REDI3x3 Working paper 56 May 2019

Wage determination in South Africa: The case of the union wage premium since the end of apartheid

Martin Wittenberg and Andrew Kerr

Abstract

South African inequality, both in overall income and in labour earnings, has increased on many measures, even accounting for changes in the characteristics of the population. This raises the question as to what has led to these changes. One of the explanations for the rise in inequality in developed economies is the collapse in union power. South African unions, by contrast, are still influential and have even been accused of increasing inequality by raising the wages of insiders at the expense of the non-unionised. We locate the analysis of the union wage premium in South Africa in the context of measurement and data quality issues, estimation methods, the evolution of unionism and the intersection between unionism and other institutions. We briefly review studies on union density and the union premium in the other BRICS economies (Brazil, Russia, India and China) also, and compare our results for South Africa to the results from these countries obtained in the literature.

The **Research Project on Employment, Income Distribution and Inclusive Growth** is based at SALDRU at the University of Cape Town and supported by the National Treasury. Views expressed in REDI3x3 Working Papers are those of the authors and are not to be attributed to any of these institutions.

Wage Determination in South Africa: The case of the union wage premium since the end of apartheid*

Martin Wittenberg
DataFirst and School of Economics
University of Cape Town

Andrew Kerr
DataFirst
University of Cape Town

May 2019

1 Introduction

South African inequality, both in overall income and in labour earnings, has increased on many measures, even accounting for changes in the characteristics of the population. This raises the question as to what has led to these changes. One of the explanations for the rise in inequality in developed economies is the collapse in union power. South African unions, by contrast, are still influential. Indeed some authors (Nattrass and Seekings 2014) have suggested that union power has contributed to increased inequality in earnings, by raising the wages of "insiders" at the expense of the non-unionised.

Several papers have investigated the size of the union wage premium in South Africa but none have investigated systematically how the union premium has evolved over the post-apartheid period. Most estimate the premium at one point in time and at best compare a handful of surveys. Surprisingly, none of the studies (even the ones that have explicitly tried to track the evolution of the union wage premium) have looked at the changes in South African unionism over this period.

In this paper we locate the analysis of the union wage premium in South Africa in the context of measurement and data quality issues, estimation methods, the evolution of unionism and the intersection between unionism and other institutions. Since this is a special issue focusing on , we also briefly review studies on union density and the union premium in the other BRICS economies (Brazil, Russia, India, China and South Africa) also, and compare our results for South Africa to the results obtained in the literature from these countries.

We begin the paper with a literature review which considers first some of the methods used in the estimation of union gaps and then turn to look at the details of the South African literature, as well as the literature for the other BRICS countries that has looked at this question. We then discuss the data that we will use to estimate the union premium in South Africa, viz. version 3.1 of the Post-Apartheid Labour Market Series (PALMS) in section 3. Section 4 presents a discussion of the methods that we will use. We draw attention to a connection made by Fortin, Lemieux and Firpo (2011) between the estimation of union effects and standard decomposition techniques. The results and discussion follow. We draw some lessons in the conclusion.

*This research has been supported by REDI3x3. Many thanks to Caroline Stapleton who provided excellent research assistance.

2 Literature Review

2.1 Big picture overview of unions

Rising inequality around the world has coincided with the decline of union power in many developed countries. Unions are generally thought to raise the wages of their members. In this section we review the literature on the union wage premium and how it is measured. We review the estimates of the union premium and union density in South Africa and the other BRICS countries.

2.2 Estimating the union wage premium

Researchers have used widely differing estimation techniques, regression specifications and samples to estimate union premia in a large number of countries. As a result the estimates are not always comparable. There are three main methods that have been employed in South Africa, and since these are also discussed in Lewis's (1986) book on the union premium, and are used in the studies in the other BRICS countries that we discuss below, we describe these three methods in the following sections.

2.2.1 Single equation estimation with a union dummy (OLS)

This involves estimating the equation

$$\ln w_i = \mathbf{x}_i \boldsymbol{\beta} + \delta U_i + \varepsilon_i \tag{1}$$

where $\mathbf{x}_i$ is a vector of characteristics determining earnings U_i is the union membership dummy variable and ε_i is an idiosyncratic error term. This method is also referred to as the "treatment effects" estimate in some parts of the South African literature, although we argue below that this is misleading terminology. Estimation is typically by OLS, although it has been estimated with sample selection corrections ("Heckman" two-step estimates), by quantile regression, interval regression and household or individual fixed effects.

2.2.2 Two equation estimation (Lewis)

The preferred estimation method, according to Lewis's (1986) magisterial review, is to estimate separate equations for union and non-union members

$$\ln w_i^u = \mathbf{x}_i \boldsymbol{\beta}^u + \varepsilon_i^u \tag{2}$$
$$\ln w_i^n = \mathbf{x}_i \boldsymbol{\beta}^n + \varepsilon_i^n \tag{3}$$

and calculate the union wage premium as[1]

$$\overline{\mathbf{x}} \left(\widehat{\boldsymbol{\beta}}^u - \widehat{\boldsymbol{\beta}}^n \right) \tag{4}$$

This approach follows intuitively from an application of the Oaxaca-Blinder decomposition method. The "raw" difference in the log wage rate can be decomposed in two ways

$$\begin{aligned} \overline{\ln w}^u - \overline{\ln w}^n &= \left(\overline{\mathbf{x}}^u - \overline{\mathbf{x}}^n \right) \widehat{\boldsymbol{\beta}}^u + \overline{\mathbf{x}}^n \left(\widehat{\boldsymbol{\beta}}^u - \widehat{\boldsymbol{\beta}}^n \right) \\ \overline{\ln w}^u - \overline{\ln w}^n &= \left(\overline{\mathbf{x}}^u - \overline{\mathbf{x}}^n \right) \widehat{\boldsymbol{\beta}}^n + \overline{\mathbf{x}}^u \left(\widehat{\boldsymbol{\beta}}^u - \widehat{\boldsymbol{\beta}}^n \right) \end{aligned} \tag{5}$$

[1]This is actually the log wage gap, which is often converted to a percentage change, but there seems little benefit to doing so.

The first term of each version of the decomposition is the "explained" part of the difference in wages, due to the difference in average characteristics. The second term is the price effect, due to different returns in the two sectors, keeping the characteristics fixed. The union premium given in equation 4 is a weighted average of the two price effects. A useful way of thinking about this is to consider a non-unionised worker with log earnings $\ln w_i^n$. This can be written as $\mathbf{x}_i\widehat{\boldsymbol{\beta}}^n + \widehat{u}_i^n$, where $\widehat{u}_i^n$ is the OLS residual. For this worker we can predict what they would earn if they were paid like a unionised worker (taking into consideration their idiosyncratic error). It would be $\mathbf{x}_i\widehat{\boldsymbol{\beta}}^u + \widehat{u}_i^n$. Their predicted "wage gain" is therefore $\mathbf{x}_i\left(\widehat{\boldsymbol{\beta}}^u - \widehat{\boldsymbol{\beta}}^n\right)$. Averaging the predicted wage gains over the non-unionised workers we get $\overline{\mathbf{x}}^n\left(\widehat{\boldsymbol{\beta}}^u - \widehat{\boldsymbol{\beta}}^n\right)$. Similarly the term $\overline{\mathbf{x}}^u\left(\widehat{\boldsymbol{\beta}}^u - \widehat{\boldsymbol{\beta}}^n\right)$ is the average predicted "wage loss" if a unionised worker were to be paid like a non-unionised worker. More positively, it is the realised gain over the hypothetical case where they were not union members. The overall premium is therefore the expected/realised wage gain resulting from being unionised, averaged across the population.

2.2.3 Switching regression (Lee)

The third approach, based on Lee (1978), adds an endogenous switch to the two equation model of the previous section. The latent variable underlying the "switch" is written as

$$I_i^* = \alpha_0 + \alpha_1\left(\ln w_i^u - \ln w_i^n\right) + \mathbf{x}_i\boldsymbol{\alpha}_2 + \mathbf{z}_i\boldsymbol{\alpha}_3 + \eta_i \tag{6}$$

$$U_i = \mathbf{1}\left(I_i^* > 0\right) \tag{7}$$

where η_i is a standard normal random variable and $\mathbf{1}(.)$ is the indicator function. After substituting in for $\ln w_i^u$ and $\ln w_i^n$ the reduced form of the "switch" is a straight-forward probit model with explanatory variables $\mathbf{x}_i$ and $\mathbf{z}_i$. Assuming that $(\varepsilon_i^u, \varepsilon_i^n, \eta_i)$ are trivariate normal, the two regressions (equations 2 and 3) can be estimated by OLS provided that a Heckman-type correction (inverse Mills ratio) is added. The formula for the union premium is now more complicated, since the predicted values for the log earnings (given the characteristics $\mathbf{x}_i$ and $\mathbf{z}_i$) will depend also on the selection terms. In principle, the predicted values of $\ln w_i^u$ and $\ln w_i^n$ can be substituted back into equation 6 to derive estimates of the structural model. In practice, however, even the estimation of "selectivity corrected" earnings equations (2 and 3) is very difficult because Heckman-type corrections are not robust in the absence of exogenous instruments $\mathbf{z}_i$ and it is hard to find plausible instruments that would influence unionisation but not the wage. Lewis's (1986) review ended up quite sceptical of the selectivity-corrected estimates in the US literature. Arguably the South African experience with selection corrections has been worse, with all studies that have attempted to do this using instruments that yield logically inconsistent models (Wittenberg 2014). Several authors have commented on the fact that the point estimates with these corrections are often quite wild (Hofmeyr and Lucas 2001, Casale and Posel 2010).

2.3 The union premium in South Africa

We turn now to an examination of some of the findings. Several themes can be identified in the South African literature. These include the question of whether the union premium is "big"and what this indicates about union power in the South African economy, the evolution of the size of the premium since the end of apartheid, the relationship between the union premium and racial discrimination and between the union premium and central bargaining and the public sector and the role of selection into union status.

2.3.1 The size of the union premium in South Africa

In Table 1 we show a range of estimates that are most comparable to each other and to the estimates that we will produce later. We have not reported "selectivity adjusted" results, except where these were the only ones that the study reported (e.g. in the case of Azam and Rospabé 2007). This is because the "selectivity adjusted" estimates are less stable than the standard OLS ones. Consequently the estimates shown in Table 1 are not necessarily the authors' preferred estimates, although they are typically in the same ballpark. For the sake of completeness, the Table does indicate all the types of estimations conducted by each study, although we only report the single equation/OLS results.

Moll (1993) estimated the premium for Black (i.e. African, Asian and Coloured) blue collar workers to be 24 per cent according to his preferred estimates (a version of the Lee procedure), using the 1985 Bureau for Market Research (BMR) Survey. He provides international comparative evidence to suggest that this is a "big" effect. Indeed, Lewis's (1986) meta-analysis of more than one hundred studies from the USA suggested a mean wage gap of around .15 log points in 1973.

Schultz and Mwabu (1998) using the Project for Living Standards and Development (PSLSD) find much bigger effects (around a 100% increase) at the 10th percentile of the conditional earnings distribution, suggesting that union power was, if anything bigger by 1993. Butcher and Rouse (2001, Table 4a, p.364) find a total union effect of 0.325 for Africans, according to their preferred estimates, but the union gap within sectors covered by an Industrial Council was 0.224. They argue that "these estimates are similar to estimates using data from other countries such as the United States and the United Kingdom in studies that attempt to control for individual heterogeneity" (Butcher and Rouse 2001, p.369).

It is interesting to note that Moll and Butcher and Rouse came to very different conclusions on whether or not the premium was "large" with very similar point estimates. Of course a glance at Table 1 will confirm that many of the estimated premia are now considerably higher than these. The solitary outlier on the down side is the study by Bhorat, Goga and Van Der Westhuizen (2012) in a specification that controls for various job characteristics, such as pension, medical aid and a written contract. We will argue below (section 4.3.1) that these are "bad controls" and that this estimate should therefore not be trusted. By contrast the figure given by Azam and Rospabé (2007) is on the astronomical side, probably because it uses one of the dubious selection corrections.

2.3.2 How has the union premium in South Africa evolved?

A key question from early in the literature was to determine how the size of the premium has been affected by the end of apartheid. Hofmeyr and Lucas (2001) look at the evolution of the premium between 1985 and 1993 and conclude that it increased in size. Banerjee, Galiani, Levinsohn, McLaren and Woolard (2008) looked at the period 1995 to 2004 and concluded that the premium ballooned (to over .35), although there was some evidence that it then decreased slightly between 2000 and 2004. Ntuli and Kwenda (2014) investigate premiums every third year between 2001 and 2010. They find a premium of 0.45 at the beginning, but it has steadily come down since then to 0.35 at the end.

2.3.3 The premium and racial discrimination in South Africa

A quick glance at Table 1 shows that most studies thus far (with the sole exception of Banerjee et al. (2008)) have estimated separate union premia for Whites and Africans (sometimes for Blacks, i.e. Africans, Coloureds and Indians together). Moll (1993) was interested in how the formal incorporation of African trade unions affected the labour market and given the racially segregated history of unionism it made perfect sense to do so. Studies of the immediate post-apartheid era were also interested in how the evolution of labour legislation would interact with racial discrimination and from this perspective it again made sense to estimate union premia separately for Whites and Africans. It is less clear why this approach has persisted. There are still differences between unions aligned broadly with the African National Congress (ANC) and other unions, but the fault lines seem less racial and more ideological and skill-based. It is also no longer the case that "African unions" are necessarily aligned with the ANC. Furthermore Black individuals (notably Indians and Coloureds, but also increasingly many Africans) have moved into the professional and managerial positions traditionally occupied by Whites. Clearly racial differences persist, but it seems odd to investigate the union premium for a subset of the labour market, rather than estimating the average level for the economy as a whole.

Similarly, the restriction of many of the studies to male samples seems strange given the increasing importance of women in the labour market (Casale and Posel 2010).

2.3.4 The union premium and centralised bargaining in South Africa

Moll (1996) and Butcher and Rouse (2001, Table 4a, p.364) draw attention to the fact that unions have the capacity to influence wages outside their immediate membership via the extension of wage agreements reached at Bargaining Councils to other parties. Both studies find that coverage by Bargaining Councils increases wages even for non-members. More recently Bhorat et al. (2012) find a Bargaining Council premium of around 0.09 using 2005 data.Unfortunately information on Bargaining Council coverage is not easily available, so this aspect of wage determination has not been systematically explored over many different datasets.

2.3.5 The union premium and the public sector in South Africa

A number of studies have noted that the introduction of the Public Service Co-ordinating Bargaining Council (PSCBC) in 1997 changed the nature of wage determination in the public sector (Bhorat et al. 2012). Casale and Posel (2010) argue that the compression of wages in public sector professional jobs, particularly nursing and teaching, has paradoxically led to a higher gender wage gap in the unionised sector than in the non-unionised one. Kerr and Teal (2015) have shown that there is a sizable public sector premium on top of the standard union premium.

2.3.6 What is the impact of selectivity?

Given the size of the estimated union coefficients shown in Table 1, it would stand to reason that employers would become more selective about whom to hire. Standard labour market theory would lead one to suspect that at least part of the premium would be due to unmeasured differences in skills. This is why so many of the studies try to control for selection econometrically. As noted above most of the Heckman style corrections are vitiated by the choice of improper or implausible instruments. Some of the studies tried to correct for selection by adding household fixed effects (e.g. Butcher and Rouse 2001). This strategy works only if the same unmeasured skills are common within a household. Some studies have used individual fixed effects with panel data (Casale and Posel 2010, Kerr and Teal 2015). They show that there is still a union premium, but a markedly smaller one. Kerr and Teal (2015) note, however, that fixed effects regressions are susceptible to the effects of measurement error in union status, which is likely to bias the coefficients towards zero.

2.4 The union premium and union density in BRIC countries

In the BRICS countries the context in which unions operate varies substantially. Russia and China both had unions under socialism that were essentially parts of the state. In Russia this situation ended in the late 1980s, whilst in China unions have been state-controlled (Clarke 2005). (Zhu, Warner and Feng 2011) argue that China's move towards being a market economy has led to some changes in the role of unions, including the emergence of "grass-roots workers' protection groups" that were independent from the state-controlled All-China Federation of Trade Unions (ACFTU), which has had to respond to the grass roots workers movements. In Brazil unions have operated in a state-controlled system that results in unions having limited freedom in their activities (Boito 1994). Union activity declined in Brazil in the 1990s but has increased in the 2000s, playing an important role in the support for the PT political party (Boito and Marcelino 2011). In India there are large trade unions associated with both major national political parties (Roychowdhury 2011). In India unions are especially powerful in public enterprises (Roychowdhury 2011), similar to South Africa, as we discuss further in our analysis below.

Estimates of the union premium using individual level data could be found for all four of the BRICS countries besides South Africa. In India (Das 2008) finds that 17% of wage workers in India were unionised, and that this varied substantially by state. He argues that, contrary to the developed country literature and similar to South Africa, there is no evidence of overall declining union density, but that the data quality are not good enough to make firm conclusions about the direction of change. (Bhandari 2010) uses household survey data to examine the union premium for Indian manufacturing workers, finding that it is around 20%. The single equation model with OLS was used.

In Brazil there are a few studies estimating the union premium and union density. (Arbache 1998) used the National Household Survey (PNAD) household survey data to show that union density declined from 22 percent in 1986 to 17 percent in 1999, whilst (de Oliveira Cruz and Naticchioni 2012) use later waves of the PNAD household survey and find that union density was around 18 percent in 2002 and 2009. This paper differed from (Arbache 1998) because it included both rural and urban areas. Both papers use the single equation union dummy method with OLS discussed above to estimate union premia. (Arbache 1998) finds a premium of around 15-20 percent between 1986 and 1999 whilst (de Oliveira Cruz and Naticchioni 2012) find a premium of around 10% in 2002 and 9% in 2009- for urban workers only. Both papers include controls for education, experience, sex, and region, whilst (de Oliveira Cruz and Naticchioni 2012) also include controls for race, occupation and industry.

In China (Mishra and Smythe 2014) estimated a union premium for a sample of 800 workers from 80 firms. The union premium was around 15% when controlling for industry, occupation, gender, experience and education. (**?**) undertook a similar exercise using data from 3200 employees from 162 public enterprises. 83% of the firms had a union present. The union premium was around 5% in 1994-1996 but was 15% in 1999-2001 in urban firms and it was 15% and 10% for rural firms. There do not seem to have been attempts to measure the union premium in China with representative household survey data.

We could not find specific union density figures for Russia, but (Lehmann and Muravyev 2009) showed that union density declined from about 80% in 1995 to 56% in 2007 the 10 members of the Commonwealth of Independent States, of which Russia was by far the largest. (Lukiyanova 2011) argues that, despite these high rates, trade unions are not influential in wage bargaining, which mostly occurs at the firm level. This accords with the one study that estimated a union premium for Russia. (Clarke 2002) used household survey data collected by Institute for Comparative Labour Relations Research (ISITO) to investigate the determinants of wages. The author reports that the union dummy included in an OLS regression with a sample of approximately 3000 individuals was not statistically significant but the coefficient was not reported.

2.5 Issues in estimating the union premia

2.5.1 Selection

As noted above, many of the studies tried to correct for sample selection into union status. The key problem is that the Heckman type corrections have invariably used the number of *other* union members in the household as instrument. Empirically it appears that the estimates are sensitive to the inclusion of this variable (Casale and Posel 2010). Unfortunately its inclusion makes the model logically inconsistent (Wittenberg 2014). As Wittenberg (2014) notes, this is a much bigger problem than lack of identification. The latter occurs when the data cannot ever pinpoint the true parameter vector within the parameter space. The former means that there is no possible vector that will ever satisfy the model. Fundamentally the problem arises because an outcome (union status) is being used to predict the probability of that outcome.

Besides union membership a number of authors have also worried about selection into employment. Hofmeyr and Lucas (2001) attempt to control for simultaneous selection into employment and unionisation via a multinomial logit model. However their diagnostics suggest that the selection into employment does not add information while selection between union and non-union status matters greatly. Ntuli and Kwenda (2014) add Heckman style correction terms for selection into employment into the probits for estimating union membership, which in turn yields a selection correction term for the union wage regression. This is unlikely to be a valid procedure.

2.5.2 Data quality

Most authors that have estimated union premia have not discussed how they dealt with missing earnings information. A notable exception is Rospabé (2002) who uses interval regression in order to deal with bracket information. This, of course, assumes that the errors are distributed normally. Azam and Rospabé (2007) want to do a sample selection correction which is not possible with interval regression, so they convert bracket information to midpoints[2].

Besides the case of bracket responses, however, there are many other data quality issues (Wittenberg 2016b). In particular there is a noticeable discontinuity between the October Household Survey (OHS) series and the Labour Force Survey (LFS) information. Furthermore there are outlier problems in the earnings information in several of the surveys, in particular 1999 and September 2000.

2.5.3 Specification

Some authors control for hours worked, whereas others do not. Some of these differences arise from differences in the quality of the hours information in the datasets being used (this is true particularly of the 1993 PSLSD). Even where the log of the hourly earnings rate is the dependent variable, it is not always the case that log hours is also among the control variables on the right hand side.

There are many commonalities in the regression specifications, as Table 1 indicates. There is some doubt among analysts whether or not to include sectoral and occupation dummies. Some studies (e.g. Schultz and Mwabu 1998) begin without these variables (and then find an even

[2]It is unlikely that this works, since the actual earnings figure will be the midpoint plus a deviation. That deviation will be added to the standard regression error as a result of which the assumptions underlying the Heckman-type correction will no longer hold.

bigger union premium) and subsequently add these controls into the specification. Some studies do not control for occupations at all, while all of the results reported in Table 1 control for industry/sector at least at the one digit level. However there is a good reason for controlling for both. As Butcher and Rouse (2001, p.361) note "Researchers typically control for industry and occupation when estimating union wage gaps in order to more closely approximate an experiment in which randomly chosen workers are made union members without changing their occupation or industry"

2.6 Surveys used

It is striking that even studies that explicitly try to track the evolution of the union premium (Banerjee et al. 2008, Ntuli and Kwenda 2014) only use a handful of surveys. The risk with such a procedure is that idiosyncratic shifts in measurement in one or two of the surveys can contaminate the conclusions reached (Wittenberg 2016b, Wittenberg 2016c).

2.7 What have we learned about the union premium in South Africa and in BRICS?

Looking at Table 1 most estimates of the union premium in South Africa since 1993 hover between 0.18 and 0.44. This suggests that union premia are on the "large" side.

Taken at face value the results for South Africa seem to support the idea that the union premium has increased since 1985 and then declined since around 2000. Nevertheless there are some concerns about the results for South Africa. These include that the results are based on a rather small subset of all the available labour market datasets, the samples, specifications and estimation methods are not always comparable, little attention has been paid to data quality issues and studies using individual fixed effects regressions suggest that part of the union premium may be due to unobserved individual characteristics. Remarkably none of the papers (not even the ones attempting to describe changes in the union premium) discuss changes in the nature of unionisation in South Africa.

In the other BRICS countries union premia seen to have been estimated less often. All the studies we reviewed used the single equation OLS model, finding premia usually between 5 and 15%, although one estimate from Brazil in the 1990s was around 20%. These are substantially smaller than in South Africa, indicating stronger union power than in the other BRICS countries.

3 Data

In our estimation of the union premium in South Africa we use version 3.1 of the Post-Apartheid Labour Market Series (PALMS) dataset (Kerr, Lam and Wittenberg 2016). This assembles the labour market information from 55 different datasets ranging from the 1993 PSLSD, the 1994 to 1999 October Household Surveys, the 2000-2007 biannual Labour Force Surveys and the Quarterly Labour Force Surveys from 2008 to 2015. We use most of these surveys, but the earnings information from 2015 had not yet been made available and was not collected/released in 2008 and 2009. Furthermore a question about union membership was not asked in 2008-Q2: 2010 so these periods are all excluded from our analysis. The sectoral, occupation and hours worked information was not comparable for 1993 so most of our analyses omit it. Information about who was in the public sector was not available in 1994 and 1995, so many of our analysis go from 1996 to 2014.

4 Methods

4.1 Dealing with data quality issues

The key issues that need to be confronted in using the PALMS data are dealing with the missing information (in particular earnings that are reported only in brackets), the outliers that contaminate information in several of the surveys, and the shifts induced by sampling changes and changes in the instrument. In this paper we use the approaches outlined in Wittenberg (2016b), viz.:

- **Outliers**

  Observations are flagged as outliers if they have a standardised residual that is larger than five in absolute value in a pooled Mincerian wage regression on the PALMS dataset. This eliminates most of the anomalous "millionaires" identified by Burger and Yu (2007), as well as some observations that seem to have anomalously low earnings.

- **Missing information and bracket responses**

  Missing earnings information or earnings that are reported in brackets are imputed multiple times (ten times overall) according to a hot deck procedure, involving matching on predicted values. Earnings data for 1996 (which exists only in brackets) is imputed by taking draws from the 1997 individual earnings distribution.

- **Sampling changes and changes in the instrument**

  One of the biggest discontinuities in the series occurs between October 1999 (the last OHS) and February 2000 (the first LFS). There are no less than three material changes: firstly, the OHS allowed individuals to report both wage income and income from self-employment activities, whereas the LFS only had one question dealing with income from one's primary activity; secondly, the LFS found many more informal and self-employment activities, in particular in agriculture (Neyens and Wittenberg 2016); thirdly, the 1999 OHS found too many high income earners, while the 2000:1 LFS found much too few (Wittenberg 2016a). Many of these changes (except for the last) are ameliorated by focusing on wage employees only, which is in any case appropriate if one is trying to measure the impact of unions.

Where we report the standard errors, these are robust standard errors correcting for clustering on enumerator area and also corrected for the multiple imputation of missing earnings, according to Rubin's rules (Wittenberg 2016b).

4.2 Union effects and causal effects

The key objective in estimating union premia is to simulate the hypothetical experiment in which a randomly chosen worker is made a member of a union. The Ruben causal model suggests that the relevant union effect (UE) would be

$$UE = E\left(Y_i^u - Y_i^n\right)$$

where Y_i^u is the level of (log) earnings of individual i if unionised and Y_i^n is the level if not. When the expectation is taken over the South African population this is obviously equivalent to the average treatment effect (ATE) i.e.

$$ATE = \frac{1}{N}\sum\left(Y_i^u - Y_i^n\right)$$

As is always the case we have a missing value problem since we only ever observe either Y_i^u or Y_i^n. If union membership was randomly assigned, then an unbiased estimate of the ATE would be the raw gap (RG), i.e. difference in average (log) earnings between the unionised and non-unionised

$$RG = \frac{1}{N_u} \sum_{U_i=1} Y_i^u - \frac{1}{N_n} \sum_{U_i=0} Y_i^n$$

where N_u is the number of unionised and N_n the number of non-union members. Writing the observed outcome as Y_i we have

$$Y_i = Y_i^n + U_i \left(Y_i^u - Y_i^n \right)$$

In this case the ATE could be estimated by the simple regression

$$Y_i = \beta_0 + \delta U_i + \varepsilon_i \tag{8}$$

The OLS specification in regression 1 can be thought of as an extension of this equation, with the additional controls $\mathbf{x}_i \boldsymbol{\beta}$. Implicit in this formulation is the assumption that

$$Y_i^n = \mathbf{x}_i \boldsymbol{\beta} + \varepsilon_i, \; Y_i^u = \mathbf{x}_i \boldsymbol{\beta} + \delta + \varepsilon_i$$

However it is clear that union membership is not randomly assigned and that the "controls" $\mathbf{x}_i$ are not mere nuisance variables. If these equations are read as Mincerian wage regressions, the assumption is that unions only affect the level of wage but not the returns (e.g. to education). This is unlikely to be the case.

Instead, we might think of writing the outcomes for each individual as

$$Y_i^n = \mathbf{x}_i \boldsymbol{\beta}^n + \varepsilon_i^n, \; Y_i^u = \mathbf{x}_i \boldsymbol{\beta}^u + \varepsilon_i^u$$

Obviously only one of these will be observed. If we estimate these two equations using the observed (log) earnings on unionised and non-unionised workers respectively, we get the two equation system given by equations 2 and 3. The Lewis procedure can therefore be thought of as a different way of estimating the ATE[3]. By definition this is

$$ATE = \frac{N_u}{N} \frac{1}{N_u} \sum_{U_i=1} \left(Y_i^u - Y_i^n \right) + \frac{N_n}{N} \frac{1}{N_u} \sum_{U_i=0} \left(Y_i^u - Y_i^n \right)$$

Substituting in the predicted values $\widehat{Y}_i^n$ for the unobserved Y_i^n when $U_i = 1$ and likewise the predicted $\widehat{Y}_i^u$ for Y_i^u when $U_i = 0$ will yield an estimate of the ATE.

Indeed Fortin et al. (2011) point out that the decomposition equation 5 consists of two terms: the former is a term controlling for the differences in characteristics, while the second can be thought of as the Average Treatment Effect on the Treated (ATT)

$$\widehat{ATT} = \overline{\mathbf{x}}^u \left(\widehat{\boldsymbol{\beta}}^u - \widehat{\boldsymbol{\beta}}^n \right) \tag{9}$$

So the standard Oaxaca-Blinder technique applied to the two equation model using the coefficient vector of the non-unionised as the base case can also be thought of as yielding a treatment effect. The key question is therefore under what circumstances these procedures will yield valid estimates.

Fortin et al. (2011, pp.13ff) discuss this in general. Indeed their application is precisely the estimation of union effects. They list the following conditions as being necessary:

[3] This is why it is misleading to label the single equation OLS estimate as the "treatment effects" estimate.

- Assumption 1: Mutually exclusive groups

  An individual is either a member of a union or not (not both and not neither). This obviously holds.

- Assumption 2: Structural form

  The individual earns according to wage structures which can be written as $m_U(\mathbf{x}_i, \varepsilon_i)$ if unionised or $m_N(\mathbf{x}_i, \varepsilon_i)$ if not. This is a very general characterisation, but the implication is that the differences in earnings between unionised and non-unionised workers can arise for only three reasons: i) differences in the wage functions m_U and m_N; ii) differences in the observed characteristics $\mathbf{x}_i$; and iii) differences in the unobserved characteristics ε_i (Fortin et al. 2011, p.16).

- Assumption 3: Simple counterfactual treatment (no general equilibrium effects)

  We assume that the counterfactual earnings function for unionised workers $m_U^C(.,.)$ is $m_N(.,.)$, i.e. we are assuming that if the individual was not unionised they would earn according to the wage structure of the non-unionised. This assumption rules out considering hypothetical counterfactuals like the earnings function if there were no unions at all, since this would most definitely change the distribution of earnings for everyone. Lewis (1986) repeatedly makes the point that at best one can estimate a union wage gap, i.e. the increment faced by a worker at the margin. With this assumption it may be possible to identify the ATT. In order to identify the ATE we would also need to assume that the counterfactual earnings function for the non-unionised $m_N^C(.,.)$ is $m_U(.,.)$. Note that this means that the estimation of Lewis wage gaps according to the method of section 2.2.2 requires stronger assumptions than the estimation of the ATT according to equation 9.

- Assumption 4: Overlapping support

  Let the set of all wage setting factors be Ω, then for all $(\mathbf{x},\varepsilon) \in \Omega$ we have $0 < \Pr(U = 1|\mathbf{x}, \varepsilon) < 1$, i.e. there are no characteristics (observed or unobserved) which guarantee that a worker is unionised or not unionised.

- Assumption 5: Conditional independence/ignorability

  For all $\mathbf{x}$, we have $U \perp\!\!\!\perp \varepsilon|\mathbf{x}$

  This is the strongest assumption and the one most likely to fail. It would be violated (for instance) if there is selection into union status (for a given set of observed characteristics) due to some of the unobserved characteristics. Given the fact that union membership is a choice variable it is unlikely to hold. If we had a good set of instruments we could aim to back out a "Local Average Treatment Effect" (see Fortin et al. 2011, p.25), but we don't have any plausible ones at present.

Our approach in this paper will therefore be more modest. We will not make major causal claims for the premia that we find. Nevertheless the discussion has been useful for crystallising what the objective of the estimations should be. Given that selection into unionism will occur when people anticipate wage gains, we think these will overestimate the true causal effects. However the values of the union premium over time should be informative about the direction in which the causal effect is changing, unless the selection processes are also changing over time.

4.3 Specification

We will parameterise the earnings functions $m_U(.,.)$ and $m_N(.,.)$ in standard Mincerian style. This raises the question as to which variables should be included. In particular we need to include the standard variables (education, a quadratic in age or experience, dummies for gender, race and location) but we also need to avoid the problem of "bad control" (Angrist and Pischke 2009, pp.64ff).

4.3.1 Bad controls

The discussion in Angrist and Pischke (2009) begins with the problem of trying to estimate the causal effect of education on earnings. It notes that occupation (blue collar vs white collar) would be a "bad control". The logic can be explained best by means of Figure 1. This is a causal graph as introduced by Pearl and discussed by Morgan and Winship (2007). The occupation variable is on the causal path from education to earnings. However we assume that some of the unobserved variables ε (e.g. ability and ambition) influence occupational choice as well as earnings (Y). In this situation, controlling for occupation will induce a correlation between education and ε. It is fairly easy to see that this has to be the case: individuals can get to high occupational status either by high education or by high ability (or both), so within the top occupations individuals with low educational qualifications must have higher ability, hence controlling for occupation induces a negative correlation between education and ability. This, however, contaminates the estimate of the impact of education on earnings. In the language of causal path diagrams, occupation is a "collider" (Morgan and Winship 2007, pp.64-67) and controlling for it opens a "back-door" path through ε which leads to inconsistent estimates of the causal impact of education.

Figure 1:

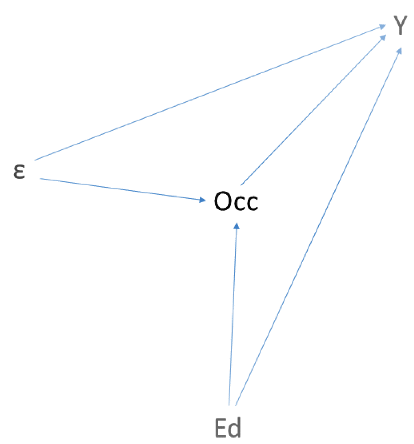

Notes:

A causal graph: occupation is a "bad control" since it is a "collider" in Pearl's terminology

While occupational dummies would be "bad controls" if we are trying to estimate the causal impact of education, we are interested in the causal impact of unions (with the caveats outlined in the previous section). The point of departure is therefore a causal diagram like the one in Figure 2. We assume that sector and occupation all influence the likelihood of joining a union, but earnings are also different by sector and occupation. This means that if we do not control for these variables, there will be many "back-door" paths from union status to earnings.

Figure 2:

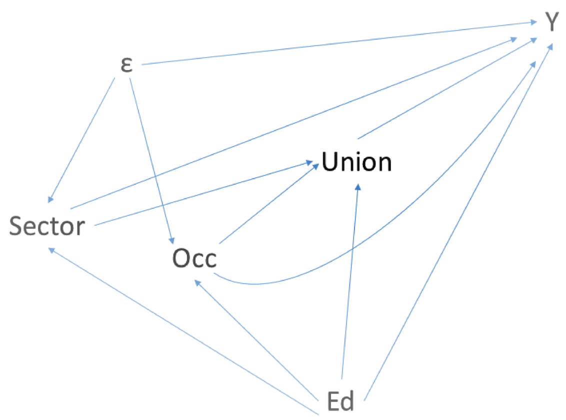

Notes:
Education, occupation and sector all open "back-door" paths to earnings and so must be controlled for if we want to estimate the impact of unions

"Bad controls" from the point of view of estimating union effects would be variables that are on the "downstream" causal path from union membership status to earnings. The job characteristics that Bhorat et al. (2012) introduce would seem to fit into this category. Whether or not a firm has medical aid or a pension plan is likely to be in part an **outcome** of union activity. Of course it is not completely determined by union pressure – but that makes it a "collider" in the language of causal graphs. Controlling for such variables will induce a correlation between union status and the unobserved characteristics of firms that may make them provide these benefits even in the absence of union activity.

4.3.2 Hours worked

One variable that might fit into this category that we nevertheless use as a control is the (log of) hours worked. This is a standard variable in Mincerian regressions. Furthermore the dependent variable of interest is the (log of the) wage rate, i.e. earnings per hour. Implicitly log hours is therefore on the right hand side of the regression (with a coefficient of one) if log total earnings is on the left hand side. Allowing the coefficient to be unconstrained permits the wage rate to change (e.g. due to overtime costs) with the hours worked.

There is a second way in which hours might feature in the analysis. The labour market can be thought of as buying and selling hours worked. From this perspective the key question of interest is not the impact of unions on the wage rate of the marginal **worker**, but the impact on the rate at which the marginal **hour of work** is sold. Empirically this means that the wage regressions need to be weighted according to the hours of work supplied.

4.4 Empirical strategy

Our empirical analysis proceeds in stages. Firstly we analyse the evolution of unionisation in South Africa. This will serve as broad backdrop for the empirical analysis that follows. In particular it transpires that changes in the public sector unionisation rate is the most striking development since the end of apartheid. Secondly we estimate the union wage premium. This proceeds in stages. We begin with the single equation approaches (estimating equation 1). We also explore the coefficients on key control variables in those Mincerian regressions and how they have evolved over the period. We then re-estimate the premia using the Lewis two-equation approach, finally doing so weighting by hours worked. Thirdly, we focus on the intersection of unionism and public sector employment.

5 Results

5.1 The evolution of unionisation

Figure 3 shows the evolution of the unionisation rate for employees using PALMS. Unionisation rates increased in the 1990s but reached the peak in 1999, after which they decreased from a maximum of around 37% to 30% by 2014. The level in the period 1997 to 1999 may be slightly overestimated, due to the better enumeration of more marginal forms of work in the Labour Force Surveys.

Figure 3:

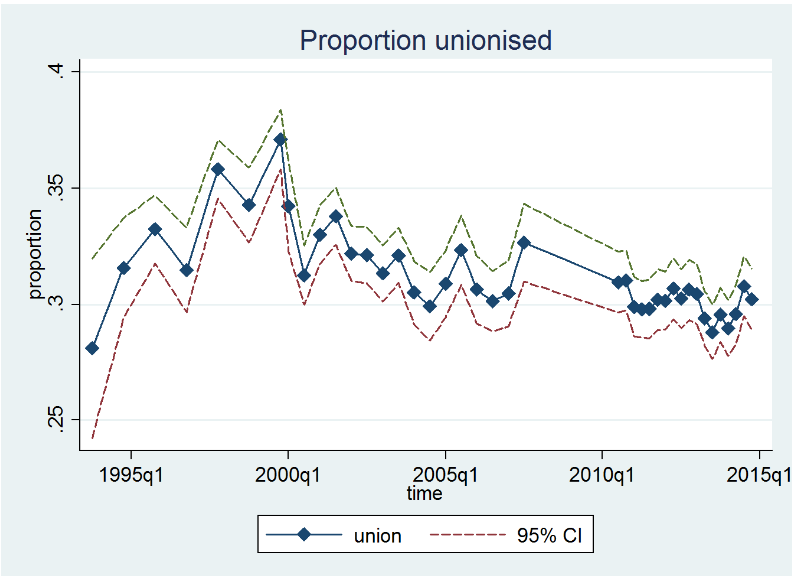

Notes: The proportion of union members has consistently been above 30% but seems to have dropped a little recently

Figure 4 shows unionisation rates by sector over the last 20 years. Unionisation rates are highest in mining, where more than 80% of workers belonged to unions by the end of 2014. Utilities and Services have unionisation rates that are above 50% by 2014 – these are both sectors with a high proportion of public employees. Utilities is a very small sector so the estimated unionisation rates vary quite a bit over the different surveys. Unionisation rates in both the utilities and services sectors have grown over the last 20 years. In contrast unionisation rates in manufacturing and transport have declined by about 10 percentage points over the last 20 years and were around 35% by the end of the period. Unionisation rates in the trade and finance industries were stable at around 20% over much of the period. The unionisation rate in construction seems to have declined and sat at around 15% by 2014, whilst the rate in agriculture was about 5%- the rate started the period at about the same level, increased but has subsequently decreased. The very low unionisation rates in domestic work reflect the difficulty in organising workers dispersed across their employers' homes.

Figure 4:

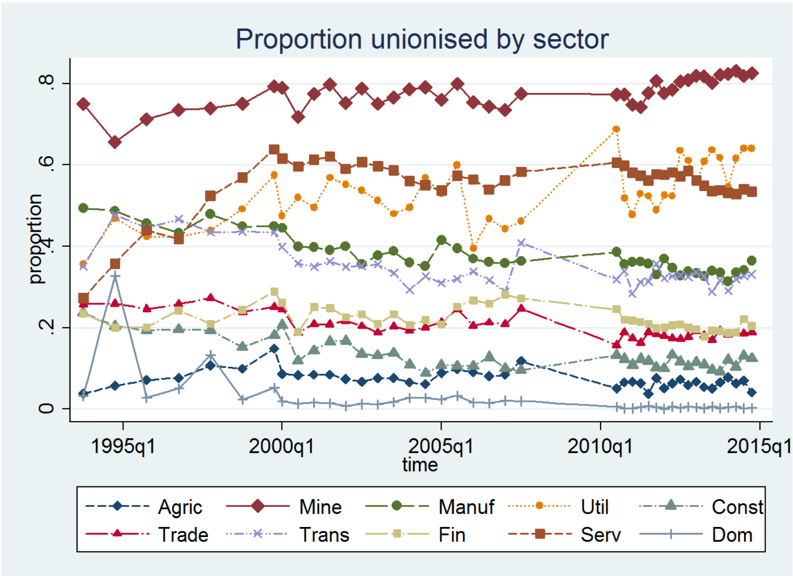

Notes: Unionisation rates in PALMS by one digit economic sector

Figure 5 shows the evolution of unionisation in the public and private sector. The public sector figures for 1996 to 1999 were calculated using the three digit sectoral codes, as discussed in Kerr and Wittenberg (2016). Figure 5 shows a massive increase in unionisation in the public sector. By contrast unionisation rates in the private sector seem to have come down somewhat. The step down after 1999, i.e. between the OHSs and the LFSs is probably again due to measurement changes more than a real drop in unionisation at that time. The outcome was that the proportion of union membership in the private sector dropped from 65% of total union membership in 1993 to only 50% in 2014.

5.2 Mincerian wage regressions: single equation estimates

We now turn to single equation estimates. The regressions include a union dummy, gender, a quadratic in age, years of education, population group dummies, province dummies, 1 digit industry code, 1 digit occupational classification, a dummy for employment in the public sector, log of hours worked and marital status. Figure 6 shows the evolution of four of the coefficients. The 95% confidence intervals are constructed taking into account the sample design of each survey as well as the multiply imputed nature of the data.

The pattern on the union premium corresponds in broad outline to the story told in Table 1 – it increases from the early 90s and seems to come down in more recent years. Nevertheless a closer examination of the coefficients shows that the pattern does not correspond to the pattern in the

Figure 5:

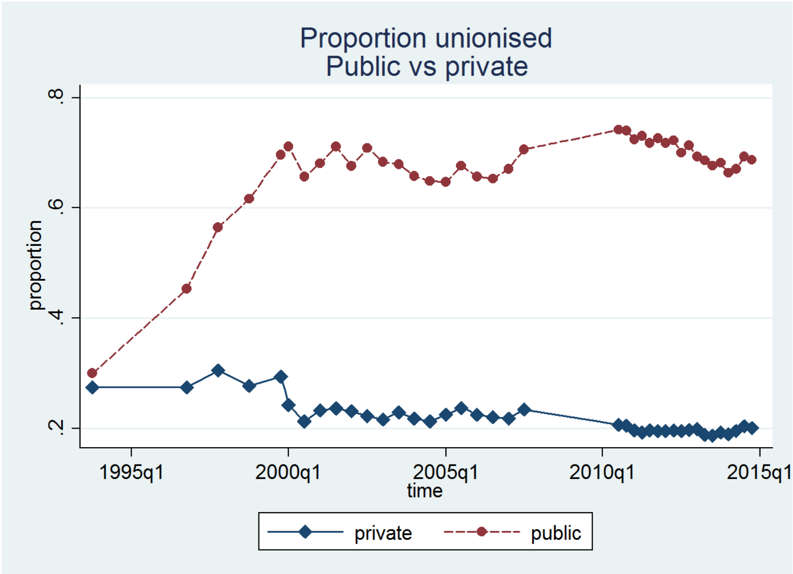

Notes: The public sector has become highly unionised since the end of apartheid

Table. In particular these estimates show an increase between September 2001 and September 2004. Furthermore the premium for September 2007 is abnormally low, so that the monotonic pattern shown by Ntuli and Kwenda (2014) is at least partially an artefact of the particular datasets used. Such anomalies appear more clearly once many more datasets are used.

The coefficients on some of the other variables as shown in Figure 6 are also interesting. The returns to education seem to have increased between 1993 and 2012 and then collapsed dramatically more recently. It is hard to believe that this could be true. Indeed there seem to be strange patterns in the earnings data in the most recent QLFSs, as noted in some of our other REDI working papers. Interestingly enough the premium earned by people classified as "White" has also dropped dramatically in the same period after being fairly stable at around 0.6 for much of the post-apartheid period. The wage penalty suffered by women has shown a more gradual erosion. Taken at face value these coefficients suggest that discrimination has decreased since 1993.

Figure 7 shows yet more coefficients to give us a fuller sense of how the processes of wage determination seem to have changed. The public sector premium seems to have decreased since 2010. The coefficient on being single when compared to the base case of being married shows no major change. It seems to hover in a band from -0.1 to -0.05 which gives a marital premium of somewhere between 55 and 10%. The penalty of being located in the Eastern Cape (when compared to the base case of being in the Western Cape) seems to have dropped over time. The most notable feature of this graph is the extremely noisy estimate in 1993, suggesting a problem with the location variables in that survey. The gap between managers and unskilled workers seems to have widened considerably over the period.

5.3 Two equation estimates

We now turn to estimating the union gap by means of the two equation system outlined in section 2.2.2. First we consider the evolution of some of the key coefficients in these regressions,

Figure 6:

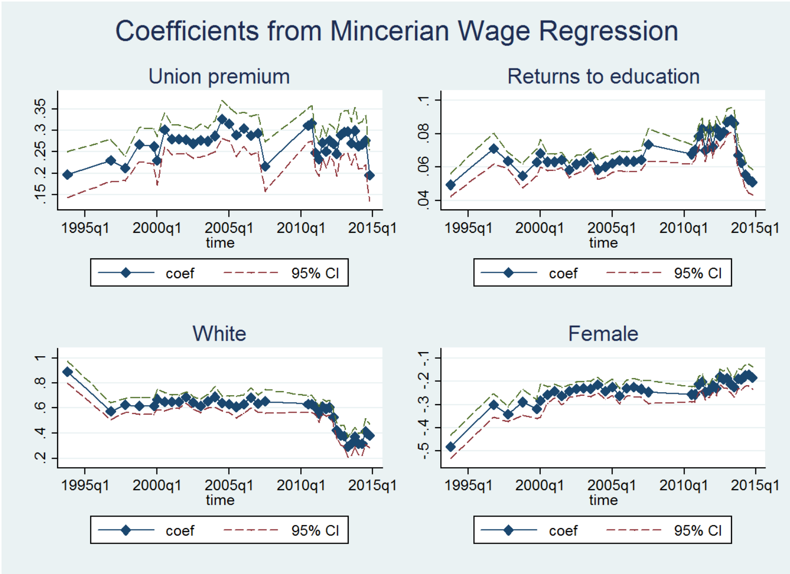

Notes: Coefficients from single equation regressions. Additional controls include a quadratic in age, province, industry, occupation, public sector, hours worked and marital status.

as shown in Figure 8. Each coefficient is extracted from a regression which has the same controls as in the single equation case, except that the equations are estimated separately for union and for non-union members.

The coefficients again show an apparent collapse in returns to education post 2012, for both union and non-union samples. Since 2007 it would appear that average returns to education have been higher in the union sector. The premium on earnings for White workers seems to have been reduced more quickly in the unionised sector, in line with the earlier findings of Azam and Rospabé (2007). The difference between male and female wages seems lower in the union sector, although it seems to have come down in both. Casale and Posel's (2010) finding that the union premium is actually bigger among women workers is not directly comparable to this result, since we are estimating separate union and non-union regressions with a simple gender dummy, whereas Casale and Posel (2010) estimate the union effects from separate regressions for men and women.

Perhaps the most interesting series of coefficients in Figure 8 is that for the public sector premium. It suggests that the premium has collapsed among non-unionised workers, while it has remained steady for unionised ones. Given the high rate of unionisation in the public sector it is likely that non-unionised public servants are somewhat unusual. It also raises the question discussed by Kerr and Wittenberg (2016) about the nature of the wage imputations for public sector workers in the QLFS.

Figure 9 provides additional information on what, if anything, is gained by estimating regressions separately by union status. The marital premium is very similar in both regressions, but interestingly the difference between Eastern Cape and Western Cape wages looks very different in the union and the non-union sector. This is most likely due to the fact that public sector wages are negotiated nationally. Nurses, teachers and policemen will be remunerated according to national scales and not so much local labour market conditions. The difference in log wages between managers and unskilled workers is a bit larger in the non-unionised sector, but the

Figure 7:

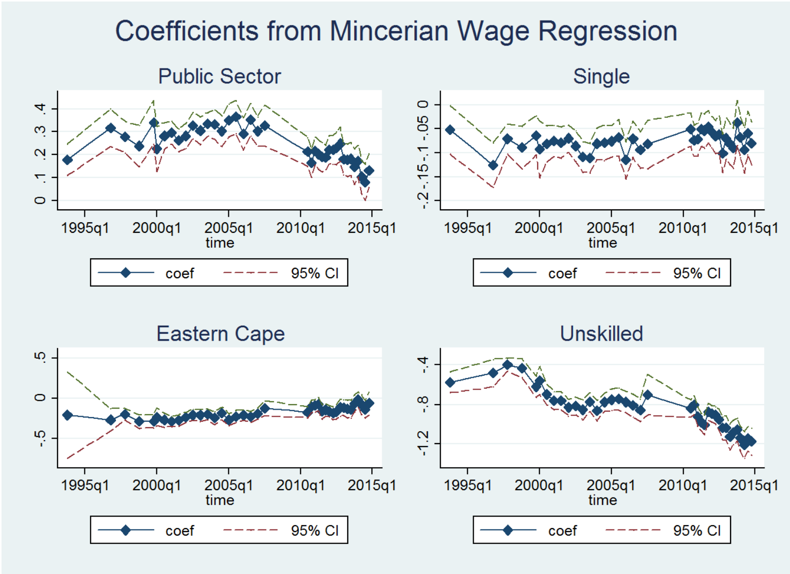

Notes: Additional coefficients from the same Mincerian regressions as in Figure 6

two sets of coefficients track each other quite closely. Over the twenty year period the gap has widened considerably.

The coefficients on log hours are curious. The dependent variable is the log of total earnings, so the expected coefficient would have been one. The fact that the coefficient on hours in the union sector is around zero suggests that remuneration is decoupled from hours worked, which seems astonishing. The pattern in the non-unionised sector shows some responsiveness, increasing over time, but with coefficients in a range which suggest that individuals that work longer hours actually have a lower wage-rate.

We turn now to consider what the two equation estimation approach suggests about the evolution of the union premium. Figure 10 presents the outcome. The line labelled "lewis" represents the estimate according to the procedure outlined in section 2.2.2. As indicated in that section, the Lewis estimate can be thought of as the average of two price gaps from Oaxaca-Blinder decompositions: one taking the non-union wage as the default and averaging the gap according to the characteristics of the unionised sample (labelled "union" in the Figure); and the other taking the union wage as the default and averaging the gap over the characteristics of the non-union sample (labelled "non-union" in the Figure). We noted in section 4.2 that the former can be thought of as the Average Treatment Effect on the Treated. In this case it would be the union gap for the unionised.

The Lewis union wage gap in Figure 10 behaves similarly to the union premium as estimated by single equation methods (as shown more clearly in Figure 11), increasing until it reaches a maximum around 2004 and then collapsing more recently. Interestingly, however, the "ATT" version of the union effect behaves differently, reaching a maximum in the post 2010 period. More troublingly, however, the three different estimates presented in Figure 10 are reasonably close together for much of the post-apartheid period, but diverge dramatically recently.

Figure 8:

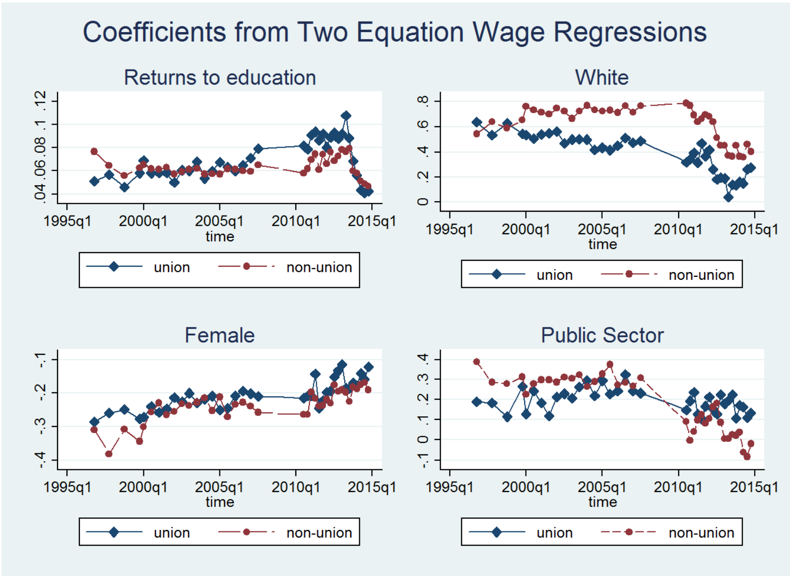

Notes: Coefficients from separate regressions for union and non-union members. Other control variables include a quadratic in age, province, industry, occupation, hours worked and marital status.

5.4 Weighting by hours worked

The final estimate of the union gap is provided by the two equation method, but weighting both the regressions and the calculations of the means by the hours of work supplied to the labour market, so that the premium is conceptually attached to an hour worked rather than the person who supplied it. As Figure 11 shows, this doesn't make a big difference to the trends observed. Indeed all three methods provide very similar estimates, except post 2012. All three estimates suggest that September 2007 and the last quarter of 2014 were strange, to say the least.

5.5 The public sector and unionism

At various stages we have observed that the high rate of unionism in the public sector may influence some of the results that we observe. Indeed it is likely that unions would operate very differently in the private and public sectors. In Figure 12 we estimate the union premium separately for the private and public sectors. The results are startling. At the beginning of the period it would appear that the union premium was higher in the private sector, but by the mid-2000s this had been reversed. The private sector premium appears to have increased a little since the early 1990s but has has come down again and seems to be smaller now than it has ever been. By contrast the union premium in the public sector seems to have ballooned. The decisive shift upwards seems to have happened post 1997, which was when the Public Service Co-ordinating Bargaining Council was introduced. Given the nature of centralised bargaining it is a bit puzzling why the union premium seems to be as hefty as it is, given that non-union members would be covered by the agreement. However in many sections of central government union dues are deducted by default and employees have to opt out (personal communication, Ingrid Woolard). In this environment it is clear that non-union members are likely to be anomalous – perhaps contract or other temporary workers. Alternatively they might be working in less

Figure 9:

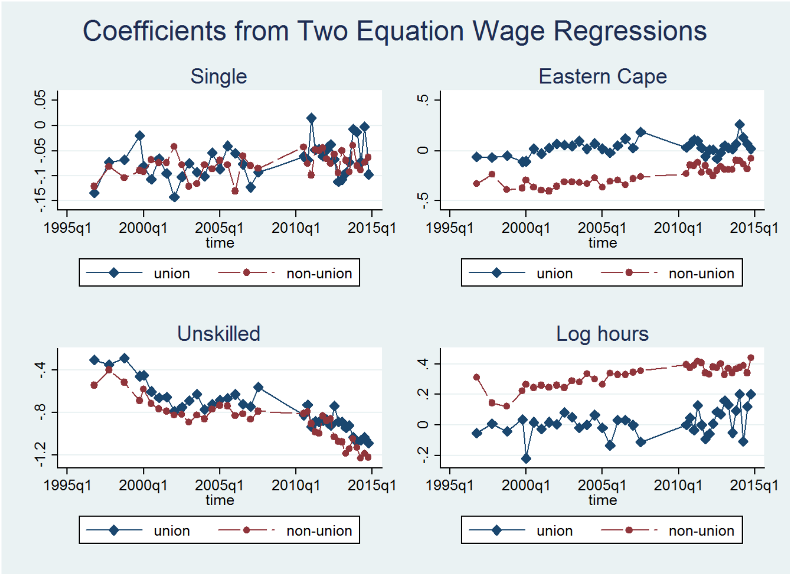

Notes: Additional coefficients from regressions estimated separately for union and non-union members. The regressions are identical to those underlying Figure 8.

"central" parts of the bureaucracy, e.g. parastatals or municipalities.

The large union premium estimated for the public sector should therefore be cautiously interpreted. It is likely to reflect both the very centralised nature of bargaining, the very different nature of the employer (one that is in formal alliance with the unions and that is not constrained by profit considerations) as well as more mundane measurement issues, such as the extent to which particularly contract workers are able to accurately identify their employer. Nevertheless it is clear that the diverging trends shown in Figure 12 are perhaps the central drivers of the union premium.

6 Discussion

Our empirical analysis yields some simple take-home messages. First, unionisation has held steady at around 30% of the labour force, but this conceals a major shift towards public sector unionism and away from private sector unions. Second, different methods of estimating the union wage premium all yield broadly similar results, except methods that employ Heckman style sample selection corrections. Third, data quality issues are important, specifically the anomalous looking September 2007 LFS, and the "collapses" in the returns to education, the racial wage premium as well as changes in the union premium post 2012. Forth, our estimates suggest that the union premium increased from the early '90s to the mid 2000s and has since come down, although the "Average Treatment Effect on the Treated" version of the union gap remains high, but these aggregate trends conceal a divergence in the evolution of the union premium in the public and private sectors. Indeed it appears that the interaction between public sector bargaining and unionism is the key driver of changes in the union premium. We elaborate a little on some of these points in the following section.

Figure 10:

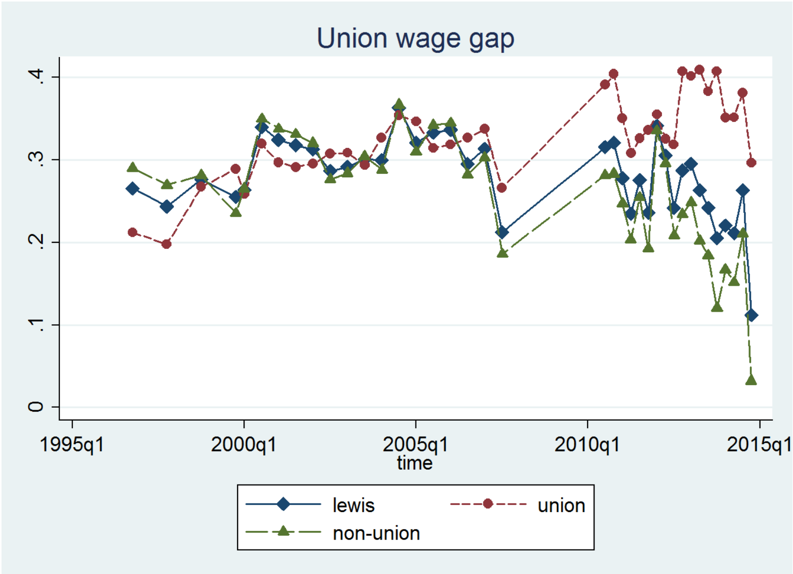

Notes: The union premium estimated according to the two equation method

6.1 Data quality issues and imputations in the QLFS

The shifts in the post 2012 QLFS data seem troubling, since it is difficult to believe that these are "real" shifts. Indeed the problem with th eentire QLFS earnings series is that the publicly relased version of the data has been fully imputed, so that it is difficult to treat these surveys in the same way as the earlier ones. Kerr and Wittenberg (2016) show that the imputation process seems to produce problems in the context of estimating the public sector premium.

More generally there are potentially problems when an imputation process does not use an important variable. In this case it is uncertain whether Statistics South Africa used union status in the imputation. If it is not used, then the implication is that the difference between union and non-union earnings would appear smaller on the imputed data than it should be. This point applies with equal force to the multiply imputed data in PALMS (relevant mainly to the OHS and LFS data). The imputation routine used by DataFirst also omitted union status as a variable, so it may be the case that we are underestimating the union premium in this paper.

Even though there may be questions about the DataFirst imputation algorithm, it has the chief virtue that it is completely transparent and could potentially be "fixed". At present it is impossible to know precisely what happened in the QLFS imputations. Looking at the output it seems clear that the process produces data that look anomalous post 2012 even when compared to the earlier QLFSs.

6.2 The union premium and the public sector

The size of the "Averate Treatment Effect on the Treated" estimate has increased since the '90s and remains high. This means simply that individuals who are unionised would earn considerably less if they were paid like non-unionised workers. Given the fact that the unionised worker is now much more likely to be in the public sector this is arguably reflective of changing wage determination processes within the public sector. Our analysis suggests that the creation of the

Figure 11:

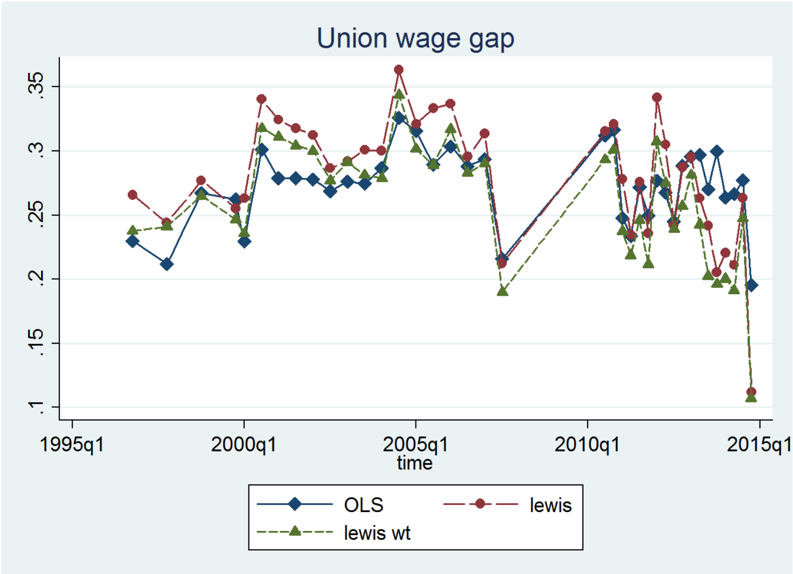

Notes: The evolution of the union wage gap according to three different methods

Public Service Co-ordinating Bargaining Council together with the strong unionisation of the public sector has been key in affecting the aggregate union premium.

7 Conclusion

In this paper we have estimated the South Africa union premium using nearly 25 years of household survey data from the post-Apartheid period. Our estimates suggest that this premium is on the "large" side, particularly when we consider the premium for the unionised (ATT). It seems especially large when compared to union premia estimated for the other BRICS countries, which generally range from 5% to 20%.

It is clear that the size of the union premium in South Africa is heavily affected by the particular nature of bargaining within the public sector. Outside this sector the premium seems to have come down since the mid-2000s. While the literature has been very concerned about the government extending centrally bargained agreements to other parties, it appears that the way in which the government has bound itself (and its employees) through the Public Service Co-ordinating Bargaining Council appears to have a massive effect on the coefficients estimated through Mincerian wage regressions. Thinking carefully about how to take account of these different processes in the econometrics seems important future work.

Our analysis has shown why it is very useful to use as many surveys as possible in the estimation. Anomalies become much more visible when looked at in context. An example is the QLFS post-2012 data, which suggests that imputations undertaken by Statistics South Africa hamper the estimation of unbiased union premia. Part of the problem in empirical analysis is to separate out the signal from the noise. Although we cannot be sure we have found all studies estimating the union premium for the other BRICS countries, it seems that studies that do exist almost all rely on only one or two surveys. Our analysis for South Africa clearly shows that relying on too few surveys invites one to become hostage to the noise.

Figure 12:

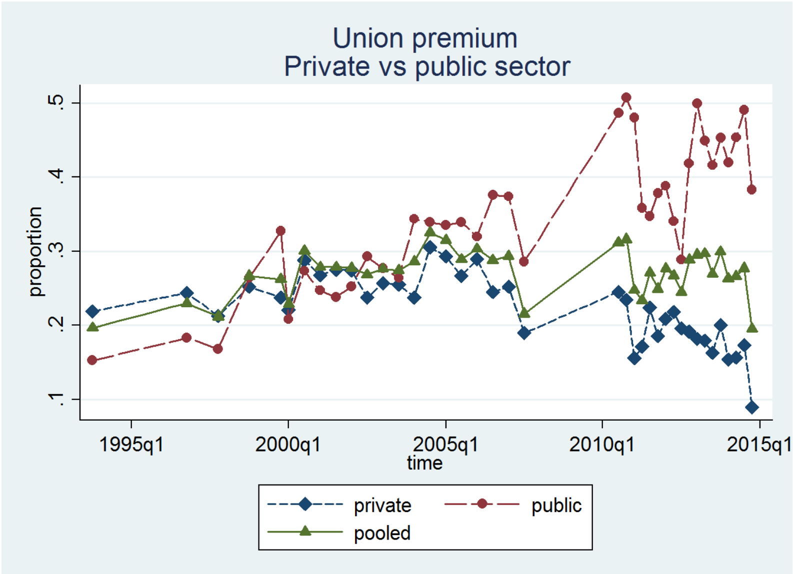

Notes: The union premium in the private and public sectors, as estimated by the single equation method. The additional variables in the regression are as in all the other cases.

References

Angrist, J. D. and Pischke, J.-S.: 2009, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, Princeton, NJ.

Arbache, J. S.: 1998, Unions And Wage Inequality In Brazil, *Technical report*.

Azam, J.-P. and Rospabé, S.: 2007, Trade unions vs. statistical discrimination: Theory and application to post-Apartheid South Africa, *Journal of Development Economics* **84**, 417–444.

Banerjee, A., Galiani, S., Levinsohn, J., McLaren, Z. and Woolard, I.: 2008, Why has unemployment risen in the new South Africa?, *Economics of Transition* **16**(4), 715–740.

Bhandari, A. K.: 2010, Does Union Membership Pay Off? Evidence from Organized Indian Manufacturing Industries, *The Indian Journal of Industrial Relations* **45**(3), 459–469.

Bhorat, H., Goga, S. and Van Der Westhuizen, C.: 2012, Institutional wage effects: Revisiting union and bargaining council wage premia in South Africa, *South African Journal of Economics* **80**(3), 400–414.

Boito, A.: 1994, The State and Trade Unionism in Brazil, *Latin American Perspectives* **21**(1), 7–23.

Boito, A. and Marcelino, P.: 2011, Decline in Unionism? An Analysis of the New Wave of Strikes in Brazil, *Latin American Perspectives* **38**(5), 62–73.

Burger, R. and Yu, D.: 2007, Wage trends in post-Apartheid South Africa: Constructing an earnings series from household survey data, *Working Paper 07/117*, Development Policy Research Unit, University of Cape Town.

Butcher, K. and Rouse, C.: 2001, Wage effects of unions and industrial councils in South Africa, *Industrial and Labor Relations Review* **54**(2), 349–374.

Casale, D. and Posel, D.: 2010, Unions and the gender wage gap in South Africa, *Journal of African Economies* **20**(1), 27–59.

Clarke, S.: 2002, Market And Institutional Determinants Of Wage Differentiation In Russia, *Industrial and Labor Relations Review* **55**(4), 628–648.

Clarke, S.: 2005, Post-socialist trade unions: China and Russia, *Industrial Relations Journal* **36**(1), 2–18.

Das, S.: 2008, Trade Unions In India: Union Membership And Union Density, *The Indian Journal of Labour Economics* **51**(4), 969–982.

de Oliveira Cruz, B. and Naticchioni, P.: 2012, Falling urban wage premium and inequality trends: evidence for Brazil, *Investigaciones Regionales - Journal of Regional Research* (24), 91–113.

Fortin, N., Lemieux, T. and Firpo, S.: 2011, Decomposition methods in economics, *in* D. Card and O. Ashenfelter (eds), *Handbook of Labor Economics*, Vol. 4a, Elsevier, Amsterdam, pp. 1–102.

Hofmeyr, J. and Lucas, R. E.: 2001, The rise in union wage premiums in South Africa, *Labour* **15**(4), 685–719.

Kerr, A., Lam, D. and Wittenberg, M.: 2016, Post-Apartheid Labour Market Series [dataset], DataFirst, University of Cape Town. Version 3.1.

Kerr, A. and Teal, F.: 2015, The determinants of earnings inequalities: Panel data evidence from KwaZulu-Natal, South Africa, *Journal of African Economies* **24**(4), 530–558.

Kerr, A. and Wittenberg, M.: 2016, Public sector wages and employment in South Africa, *unpublished working paper*, REDI.

Lee, L.-F.: 1978, Unionism and wage rates: A simultaneous equations model with qualitative and limited dependent variables, *International Economic Review* **19**(2), 415–433.

Lehmann, H. and Muravyev, A.: 2009, How Important Are Labor Market Institutions for Labor Market Performance in Transition Countries?, *IZA Discussion Papers 4673*, Institute of Labor Economics (IZA).

Lewis, H. G.: 1986, *Union relative wage effects: A survey*, University of Chicago Press, Chicago.

Lukiyanova, A.: 2011, Effects Of Minimum Wages On The Russian Wage Distribution, *Basic research program working papers*, National Research University Higher School of Economics (HSE).

Mishra, V. and Smythe, R.: 2014, Technological Change and Wages in China: Evidence from Matched Employer–Employee Data, *Review of Development Economics,* **18**(1), 123–138.

Moll, P.: 1993, Black South African unions: Relative wage effects in international perspective, *Industrial and Labor Relations Review* **46**(2), 245–261.

Moll, P.: 1996, Compulsory centralization of collective bargaining in South Africa, *American Economic Review* **86**(2), 326–329.

Morgan, S. L. and Winship, C.: 2007, *Counterfactuals and Causal Inference: Methods and Principles for Social Research*, Cambridge University Press, New York.

Nattrass, N. and Seekings, J.: 2014, Job destruction in Newcastle: minimum wage-setting and low-wage employment in the South African clothing industry, *Transformation: Critical Perspectives on Southern Africa* **84**, 1–30.

Neyens, E. and Wittenberg, M.: 2016, Changes in self-employment in the agricultural sector, South Africa: 1994-2012, Working Paper, DataFirst University of Cape Town.

Ntuli, M. and Kwenda, P.: 2014, Labour unions and wage inequality among African men in South Africa, *Development Southern Africa* **31**(2), 322–346.

Rospabé, S.: 2002, How did labour market racial discrimination evolve after the end of Apartheid?, *South African Journal of Economics* **70**(1), 185–217.

Roychowdhury, S.: 2011, Public Sector Restructuring and Democracy: The State, Labour and Trade Unions in India, *International Labour Review* **150**(1-2), 127–143.

Schultz, T. P. and Mwabu, G.: 1998, Labor unions and the distribution of wages and employment in South Africa, *Industrial and Labor Relations Review* **51**(4), 680–703.

Wittenberg, M.: 2014, Union selection effects – some inconsistent models, *Working Paper 133*, SALDRU.

Wittenberg, M.: 2016a, The top tail of South Africa's earnings distribution: Evidence from the Pareto distribution, *unpublished working paper*, REDI.

Wittenberg, M.: 2016b, Wages and wage inquality in South Africa 1994-2011: Part 1 – wage measurement and trends, *South African Journal of Economics* **in press**. http://dx.doi.org/10.1111/saje.12148.

Wittenberg, M.: 2016c, Wages and wage inquality in South Africa 1994-2011: Part 2 – inequality measurement and trends, *South African Journal of Economics* **in press**. http://dx.doi.org/10.1111/saje.12147.

Zhu, Y., Warner, M. and Feng, T.: 2011, Employment relations "with Chinese characteristics": The role of trade unions in China, *International Labour Review* **150**(1-2), 127–143.

Table 1: Union wage premia estimated in the literature

Study	Data	Sample	Estimation	Controls	Effect[a]
Moll (1993)	BMR '85	Urban black blue collar	OLS, Lewis, Lee	Education, Potential Experience, Tenure, Skill, Sector, Location, Marital Status, Race, Language, Gender	0.095[b]
	CSS '85	White male full-time skilled/semi-skilled	"	"	0.113[b]
Moll (1996)	OHS '94	?	?	?	0.11,0.23
Schultz & Mwabu (1998)	PSLSD '93	African men	OLS, Quantile Reg	Education spline, Potential experience, industry, region	0.191
		White men	"	"	-0.097
Butcher & Rouse (2001)	OHS '95	Africans	OLS	Female, Head of Household, Age, Married, Occupation X Education, Industry, Province	0.184
	"	Whites	"	"	0.096
Hofmeyr & Lucas (2001)	BMR '85	African urban male	OLS, Lewis, Lee	Education, Potential Experience, Marital Status, Sector, Occupations, Location	0.064
	PSLSD '93	African urban male	"	"	0.198
Rospabé (2002)	OHS '99	African males	Interval Reg	Schooling spline, Potential Experience, Tenure, Marital Status, Urban, Skill, Sector, Province	0.253
	"	White males	"	"	-0.011
Azam & Rospabé (2007)	OHS '99	African males	Heckman, Lee	Schooling, Age, Tenure, Marital Status, Occupation Level, Sector, Formal, Urban, Province	0.6954
	"	White male	"	"	0.0294
Banerjee *et al* (2008)	OHS '95	All males	OLS, Matching	Education, Race, Age, Occupation, Sector, Province, Urban	0.19
	OHS '98	"	"	"	0.27
	LFS '00:2[c]	"	"	"	0.38
	LFS '04:1[c]	"	"	"	0.35

Table 1: Union wage premia estimated in the literature

Study	Data	Sample	Estimation	Controls	Effect[a]
Casale & Posel (2010)	LFS '03:2	African women	OLS, Quantile Reg, Heckman, Lee	Education, Age, Tenure, Occupation, Sector, Public Sector, Formal, Firm Size	0.216
	"	African men	"	"	0.183
	LFS panel '01-'04	African women	Panel Fixed effects	"	0.083
	"	African men	"	"	0.082
Bhorat *et al* (2012)	LFS '05:2	African	OLS, Heckman	Gender, Education Spline, Experience, Marital Status, Head of Household, Occupation, Sector, Hours, Metro, Province, Self-Employed	0.1967
	"	"	"	" + non-wage benefits	0.0561
Ntuli & Kwenda (2014)	LFS '01:2	Formal Non-Agric, African men	Lee	Education, Age, Marital Status, Tenure, Firm Size, Sector, Occupation, Urban, Province	0.4492[b]
	LFS '04:2	"	"	"	0.4324[b]
	LFS '07:2	"	"	"	0.3893[b]
	QLFS '10:3	"	"	"	0.3535[b]
Kerr & Teal (2015)	KIDS '93, '98, '04	African KZN	pooled OLS, panel fixed effects	Education, Age, Marital Status, Occupation, Sector, Public Sector	0.276[d] 0.213[e]

Notes:
a. The OLS estimates are given, where they are available.
b. Percentage difference p converted to log wage difference as $ln(p+1)$
c. Only survey year given in regression table, but this survey was used elsewhere in the paper
d. Private union coefficient e. Difference between public union and public non-union coefficients

Sources:
Moll (1993) Table 5, p.256, last column and Table 9, p.259; Moll (1996) p.328;
Schultz & Mwabu (1998), Table 5 and Table 6, pp.690-1;
Butcher & Rouse (2001), Table 2a col.5 and Table 2b col.5, pp.359-60;
Hofmeyr & Lucas (2001), Table 5, p.709; Rospabé (2002), Table 3, p.205;
Azam & Rospabé (2007) Table 2A, p.439; Banerjee *et al* (2008), Table 9, p.728;
Casale & Posel (2010) Table 8, p.54; Bhorat *et al* (2012), Table 2 cols3-4, p.408;
Ntuli & Kwenda (2014) Table 6, p.338; Kerr & Teal (2015) Table 7, col 5, p.541

The **Research Project on Employment, Income Distribution and Inclusive Growth (REDI3x3)** is a multi-year collaborative national research initiative. The project seeks to address South Africa's unemployment, inequality and poverty challenges.

It is aimed at deepening understanding of the dynamics of employment, incomes and economic growth trends, in particular by focusing on the interconnections between these three areas.

The project is designed to promote dialogue across disciplines and paradigms and to forge a stronger engagement between research and policy making. By generating an independent, rich and nuanced knowledge base and expert network, it intends to contribute to integrated and consistent policies and development strategies that will address these three critical problem areas effectively.

Collaboration with researchers at universities and research entities and fostering engagement between researchers and policymakers are key objectives of the initiative.

The project is based at SALDRU at the University of Cape Town and supported by the National Treasury.

Tel: (021) 650-5715

www.REDI3x3.org

REDI3X3

Research Project on Employment, Income Distribution & Inclusive Growth

Towards an integrated response to unemployment, inequality and poverty