage gap does not support the view that increases in inequality

[17] Since the general trends are the same for hourly and monthly earnings, kernel density estimates of the log of hourly earnings are shown in the Appendices.

[18] See Cahuc & Zylberberg (2004:290) for a survey of international evidence of wage differences between men and women, as well as explanations for the general finding that women's wages are systematically lower than those of men.

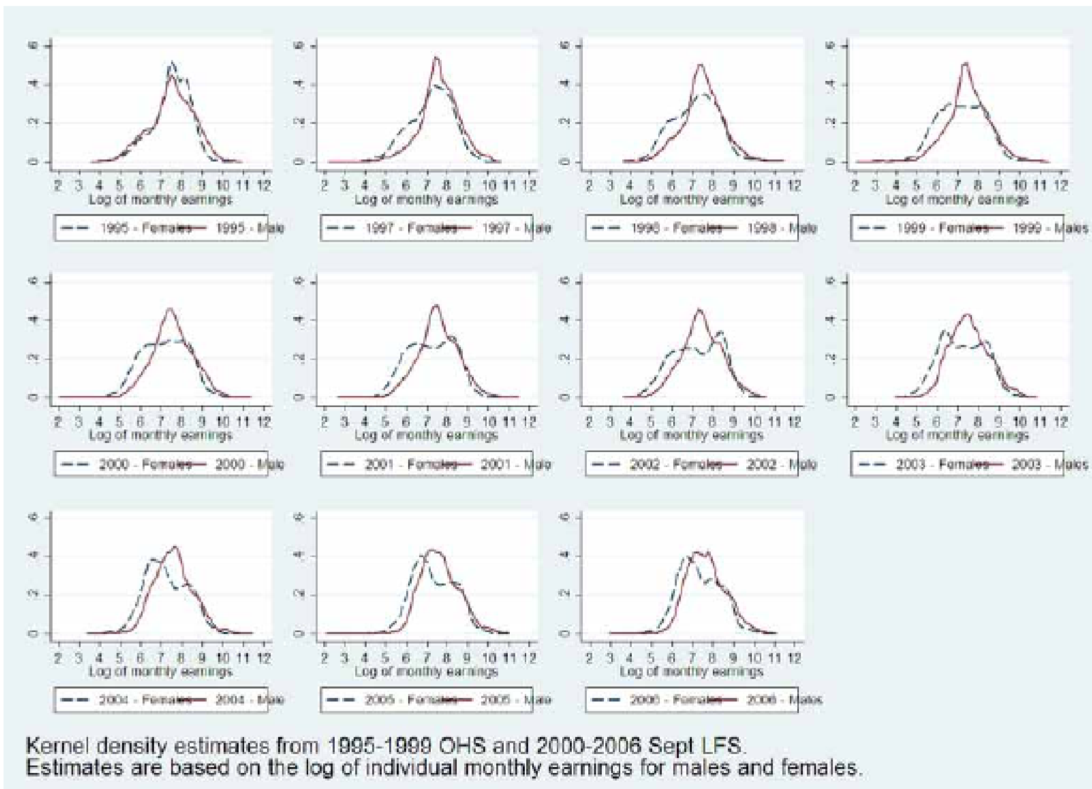

Kernel density estimates from 1995-1999 OHS and 2000-2006 Sept LFS.
Estimates are based on the log of individual monthly earnings for males and females.

Figure 3: Shifts in real earnings distributions 1995-2006: Male vs Female (monthly earnings)

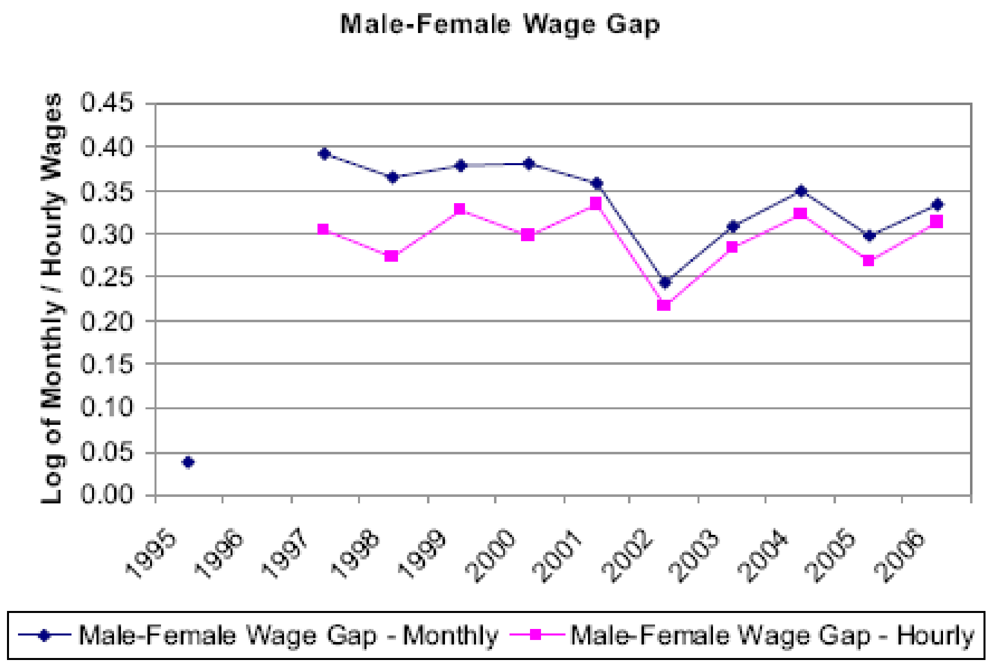

Figure 4: Shifts in real earnings distributions 1995-2006: Male-Female Wage Gaps