



Modelling Economic Policy Issues

Utility of inequality sensitive measures of the gender wage gap: Evidence from South Africa

Adeola Oyenubi^{a,*}, Jacqueline Mosomi^b^a School of Economics and Finance, University of the Witwatersrand, South Africa^b SALDRU, University of Cape Town, South Africa

ARTICLE INFO

JEL classification:

J01
J31
J71
I31
C43

Keywords:

Gender wage gap
Metric entropy distance
Stochastic dominance
Wage distributions
Counterfactual analysis
Inequality, Labour markets
South Africa

ABSTRACT

We examine the trend in the gender wage gap in South Africa taking into consideration the high level of general inequality in the country. To do this, we utilize the information-theoretic approach of the Generalised Entropy (GE) measures of divergence supplemented by stochastic dominance analysis and tests. Under this approach, male and female wage distributions are summarized by suitable evaluation functions and then the difference between the evaluations is computed. Unlike conventional measures of the gender wage gap based on the mean and quantiles, this approach does not assume rank invariance and makes explicit the evaluation functions used. We find that the rate of convergence of the gender wage gap depends on the measure used to aggregate the gap across the entire distribution. Comparison of the aggregate measures of the gender wage gap (mean, median and the entropy measure) reveals that the convergence of the gender wage gap is faster under the entropy measure. This is because unlike the mean, the entropy measure captures the dispersion of wages across individuals. Stochastic dominance test results reveal that male wages dominate female wages mostly in the second-order sense i.e., only for evaluation functions that are increasing in wages and concave. Finally, decomposition results show that across measures, the persistent gender wage gap is attributable to wage structure effects as opposed to composition effects. Therefore, to effectively address the persistent gender wage gap in the middle of the income distribution, policies should be geared towards the returns to characteristics at the median. The implication of these findings is that measures of the gender wage gap that ignore the overall inequality are likely to skew the estimate of the gender wage gap.

1. Introduction

The measurement of the gender wage gap and its trend over time is one of the measures of progress towards a gender egalitarian society. Gender equality is not only a sustainable development goal but equal participation and remuneration for men and women contributes to accelerated economic growth. The trend in the gender wage gap allows policy makers and researchers to evaluate the impact of various laws and policies enacted to reduce gender discrimination.

International literature on the gender wage gap finds that the wage structure is extremely important in explaining why the trend in the gender wage gap differs between countries (Blau and Kahn, 1996; Fortin and Lemieux, 2000; Blau, 2012; Carrillo et al. 2014). This

* Corresponding author.

E-mail address: adeola.oyenubi@wits.ac.za (A. Oyenubi).

stems from the argument that the labour market remains segregated by gender with women being overrepresented at the bottom and underrepresented at the top of the wage distribution. Gendered occupational segregation, part-time work, and the need for flexible hours due to the disproportionate unpaid care work shouldered by women are some of the factors attributed to the persistent gender differences in wages (Goldin, 2014; Blau and Kahn, 2017). In countries where the wage distribution is more widespread, gender wage inequality is on average higher and especially worse at the bottom of the wage distribution (Carrilo et al. 2014). Using microdata from several industrialised countries, Blau (2012) reports that the higher gender wage gap in the United States of America (USA) compared to Sweden and Australia was due to higher wage inequality in the country. By comparing welfare states to other states, Mandel and Shalev (2009) show that the gender wage gap is both a function of where women fall in class hierarchies and within and between class inequality in a country.

Many studies have analysed the gender wage gap in South Africa (Casale, 2004; Ntuli, 2007; Shepherd, 2008; Muller, 2009; Bhorat and Goga, 2013; Mosomi, 2019a; Casale et al., 2021). The consensus from these studies is that the gender wage gap persists in the South African labour market and while the gap was previously wider at the bottom of the distribution, the narrowing of the gap has also been fastest in this part of the distribution and at the mean with the median gender wage gap remaining stubbornly high and stagnant in the post-apartheid period (Mosomi, 2019a). Also, the gender wage gap is attributable to wage structure effects and not compositional factors e.g., education given that women in the South African labour market are on average more educated than men (Mosomi, 2019a). Methodologically, these studies mainly utilize mean or quantile decompositions which we refer to as conventional methods. Conventional methods calculate the gender wage gap (the difference between male and female wage distributions) by computing the gap at each quantile and then summarizing the result using a suitable evaluation function. The evaluation functions either impose equal weight on all observations (e.g., mean) or apply zero weight to some observations (e.g., 10th percentile, median, 90th percentile).

A critique of these conventional approaches is that the weighting schemes do not adequately capture the dispersion in distributions especially in high inequality country contexts (Maasoumi and Wang, 2019). Further, conventional measures (mean and quantiles) implicitly assume rank invariance i.e., a woman's relative rank is preserved when endowed with a man's skills set or market returns. As noted by Maasoumi and Wang (2019), this assumption might not hold in practice. Therefore, while information about the trend in the gender wage gap at the mean and at different quantiles provides useful information on gender inequality, it is still incomplete information if overall wage inequality is not considered. Also note that when rank invariance is not satisfied, it is possible that differences in outcome distributions can be biased due to individuals moving up and down the distribution (Dong and Shen, 2018).

Given this shortcoming of the conventional measures of the gender wage gap and South Africa's high inequality, we utilize the information-theoretic approach of the Generalised Entropy (GE) proposed by Maasoumi and Wang (2019) to analyse the gender wage gap in South Africa. Unlike conventional methods, the proposed information-theoretic framework first summarizes each wage distribution using an appropriate evaluation function and then computes the difference between the evaluations. Following Maasoumi and Wang (2019) we use the entropy measure and stochastic dominance tests to characterize the trend in gender wage gap in the South African context. This approach has several advantages in that, first, it is more transparent because the choice of evaluation function is made explicit. Second, it does not assume rank invariance. Lastly, the approach accommodates a situation where one is interested in comparing distributions over a large class of evaluation functions using statistical tests of stochastic dominance rankings.

We carry out the analysis using the Post-Apartheid Labour Market Series (PALMS) dataset (Kerr et al., 2017). PALMS is a stacked dataset comprising of nationally representative cross-sectional labour force surveys conducted in South Africa since 1993. The aim of this analysis is to explore summary measures of the gender wage gap that provide information about the general trend in a way that is sensitive to overall inequality.

South Africa is an interesting case for this analysis for several reasons. First, South Africa is one of the most unequal¹ countries in the world. Moreover, labour market income contributes over 70 % to this high inequality (StatsSA 2019) therefore it would be difficult to analyse gender inequality in the labour market without considering overall wage inequality. With increasing literature on the rise of income inequality driven by top income earners (Piketty and Saez, 2013; Fortin et al., 2017; Chatterjee et al., 2020), it becomes crucial to consider the influence of overall inequality on the gender wage gap. In addition, and specific to South Africa, discriminatory laws implemented during apartheid² make the labour market highly segregated by race and gender. On the one hand, Black men and women tend to be concentrated in elementary occupations and domestic work while White women tend to be concentrated in professional and clerking occupations. White men tend to be concentrated in managerial and professional occupations. Because of this intersectionality, women in the South African labour market are not homogenous that is, while most Black women are concentrated in the lowest paying jobs, there are some women in well-paying professional jobs explaining the heterogeneity of the gender wage gap. The entropy measure will therefore allow us to understand the trend in the gender wage gap in a way that is sensitive to the underlying level of inequality thus allowing us to say something about progress in terms of gender wage inequality.

Second, to try and reverse the legacy of apartheid, the South African government enacted anti-discrimination laws to promote race and gender equality. For example, aiming to promote equal opportunity and treatment in the labour market, the government enacted the *Employment Equity Act (EEA)* in 1998. Further, the EEA gave way to the implementation of affirmative action for previously

¹ According to the most recent World Bank figures, South Africa's Gini coefficient (the most used measure of income distribution) is the highest in the world at 63 % (World Bank, 2022) and recent research shows that since 1994, wealth ownership of the top 10 percent has fluctuated between 80 % and 90 % (Chatterjee et al., 2022).

² Apartheid laws such as the Bantu Education, the Colour Bar and Job Reservation ensured that black Africans were relegated to the lowest paying jobs, the effects of which are still apparent in the post-apartheid labour market.

marginalized groups i.e., women, people with disabilities and black people (African, Coloured, and Indian individuals) (Leibbrandt et al., 2010). We expect that these policies³ should impact both the overall wage structure (compress the wage structure) and narrow the gender wage gap overtime. Our summary measures would therefore give an indication if these policies have impacted the trend in the gender wage gap over time.

Third, South Africa has some of the highest trade union rates and strongest bargaining councils in the world even though the unionisation rates have been declining over time (Kerr and Wittenberg, 2021). While unions serve to compress the wage structure and therefore reduce overall wage inequality, workers who are not part of trade unions or bargaining councils mostly women in low paying jobs find themselves at a disadvantage. The stagnation of wages at the median has been partly attributed to the fact that these workers are least likely to be part of a trade union (Wittenberg, 2017). With regards to the gender wage gap, Mosomi (2019a) attributes the narrowing of the gender wage gap at the bottom of the wage distribution to the implementation of minimum wage laws in low paying industries where women are overrepresented such as domestic work. Given South Africa's unique labour market characteristics we expect our results to add to the ongoing debate on how the wage structure affects the trend in the wage gap.

In this analysis we first quantify the gender wage gap using the entropy metric contrasting the results with the mean and median. We then conduct stochastic dominance rankings and tests as a check to the entropy measure. Consistent with existing literature (see Mosomi (2019a)), our results show that the gender wage gap in South Africa has narrowed since 1995 irrespective of approach. Additionally, like Mosomi (2019a), results from the conventional measures show considerable variance in their characterization of the trend in the gender wage gap. For example, we find that unlike other quantiles, the median gender wage gap was stagnant between 1995 and 2015 (indicating that differences in evaluation functions matter).

Comparing the aggregative measures of the wage gap, we find that the rate of convergence implied by the entropy measure is faster than the one implied by the mean over the entire period (1995–2015).⁴ In other words, when the evaluation function is sensitive to inequality the rate of convergence is faster than when inequality is ignored an indication that for South Africa, if inequality was declining instead of increasing, the gender wage gap would be narrower. Specifically, unlike the mean gap, the entropy measure satisfies the Pigou-Dalton principle (Maasoumi and Wang, 2019) (i.e., a one-rand reduction in gap at lower wages is relatively more valuable than at higher wages). Given the (positive) relationship between general inequality and the gender wage gap (Blau, 2012), this suggests that a lower level of general inequality would have resulted into an even narrower gender wage gap. The stochastic dominance analysis agrees with the importance of accounting for inequality in evaluating the gender wage gap. While the distribution of male wages dominates that of female wages, this is the case only in the second order sense (i.e., only for evaluation functions that are sensitive to inequality). This suggests that when it comes to the gender wage gap in South Africa, dispersion of the wage distribution matters and ignoring this aspect may skew the result. The trend in the stochastic dominance rankings also suggests that the gap has narrowed over time.

Lastly, using the reweighting method by Firpo et al. (2009), we identify counterfactual distributions and use these to decompose the wage gap at the mean and at the distributional level using the other measures discussed. Like Mosomi (2019a) we also conclude that the difference between the distributions is largely driven by the wage structure effect. This implies that what is responsible for the gap does not vary with the choice of measure.

We find that (after accounting for the implausible rank invariance assumption and making explicit the choice of evaluation function), there is a narrowing of the (aggregate) gender wage gap in South Africa. Further, this narrowing is sensitive to inequality aversion, suggesting that interventions put in place by the South African government are more effective at curbing gender inequality at the extremes of the wage distribution. Specifically, the faster rate of convergence of the gender wage gap under the entropy measure (compared to the mean) is a function of the evaluation functions used. Our results show that the more extreme quantiles show more convergence, while the median gap has remained stubbornly high. It is therefore not surprising that aggregate measures that put more weights at the extremes show a faster rate of convergence.

The implication of these results is that policies that specifically target the middle of the wage distribution (majority of the population) may be necessary to further narrow the gender wage gap in the South African context. These include policies targeted at improving the quality of education for the majority population and not just the privileged who live in areas with good schools or who can afford private education. While education for the poor is free, the quality is poor which undermines the returns to skills. Dismantling gender stereotypes and equalizing the quality of education will be a first step towards reducing occupational segregation and the persistent gender wage gap at the median.

The rest of the paper is organized as follows, in Section 2 we discuss the data and the methods in detail. In Section 3 we present the results and Section 4 concludes.

³ Other policies include the Labour Relations Act 66 of 1995 and the Black Economic Empowerment Act, number 53 of 2003.

⁴ This is consistent with Maasoumi and Wang (2019)'s result based on United States data which shows that the convergence under the entropy measure is faster than all conventional measures. In our case, we find that the trends at the quantiles are a lot more volatile in the South African data and the rate of convergence under the entropy measure is faster than the mean measure not necessarily other quantiles.

2. Data

In this study we utilize the Post-Apartheid Labour Market Series (PALMS) 1993–2015 (Kerr et al., 2017). PALMS is a dataset comprising South African labour market surveys⁵ stacked together. These surveys are the 1993 Project for Statistics on Living Standards and Development, the October Household Surveys (1994–1999), the Labour Force Surveys (2000–2007), and the Quarterly Labour Force Surveys (QLFS), collected since 2008. The PALMS dataset contains the longest-running series of data on the post-apartheid labour market and is therefore particularly well suited for analysing labour market outcomes. Our sample is restricted to wage employed individuals aged between 15 and 65. The restriction to wage earners is to ensure comparability of earnings as there is controversy over self-employment self-reported earnings.

The main variable of interest is hourly wages whereby nominal Rands are deflated by the Consumer Price Index to obtain real wages (in constant 2015 South African Rands). The hourly wage variable is constructed by dividing real monthly earnings by monthly hours. Monthly hours are calculated by multiplying hours worked in the last week by average weeks in a month. Other variables of interest contained in PALMS include occupation, industry, age, marital status, years of education, race, province, and union status.

The dataset also contains bracket weights which take into account earnings given in brackets. These weights are also calibrated to consistent demographic and geographic estimates over time to improve the consistency of representation of the surveys over time (Branson and Wittenberg, 2014).

3. Methods

The measure of the gender wage gap is one of the many indicators of the progress towards achieving gender equality. In this analysis, we utilize the entropy measure which also satisfies the condition of being a metric⁶ as a measure of the gender wage gap in the South African context. Our analysis can therefore be linked to the rich literature on inequality measures and entropic functions that can be used to rank distributions in terms of their level of inequality.

Maasoumi and Wang (2014, 2019) argue that the gap between two distributions can be interpreted as the distance between the entropies (inequality measure) of the distributions. Entropies characterize and quantify the divergence between any distribution and a uniform distribution. Therefore, the contrast between entropies of distributions of male and female wages is the divergence between the two distributions (since the uniform distribution cancels out). The additional requirement that the measure should be a metric (rather than a divergence measure) is to make sure that it accommodates consistent assessment between multiple (more than two) distributions (i.e., satisfy the triangularity property of a metric).

Like the mean, the entropy metric is an aggregative measure of the gap that summarizes the information from all quantiles of the distributions being compared. Unlike the mean, this measure does not use equal weight to aggregate differences at the quantiles. Even though the entropy measure relies on a specific weighting scheme or evaluation function, a reason to consider this measure in addition to alternatives that are based on moments of distributions is that the weighting function (or preference function) underlying the entropy metric considers dispersion in wage distributions. For example, the mean will be invariant under a mean preserving spread⁷ even though such spreads induce changes in inequality. A similar argument is made by Maasoumi and Wang (2014) about the median. Intuitively, the entropy measure is a complex weighted measure of the differences at the quantiles (Maasoumi and Wang, 2014), only this weighting scheme is sensitive to inequality. We note that there is no universally accepted measure because all the measures discussed so far are subjective, albeit to varying degrees. The point we stress is that the weights adopted under the entropy measure are sensitive to inequality which is crucial in the South African context.

We follow the exposition of the method in Maasoumi et al. (2014) to estimate the gender wage gap. Consider a policy maker's preference presented by an evaluation function (EF) (Maasoumi and Wang, 2014) given by

$$EF(x) = \int_0^1 p'(\tau) F^{-1}(\tau) d\tau \quad (1)$$

where F is the CDF of wages and $p(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ is a twice continuously differentiable individual utility function (or preference or weighting function) and $p'(\cdot) \geq 0$ is the weight assigned to $F^{-1}(\tau)$ at the τ^{th} quantile. EF is therefore a weighted average of the quantiles of the distribution of wages. Ranking distributions with different measures is equivalent to comparing evaluation functions (see Abadie (2002) for a similar argument). When the evaluation is based on a single or a subset of quantiles, this is equivalent to assigning $p'(\cdot) = 0$ to all other quantiles. Such a weighting scheme is not consistent with any (reasonable) preference function with

⁵ Each survey design consisted about 30,000 households in 3000 clusters except for the OHS 1994 which consisted of 30,000 holds and 1000 clusters, the OHS 1996 which consisted of 16,000 households in 1600 clusters and the OHS 1998 which consisted of 20,000 households in 2000 clusters. The PSLSD 1993 and the LFS 2000 were also smaller.

⁶ That is, it is a true measure of distance in that it satisfies the triangularity rule. This contrasts with divergence measures that may not be consistent when there are three or more distributions (this is similar to the difference between correlation and covariance where the former is a metric, but the latter is not).

⁷ In probability and statistics, a mean-preserving spread (MPS) is a change from one probability distribution A to another probability distribution B, where B is formed by spreading out one or more portions of A's probability density function or probability mass function while leaving the mean (the expected value) unchanged.

heterogeneous outcomes (Maasoumi et al. 2014). When the preference function is the mean, equal weight is assigned to all earnings in the distribution. However, this is not how policy interventions are usually targeted especially in contexts where inequality is important. Even though the evaluation function underlying the entropic metric is subjective, unlike the conventional measures it does allocate some weight to all quantiles (in contrast to quantile measures) and these weights are not uniform (in contrast to the mean measure).

3.1. Estimating the gender wage gap: entropy measure

Let $\ln(w_f)$ and $\ln(w_m)$ represent log earnings for female and male workers respectively. In a particular year, we observe a random sample of N_f female and N_m male workers, where $N_f + N_m = N$. The vector $\{\ln(w_f)\}_{i=1}^{N_f}$ contains N_f observations denoted by $D_i = 1$; similarly $\{\ln(w_m)\}_{i=1}^{N_m}$ is a vector of N_m observations denoted by $D_i = 0$. Let $F_1(y) \equiv \Pr[\ln(w_f) \leq y]$ represent the cumulative density function (CDF) of $\ln(w_f)$ and $f_1(y)$ be the corresponding probability density function (PDF); $F_0(y)$ and $f_0(y)$ are similarly defined for $\ln(w_m)$.

Under the above notation the mean gender wage gap is defined as

$$E(\ln(w_f)) - E(\ln(w_m)) \quad (2)$$

The gender wage gap at the τ th quantile is

$$Q_\tau(\ln(w_f)) - Q_\tau(\ln(w_m)) \quad (3)$$

where the τ th quantile of F_1 (for example) is the smallest value $Q_\tau(\ln(w_f))$ such that $F_1(\ln(w_f)) = \tau$.

The gender wage gap is therefore some function of the distribution being compared. While one can define the “gap” between distributions as a function of the difference in the densities at specific points of the distribution (the first approach). From the decision theoretic standpoint, several axioms embody what one may consider as a suitable measure of the distance between distributions (see Maasoumi and Wang (2014, 2019) for details). The gap at any quantile or at the mean is a linear weighted function of the quantile gaps. Linear functions of quantiles imply infinite substitutability of a Rand (South African currency) at all wage levels. Alternative functions will reflect different types of weights indicating the degree of aversion to inequality. The entropy between two distributions is one example of a weighting scheme that is not linear and is sensitive to inequalities.

The proposed entropy metric is a normalization of the “Bhattacharya-Matusita-Hellinger” measure of distance (Granger et al., 2004). It is the one member of the GE family of divergence measures that is a metric. Therefore, one can define a distributional wage gap as

$$S_\rho = \frac{1}{2} \int_{-\infty}^{+\infty} \left(f_1^{\frac{1}{2}} - f_2^{\frac{1}{2}} \right)^2 dy \quad (4)$$

where f_1 and f_2 are the densities of the distributions being compared. It is possible to test the hypothesis $H_0 : (f_1 = f_2) \Leftrightarrow S_\rho = 0$, otherwise under $H_1 : S_\rho > 0$.⁸ The entropy measure (S_ρ) is implemented with Stata (“Srho” package, Maasoumi and Wang 201).⁹ It can also be implemented with R Statistical Package using “npdneqtest” (Li et al., 2009). The implementation of S_ρ in this study follows the implementation Maasoumi and Wang (2019). Like these authors, we use the Gaussian kernel and a robust version of the “normal reference rule-of-thumb” bandwidth $\left(= 1.06 \min(\sigma, \frac{IQR}{1.349}) n^{-\frac{1}{5}} \right)$ where σ is the standard deviation of the variable whose density is being estimated and IQR is the interquartile range.

Note that there are competing normalizations for the entropy function indexed by a parameter k that measures inequality aversion (Maasoumi and Wang, 2019). At $k = \frac{1}{2}$ one obtains the entropic distance metric which is the measure used in this study (varying k corresponds to different levels of inequality aversion in measuring the wage gap).¹⁰

3.2. Stochastic dominance tests

As noted above, the entropy metric is still subjective because it reflects the social welfare function of the generalized entropy function. However, when distributions cross (especially at lower tails) different measures will differ subjectively in their rankings depending on the evaluation function (Maasoumi and Wang, 2019). To have a robust comparison of the different distributions relative to large classes of welfare functions, we examine stochastic dominance rankings.

Stochastic (or prospect) dominance is an alternative to the above (strong) subjective measures. This is because it considers the

⁸ The measure has the following properties (i) It is well defined for both continuous and discrete variables (ii) it is normalized to $[0, 1]$ (iii) It is well defined and applicable when the variable of interest is multidimensional (iv) It is a metric, that is, it is a true measure of “distance” and not just of divergence (v) it is invariant under continuous and strictly increasing transformations

⁹ We are grateful to the authors for sharing their code with us.

¹⁰ Maasoumi and Wang (2019) used $k = 0.1, \dots, 0.9$ and their results show that the gender wage gap increases monotonically with the level of inequality aversion.

existence of a uniform (weak) ordering between distributions of outcomes that are robust to a very large class of decision functions (or preferences). If male wages first-order dominate female wages, then we can use the mean to characterize this dominance. However, if the distribution function crosses (for example at the lower tail) there is no first-order dominance. Using the mean to rank the distributions under this condition will be very subjective because this choice ignores the fact that outcomes are heterogeneous (in this context, the distribution with the higher mean does not necessarily represent a better prospect in terms of the distribution). For example, a social planner that places a lot of weight on the lower tails may not agree with the mean ranking under this situation. Stochastic dominance, on the other hand, enables a uniform ranking of distributions over a large class of preference functions.

Stochastic dominance can be used to rank distributions over a large class of preference functions because it is indexed by $p(\cdot)$. For example, if it is found that distribution (A) first-order dominates distribution (B), this implies that (A) is a better prospect under any evaluation (or preference) function that is increasing in wages. Note that this includes the mean and the quantiles. Furthermore, if distribution (A) second order dominates distribution (B), this implies that (A) is a better prospect than (B) for any preference function that is increasing in wages and concave. Concavity here implies aversion to higher dispersion. The implication of this is that stochastic dominance allows ranking of distributions under a wider set of preference functions.

Formally, first order stochastic dominance corresponds to a class (denoted by U_1) of all increasing von Neumann-Morgenstern type social welfare functions u such that welfare is increasing in wages (i.e., $u' > 0$). Second order stochastic dominance corresponds to the class of social welfare functions in U_1 such that $u'' \leq 0$ (i.e., concavity), denoted by U_2 . Concavity implies an aversion to higher inequality in wages across male and female workers. The measures are constructed as follows.

First Order Dominance: the wage distribution of male employees First Order Stochastically dominates that of female employees $\ln(w_m)$ FSD $\ln(w_f)$ if and only if

- $E[u(\ln(w_{mi}))] \geq E[u(\ln(w_{fi}))]$ for all $u \in U_1$ with strict inequalities for some u ;
- Or $F_m(y) \leq F_f(y)$ for all y with strict inequality for some y .

Second Order Dominance: the wage distribution of male employees Second Order Stochastically dominates that of female employees $\ln(w_m)$ SSD $\ln(w_f)$ if and only if

- $E[u(\ln(w_{mi}))] \geq E[u(\ln(w_{fi}))]$ for all $u \in U_2$ with strict inequalities for some u ;
- Or $\int_{-\infty}^y F_m(y) dy \geq \int_{-\infty}^y F_f(y) dy$ for all y with strict inequality for some y .

One advantage of this approach to the wage gap measurement is that Stochastic Dominance rankings are robust to the wage distribution or/and the weights assigned to the subgroups within the distribution (Maasoumi and Wang, 2014).

The Stochastic Dominance tests are based on the generalized Kolmogorov-Smirnov (KS) test discussed in Linton et al. (2005) and Maasoumi and Heshmati (2000). The KS test statistics for FSD and SSD are based on

$$d = \sqrt{\frac{N_f N_m}{N_f + N_m}} \min \left\{ \sup [F_m(y) - F_f(y)], \sup [F_f(y) - F_m(y)] \right\} \quad (5)$$

$$s = \sqrt{\frac{N_f N_m}{N_f + N_m}} \min \left\{ \sup \int_{-\infty}^y [F_m(y) - F_f(y)] dt, \sup \int_{-\infty}^y [F_m(y) - F_f(y)] dt \right\} \quad (6)$$

In reporting the empirical test, we denote $\sup [F_m(y) - F_f(y)]$ as $d_1 \max$ and $\sup [F_f(y) - F_m(y)]$ as $d_2 \max$. We report both $d_1 \max$ and $d_2 \max$ along with the test statistic d for clarity of interpretation. $s_1 \max$ and $s_2 \max$ are similarly defined. The tests statistics are based on the sample counterparts of d and s using empirical CDFs given by

$\widehat{F}_0(y) = \frac{1}{N_f} \sum_{i=1}^{N_f} I(\ln(w_{mi})) \leq y$, where $I(\cdot)$ is an indication function, and N_f is the sample size as defined earlier. $\widehat{F}_1(y)$ is similarly defined. The underlying distributions of the statistics are unknown and depend on the data. Following Maasoumi and Wang (2014) we use a simple bootstrap technique based on 200 replications to obtain the empirical distribution of the test statistic. This approach estimates the probability of the statistics falling in any desired interval, as well as p -values. In our case, if the probability of d lying in the non-positive interval (i.e., $\Pr[d \leq 0]$) is large, say 0.9 or higher and $\widehat{d} \leq 0$, then we can infer FSD. We can infer SSD based on s and $\Pr[s \leq 0]$ in a similar fashion. For example, if we observe FSD (SSD) and $\Pr[d \leq 0] (\Pr[s \leq 0])$ is 0.95, then it means the test statistic is significant at 5 % level (p -value = 0.05)

3.3. Decompositions

To understand the driving forces of the gender wage gap we decompose the entropy measure raw gap into the compositional (differences in observable characteristics e.g., education, occupation, marital status) and wage structure effects (returns to the observable characteristics).

Formally, let the female wage structure be denoted by g_1 and the male one be denoted by g_0 . Individual wages are determined by observed characteristics X_i ¹¹ and unobserved characteristics ε_i via unknown wage structure functions g_d i.e.,

$$\ln(w_{iD}) = g_D(X_i, \varepsilon_i) \text{ for } D = 0 \text{ or } 1 \quad (7)$$

We assume wages are determined non-parametrically which avoids the imposition of a distributional assumption or a specific functional form. This allows for flexible interaction between X_i and ε_i . Following (Maasoumi and Wang, 2014) we assume that the unknown joint distribution $(\ln(w), X_i, \varepsilon_i)$ exists. Under the unconfoundedness assumption, differences in wages between the two groups comes from the composition effect i.e., differences in X_i and wage structure effect i.e., differences in $g_D(\cdot)$

With observed data we can identify the conditional distribution of female wages $\ln(w_{if})|X, D = 1 \sim F_{f|X}$ and the conditional distribution of male wages $\ln(w_{im})|X, D = 0 \sim F_{m|X}$. Firpo et al. (2009) show that under the following assumptions

1. Unconfoundedness: Suppose $(\ln(w), D, X)$ have a joint distribution we assume $(\ln(w_{11}), \ln(w_{10}))$ and D are jointly independent conditional on X
2. Common Support: For all $x \in X$, $0 < \Pr(D = 1|X = x) < 1$

the counterfactual distribution $\ln(w_{im})|X, D = 1$ can be identified using the propensity score reweighting method. This distribution represents the distribution of wages for men given that they have the same characteristics as women. If we denote ν as a function of the conditional joint distribution $(\ln(w_{10}), \ln(w_{11}))|D$ i.e., $\nu : F_\nu \rightarrow \mathbb{R}$, where F_ν belongs to a class of distributional functions that satisfy $|\nu(F)| \leq \infty$. Note that the mean, variance, quantiles, and the entropy metric are all examples of ν .

Under this specification the wage gap between two groups can be written in terms of ν :

$$\Delta_o^\nu = \nu(F_1) - \nu(F_0) = \nu_1 - \nu_0 \quad (8)$$

where Δ_o^ν is the overall wage gap based on the functional ν . The above equation can further be decomposed into two parts

$$\Delta_o^\nu = (\nu_1 - \nu_C) + (\nu_C - \nu_0) = \Delta_S^\nu + \Delta_X^\nu \quad (9)$$

Note that when $\nu_C = \ln(w_{im})|X, D = 1$ the term Δ_S^ν is the wage structure effect that represents the effect of changing $g_1(\cdot)$ to $g_0(\cdot)$ while holding the characteristics $(X_i, \varepsilon_i)|D = 1$ constant.¹² The other term Δ_X^ν is the composition effect, which is the effect of changing the distribution of characteristics from $(X_i, \varepsilon_i)|D = 1$ to $(X_i, \varepsilon_i)|D = 0$, while keeping the wage structure $g_0(\cdot)$ constant.

Following Maasoumi and Wang (2014) we note that the second term of Eq. (8) will require using the distribution of wages for men $\ln(w_{10})$ as the benchmark. For ease of interpretation, we chose to use $\ln(w_{iw})$ as the benchmark in both cases. Therefore, we construct 2 counterfactuals

- Counterfactual outcome 1: $(\ln(w_i))^{C1} = g_0(X_i, \varepsilon_i)|D = 1 = \nu_{C1}$
- Counterfactual outcome 2: $(\ln(w_i))^{C2} = g_1(X_i, \varepsilon_i)|D = 0 = \nu_{C2}$

Like (8), the difference between counterfactual outcome 1 and the distribution of wages for female workers represents the wage structure effect, however unlike (8), the composition effect is the difference between the female distribution of wages and counterfactual outcome 2.

The counterfactual distributions 1 and 2 can be identified using propensity score reweighting

$$F_{C1}(y) = \mathbb{E}[\omega_{C1}(D, X) \cdot I\{Y \leq y\}] \quad (10)$$

$$F_{C2}(y) = \mathbb{E}[\omega_{C2}(D, X) \cdot I\{Y \leq y\}] \quad (11)$$

where $y = \ln(w_i)$, D is the group dummy taking 1 for female workers and I is an indication function. $\omega_{C1}(D, X) = \left(\frac{p_1(x)}{1-p_1(x)}\right)\left(\frac{1-D}{p}\right)$ and $\omega_{C2}(D, X) = \left(\frac{p_0(x)}{1-p_0(x)}\right)\left(\frac{D}{1-p}\right)$. p is the proportion of female workers, $p_1(x) = \Pr\{D = 1|X = x\}$ represent the probability of being in group 1 given X , $p_0(x) = \Pr\{D = 0|X = x\}$ is defined in a similar way.

Using these counterfactuals, we estimate the wage structure effect and composition effect for ν that is defined as the mean, quantiles, entropy metric and the KS statistic as defined in Eqs. (4) and (5).

4. Results

4.1. The trend in the gender wage gap

In this section we present the raw log wage differential between men and women from different measures (mean, median, 10th

¹¹ These include potential experience, level of education, location, level of education, race, marital status, occupation an industry

¹² Recall that ν_1 refers to the distribution of wages for female workers.




Fig. 1. Normalized raw gender wage values for different measures of the gap.

Note: Sample contains individuals aged between 15 and 65 years of age.

Source: Authors own calculation based on Post-Apartheid Labour Market Surveys (PALMS).

percentile, 25th percentile, 75th percentile, 90th percentile and the entropy measure). The entropy measure (S_p) is a normalized metric taking values between 0 and 1, so for ease of interpretation we report the values multiplied by 100. In addition, the entropy measure (S_p) and other conventional measures are not directly comparable therefore we normalized all measures by setting the value at 1995 equal to 1 to easily track changes over time.

Fig. 1 plots the normalized raw gender wage gap values and shows that compared to the wage gap in 1995, there has been a decline of the wage gap over time for all measures except the median. The results from the figure suggest that there was some convergence of the gender wage gap in the upper quantiles (75th and 90th) between 1995 and 2005 however this trend has reversed in recent years. After 2009, the gender wage gap at the 10th and 25th quantile continued narrowing while it widened at the 75th and the 90th quantile. And, as has been documented in the literature (Mosomi 2019a; Bhorat and Goga 2013; Ntuli 2007), while initially the wage gap was highest at the bottom of the wage distribution, this has changed in recent years.

The decline of the gender wage gap coincides with the introduction of several labour market policies which were aimed at reducing employment discrimination and improving working conditions. The Basic Conditions of Employment Act 75 of 1997 for example gave way to the setting up of minimum wages in the low paying industries where women are overrepresented while the Employment Equity Act 55 of 1998 gave way to the introduction of affirmative action. The puzzle however is that, despite all these policies, the median gap displays a very peculiar trend with the gap remaining higher than the 1995 value and showing no declining trend.

The main contribution of our analysis is that, amongst all the summary measures, the entropy measure which considers overall inequality shows a faster convergence of the gender wage gap in South Africa as shown in Fig. 2. Fig. 2 compares the three summary measures, the mean, the median and the entropy measure. The entropy measure shows a quicker convergence because it is sensitive to inequality between the wage distributions. This is achieved by placing more weights at the extremes of the distributions. This is in contrast of the mean which places equal weight to all the observations and the median that places all the weight at the 50th percentile.

We note that while the declining trend of the entropy metric is qualitatively similar to the one displayed by the mean, our results show that the annualized rate of convergence displayed by the entropy metric is much larger than the one displayed by the mean. That is, while the mean suggests that the wage differential converges by approximately 4.3 % per annum,¹³ the entropy measure suggests that the annual rate of convergence is 10.23 % per annum over the period under consideration. Our results are similar to those reported by Maasoumi and Wang (2019) in that the annualized rate of convergence under the entropy measure is faster than all the conventional measures (mean and quantiles) in the USA. However, in the case of South Africa, the annualized rate of convergence at the 10th and 90th quantiles (10.9 and 12.16 % respectively) is faster than the convergence under the entropy measure while convergence at the other quantiles are slower. This suggests that the rate of convergence is faster at the tails in South Africa highlighting the importance of dispersion in the wage distribution.

¹³ This is obtained by $\left(1 - \left(\frac{\text{gap at the end}}{\text{gap at the beginning}}\right)^{\frac{1}{10}}\right) \times 100$




Fig. 2. The raw gender wage gap at the mean, median and the entropy measure.

Source: Authors own calculation based on Post-Apartheid Labour Market Surveys (PALMS) Note: Sample contains individuals aged between 15 and 65 years of age.

4.2. The stochastic dominance test results

We present the results of the stochastic dominance ranking and tests in Table 1. The theoretical basis for the stochastic dominance tests is particularly relevant in this case study since the decline in the gender wage gap seems to be more pronounced at the tails. First-order stochastic dominance (FSD) implies that one distribution is “better” than the other for all evaluation functions that are increasing in wages while second-order stochastic dominance (SSD) implies a similar ordering for utility functions that are increasing in wages and concave. The last two rows of Table 1 show whether FSD or SSD is observed and an asterisk is included when the p -value¹⁴ is greater than or equal to 0.9 i.e., when the observed difference is significant.

Results show that the male wage distribution first-order dominates the female wage distribution only once in the entire series (2000: Q1 (LFS 01:1)) and this first-order ranking is not significant. In contrast, as shown in the last row of Table 1, a second-order stochastic dominance in favour of the male wage distribution occurs more frequently (13 times) and of the 13 times, the dominance is significant 3 times. The absence/scarcity of FSD in Table 1 may look odd at first, but it underscores the point that evaluation functions are important when ranking distributions. Fig. 1 clearly shows that the gender wage gap is heterogeneous across quantiles. A quantile is an evaluation function of the type¹⁵ u_1 , that is, it increases in wages. Fig. 1 shows that there are instances where women at the 75th percentile are earning better than men. The absence of FSD therefore confirms that in the South African case, male wages are not always higher¹⁶ than female wages for all evaluation functions of the type u_1 .

In contrast, when the evaluation function is of the type¹⁷ u_2 i.e., increasing in wages and concave, then the frequency of dominance observed increased. Therefore, when the evaluation function is not sensitive to inequality it may not pick up these nuances at different parts of the male and female wage distributions.

Our result is in stark contrast to what was observed by Maasoumi and Wang (2014) (using data from the USA between 1976 and 2013). The authors find significant FSD in favour of male workers throughout the period considered. One explanation for this is that inequality¹⁸ is higher in South Africa compared to the USA. Meaning that the influence of overall inequality on the trend in the gender

¹⁴ Recall that if the probability of d lying in the non-positive interval (i.e., $Pr[d \leq 0]$) is large, say 0.9 or higher and $\hat{d} \leq 0$, then we can infer FSD with a high degree of statistical significance. However, if $\hat{d} \leq 0$ but the p -value is less than 0.9 we still observe FSD but the dominance is not significant. We can infer SSD based on s and $Pr[s \leq 0]$ in a similar fashion.

¹⁵ Recall that this refers to social welfare functions that are increasing in wages.

¹⁶ This is illustrated more clearly by cumulative distribution function plots in Figures A1–A3 in the appendix showing that while the female wage distribution is generally to the left of the male wage distribution, they do cross most of the time during the period under consideration.

¹⁷ Recall that this refers to social welfare functions that are increasing in wages and concave.

¹⁸ According to the World Bank poverty and inequality platform, the Gini index for the USA was 41.5 in 2014 while that for South Africa was 63 in the same year.

Table 1
Stochastic Dominance Tests.

	$d_1\max$	$d_2\max$	D	$\Pr(d \leq 0)$	$s_1\max$	$s_2\max$	s	$\Pr(s \leq 0)$	First order	Second order
OHS 1995	5.28	0.01	0.01	0.06	1059.13	0.00	0.00	0.42	No	Yes
OHS 1997	6.75	0.11	0.11	0.00	1045.26	−0.01	−0.01	0.70	No	Yes
OHS 1998	4.78	0.50	0.50	0.00	635.67	0.01	0.01	0.38	No	No
OHS 1999	7.91	0.01	0.01	0.05	829.98	0.00	0.00	0.45	No	Yes
LFS 00:1	6.07	0.25	0.25	0.14	833.58	−0.10	−0.10	0.95	No	Yes*
LFS 00:2	7.74	0.28	0.28	0.01	721.83	0.01	0.01	0.25	No	No
LFS 01:1	7.52	0.00	0.00	0.06	1002.79	0.00	0.00	0.82	Yes	Yes
LFS 01:2	8.58	0.02	0.02	0.02	950.63	0.03	0.03	0.35	No	No
LFS 02:1	9.17	0.02	0.02	0.02	1191.04	0.03	0.03	0.16	No	No
LFS 02:2	7.12	1.61	1.61	0.00	707.41	−0.03	−0.03	1.00	No	Yes*
LFS 03:1	6.25	1.40	1.40	0.00	773.20	0.00	0.00	0.26	No	Yes
LFS 03:2	6.80	0.56	0.56	0.00	768.49	0.02	0.02	0.19	No	No
LFS 04:1	7.11	1.51	1.51	0.00	675.57	0.02	0.02	0.14	No	No
LFS 04:2	7.50	0.54	0.54	0.00	770.76	0.02	0.02	0.03	No	No
LFS 05:1	7.78	1.81	1.81	0.00	843.08	0.06	0.06	0.27	No	No
LFS 05:2	6.71	0.91	0.91	0.00	549.51	−0.01	−0.01	0.38	No	Yes
LFS 06:1	6.84	0.09	0.09	0.00	823.45	0.05	0.05	0.00	No	No
LFS 06:2	6.79	0.02	0.02	0.00	942.99	0.13	0.13	0.13	No	No
LFS 07:1	7.02	1.33	1.33	0.00	655.32	0.00	0.00	0.58	No	Yes
LFS 07:2	5.54	0.51	0.51	0.00	698.39	0.37	0.37	0.05	No	No
QLFS 2010:1	5.88	0.02	0.02	0.01	636.40	0.02	0.02	0.28	No	No
QLFS 2010:2	5.49	0.05	0.05	0.00	658.08	0.00	0.00	0.55	No	Yes
QLFS 2010:3	5.65	0.13	0.13	0.02	606.77	0.02	0.02	0.31	No	No
QLFS 2010:4	5.19	0.01	0.01	0.34	576.85	0.00	0.00	1.00	No	Yes*
QLFS 2011:1	5.26	0.03	0.03	0.00	656.48	0.01	0.01	0.28	No	No
QLFS 2011:2	5.47	0.68	0.68	0.00	512.22	0.27	0.27	0.03	No	No
QLFS 2011:3	4.79	0.02	0.02	0.04	509.50	0.00	0.00	0.50	No	Yes
QLFS 2011:4	5.50	0.08	0.08	0.00	665.81	0.03	0.03	0.13	No	No
QLFS 2012:1	6.07	0.16	0.16	0.00	632.92	1.45	1.45	0.03	No	No
QLFS 2012:2	5.80	0.24	0.24	0.00	561.16	0.00	0.00	0.43	No	Yes
QLFS 2012:3	5.02	0.22	0.22	0.00	567.08	6.16	6.16	0.00	No	No
QLFS 2012:4	4.95	0.11	0.11	0.00	487.15	3.11	3.11	0.00	No	No
QLFS 2013:1	5.37	0.25	0.25	0.00	513.73	9.32	9.32	0.00	No	No
QLFS 2013:2	4.77	0.22	0.22	0.00	523.68	5.16	5.16	0.03	No	No
QLFS 2013:3	6.00	0.37	0.37	0.00	496.07	11.97	11.97	0.00	No	No
QLFS 2013:4	5.25	0.43	0.43	0.00	440.00	15.92	15.92	0.00	No	No
QLFS 2014:1	5.44	0.40	0.40	0.00	564.95	12.10	12.10	0.00	No	No
QLFS 2014:2	5.04	0.09	0.09	0.00	541.75	2.36	2.36	0.02	No	No
QLFS 2014:3	5.61	0.17	0.17	0.00	579.38	5.54	5.54	0.08	No	No
QLFS 2014:4	5.90	0.14	0.14	0.00	709.99	3.51	3.51	0.00	No	No
QLFS 2015:1	4.63	0.50	0.50	0.00	364.98	17.09	17.09	0.00	No	No
QLFS 2015:2	4.95	0.15	0.15	0.00	474.22	3.49	3.49	0.00	No	No
QLFS 2015:3	4.61	0.34	0.34	0.00	584.30	9.51	9.51	0.05	No	No
QLFS 2015:4	4.42	0.21	0.21	0.00	513.06	8.83	8.83	0.00	No	No

Notes: Order columns specify the order of dominance observed * indicates if the observed dominance is significant.

Source: Authors own calculation based on Post-Apartheid Labour Market Surveys (PALMS) Note: Sample contains individuals aged between 15 and 65 years of age.

wage gap is stronger and a measure that is sensitive to inequality will capture this more strongly in the South African context. Another reason for this observation is that the wealthiest sub-groups white males (and to a lesser extent black men) are less likely to be captured by the kind of dataset¹⁹ utilised in this study (household surveys).

In summary, we find that the choice of evaluation function matters and that there is a case to be made for considering evaluation functions that are consistent with a larger class of preference functions while not neglecting the narrower evaluation functions when their preference functions are made explicit. While the results largely agree across evaluation functions our argument is that there are subtle differences. For example, we have shown that the rate of convergence can be very different under different evaluation functions. In the next section, we discuss results from the decomposition analysis.

4.3. Decomposition results

In this section, we discuss the decomposition of the results of the entropy measure. Stochastic dominance decomposition results are

¹⁹ The 2019 Employment equity report shows that white men still dominate managerial positions in the private sector while black men dominate managerial positions in the public sector. The point here is that the gap may be larger than what is observed because the richer sub-group are less likely to be captured in the surveys, so that results from surveys do not capture the full scope of inequality in South Africa.

displayed in Tables A1 and A2 in the appendix. Recall that the returns (or wage structure) effect represents the difference between the female wage distribution and the counterfactual wage distribution ($c1$) where the male wage distribution is weighted to have an identical distribution of observable characteristics as those of female workers (counterfactual 1) (that is, women paid as if they were men) (note that the difference is $\ln(w_m) - \ln(w_i)^{c1}$). The estimates of the wage structure effect represent the portion of the wage gap that is not explained by observed characteristics, i.e., inequality in pay structure (or unexplained gap). The composition effect is the component of the wage gap that is explained by differences in characteristics i.e., the difference between the wage distribution of female workers and the distribution of female workers weighted by male characteristics (counterfactual 2 where men are paid as if they were women) (note that the difference is $\ln(w_f) - \ln(w_i)^{c2}$).

The decomposition results of the entropy measure are presented in Fig. 3. The results show that the return effect (wage structure effect) is always positive meaning that wages under the male wage structure but with female characteristics are generally better than the actual female wages. In other words, if male workers had identical characteristics to female workers (*counterfactual1*) they would still earn more than female workers. This result is consistent with both local and international literature (Blau and Kahn, 2017; Mosomi, 2019a,b) and suggests that male workers enjoy a premium that is largely unexplained (or due to unobserved characteristics). The return effect has however been declining over time consistent with declining wage gap over time.

The composition effect is also positive suggesting that wages under the female wage structure weighted by male characteristics are lower than actual female wages. The implication here is that women on average have better observable human capital characteristics compared to men. In fact, if women had the same observable characteristics as men, their wages would be worse.

In summary, our results are consistent with the literature in that the return effect explains a larger proportion of the wage gap and that both the wage structure and the composition effects have dwindled in recent years indicating that the wage gap and its components have reduced over time.

5. Limitations of the study

It is important to note however that even though this is the best suited data for a trend analysis of the gender wage gap, there exists some data quality issues which might complicate the interpretation of our results across different household surveys. One is that since 1993, there have been several changes to the survey instrument as Statistics South Africa (Stats SA) tried to improve data collection and this may have resulted in some inconsistencies between household surveys over time. Perhaps this explains some of the volatility observed at the quantiles.

Second and related to the first point is that due to the data quality issues, the choice of baseline year matters when carrying out a trend analysis using South African Labour Force surveys. For example, the early OHSs (1994–1998) have been regarded as poor baselines due to the under sampling of low-income work (domestic work, mining, and subsistence agriculture) (Kerr and Wittenberg, 2015). Mosomi (2019b) showed how the under sampling of domestic workers in 1994 and the inconsistency in the classification of domestic workers in 1995 led to insignificant gender wage gaps in these years. This inconsistency has however been corrected in the




Fig. 3. Entropy measure decomposition results.

Note: Sample contains individuals aged between 15 and 65 years of age.

Source: Authors own calculation based on Post-Apartheid Labour Market Surveys (PALMS).

version of the data used for this analysis (see Mosomi (2019b) for details).

Third, studies on wage trends and inequality report that the most recent Quarterly Labour Force Surveys (QLFSs) exhibit many breaks stemming from imputations done on the earnings variable (Kerr and Wittenberg, 2021). This is an ongoing research issue and thus trends after the second quarter of 2012 should be interpreted with caution.

Finally, unlike Maasoumi and Wang (2019), we do not control for selection because of data restrictions. Appropriate exclusion restrictions are in practice hard to find and according to Puhani (2000), the lack of appropriate exclusion restrictions may generate collinearity issues, resulting in unreliable coefficients and inflated standard errors. Additionally, studies that have controlled for selectivity bias report that the coefficients for λ are mostly not significant (Hinks, 2002; Ntuli, 2007; Shepherd, 2008). Therefore, we note that our result is valid for the wage employed but perhaps not the population.

6. Conclusion

This paper reexamines the trend in the gender wage gap using entropy distance and stochastic dominance tests. This approach is motivated by the need for measures that capture inequality in distributions while tracking the trend in the gender wage gap. The entropy distance and stochastic dominance tests summarize each distribution by suitable evaluation functions and then compute the difference between the evaluations. Unlike conventional approaches based on the mean or quantiles, these approaches do not assume rank invariance and makes explicit the evaluation functions that are utilized. We contrast and compare the results under this new approach to results under conventional methods in the South African context given the persistent high level of inequality and the fact that the government has introduced many policies to reduce the inequality levels.

A unique contribution of this study is that the rate of convergence of the gender wage gap is faster under the entropy measure in comparison to the mean measure. We argue that this is due to differences in evaluation (or weighting) functions. The mean gap is insensitive to inequality (i.e., imposes equal weights on the quantiles) while the entropy gap uses weights that are sensitive to inequality. By assigning more weight at the extremes, the entropy measure displays a faster convergence of the gender wage gap in South Africa. Assigning equal weight across the distribution will bias the results given that at the top and bottom of the wage distribution (extreme quantiles) there is faster convergence while the median wage gap remains high throughout the period under consideration. Stochastic dominance tests (that are based on more general evaluation functions) confirm that the gap has narrowed. However, the Stochastic dominance results show that this dominance is only for evaluation functions that are sensitive to inequality.

The sensitivity of the gender wage gap to overall inequality both in general and in the South African context has policy implications. Our results suggest that policies that are designed to address inequality can have differential effects across the income distribution. While minimum wage policies benefit lower income earners, employment equity and affirmative action laws are likely to benefit highly educated women. In the South African context, this explains why the wage gap is heterogeneous with some women having an earning advantage (over comparable males), especially when they work in male dominated sectors (Roberts and Schöer, 2021; Posel et al., 2023). Affirmative action laws benefited highly educated women who happen to be majority white due to inequalities in education because of apartheid policies while minimum wage laws have impacted the gender wage gap at the bottom.

It however appears that these policies do not address inequality at the middle of the income distribution. To address the gap at the median, there is need to level the ground when it comes to the quality of education. As long as students get unequal education, those with the access will have an advantage. Also given the role occupational sorting plays in the persistence of the gender wage gap (Blau and Kahn, 2017), addressing contextual factors including gender norms and stereotypes that influence the subject choice in school will be necessary to further narrow the gender wage gap.

Finally, addressing structural issues such as how women's work is remunerated will help narrow the gap even further. The recent COVID-19 pandemic has brought to the fore how occupational segregation not only plays an important role for gender wage gap, but also for welfare in the case of a shock. Women are concentrated in care and services occupations which are considered essential but lowly paid (teachers, nurses, child minders). Due to the occupations women are concentrated in, they have been disproportionately affected by the pandemic both in terms of job losses and exposure to health risks (Oyenubi, 2023; Mosomi and Thornton, 2022). There is need for policies that reduce the burden of care for women and fair remuneration in industries where women are concentrated.

Availability of data and materials

The data is publicly available from <https://www.datafirst.uct.ac.za/>.

Funding

This article was not funded.

CRediT authorship contribution statement

Adeola Oyenubi: Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing. **Jacqueline Mosomi:** Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no conflict of interest.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.eap.2023.12.017](https://doi.org/10.1016/j.eap.2023.12.017).

Appendix

Table A1
Stochastic dominance ranking and tests (wage structure effect).

	d_1 max	d_2 max	d	$Pr(d \leq 0)$	s_1 max	s_2 max	s	$Pr(s \leq 0)$	First order	Second order
OHS 1995	14.15	0.01	0.01	0.23	2059.10	0.00	0.00	1.00	No	Yes*
OHS 1997	17.54	0.03	0.03	0.02	2447.67	0.00	0.00	0.99	No	Yes*
OHS 1998	13.42	0.03	0.03	0.05	1864.23	0.00	0.00	0.95	No	Yes*
OHS 1999	17.15	0.00	0.00	0.56	2035.69	0.00	0.00	0.90	Yes	Yes*
LFS 00:1	11.76	0.23	0.23	0.15	1481.64	19.64	19.64	0.28	No	No
LFS 00:2	15.76	0.10	0.10	0.06	1684.28	0.00	0.00	0.42	No	Yes
LFS 01:1	18.18	0.01	0.01	0.14	2675.93	0.00	0.00	0.97	No	Yes*
LFS 01:2	17.83	0.02	0.02	0.01	2287.21	0.01	0.01	0.09	No	No
LFS 02:1	18.61	0.00	0.00	0.66	2706.20	−0.01	−0.01	0.98	Yes	Yes*
LFS 02:2	18.03	0.01	0.01	0.10	2152.85	0.00	0.00	0.63	No	Yes
LFS 03:1	12.77	0.02	0.02	0.01	2229.13	0.04	0.04	0.38	No	No
LFS 03:2	14.74	0.04	0.04	0.00	2074.64	2.41	2.41	0.01	No	No
LFS 04:1	15.09	0.05	0.05	0.00	2197.38	1.48	1.48	0.12	No	No
LFS 04:2	17.01	0.03	0.03	0.02	2086.92	0.14	0.14	0.35	No	No
LFS 05:1	17.69	0.01	0.01	0.14	2283.83	−0.01	−0.01	0.76	No	Yes
LFS 05:2	17.08	0.06	0.06	0.00	1887.46	−0.01	−0.01	0.94	No	Yes*
LFS 06:1	17.34	0.02	0.02	0.12	1842.47	0.00	0.00	0.69	No	Yes
LFS 06:2	17.66	0.003	0.003	0.04	2534.07	0.03	0.03	0.16	No	No
LFS 07:1	17.31	0.00	0.00	0.10	2429.27	0.00	0.00	0.27	Yes	Yes
LFS 07:2	17.81	0.09	0.09	0.02	2310.09	0.01	0.01	0.17	No	No
QLFS 2010:1	8.80	0.13	0.13	0.00	1094.94	0.01	0.01	0.03	No	No
QLFS 2010:2	8.68	0.11	0.11	0.00	1118.76	1.80	1.80	0.17	No	No
QLFS 2010:3	8.77	0.01	0.01	0.17	1112.78	0.00	0.00	0.98	No	Yes*
QLFS 2010:4	8.69	0.00	0.00	0.86	1123.43	0.00	0.00	0.98	Yes	Yes*
QLFS 2011:1	8.00	0.05	0.05	0.00	1044.05	0.79	0.79	0.20	No	No
QLFS 2011:2	7.53	0.01	0.01	0.15	858.10	0.00	0.00	0.54	No	Yes
QLFS 2011:3	7.73	0.01	0.01	0.16	988.32	0.05	0.05	0.53	No	No
QLFS 2011:4	8.19	0.05	0.05	0.00	1100.12	1.68	1.68	0.03	No	No
QLFS 2012:1	8.47	0.08	0.08	0.00	1005.31	2.08	2.08	0.02	No	No
QLFS 2012:2	8.89	0.04	0.04	0.00	977.47	1.50	1.50	0.18	No	No
QLFS 2012:3	6.76	0.16	0.16	0.00	762.17	6.39	6.39	0.00	No	No
QLFS 2012:4	7.24	0.02	0.02	0.00	732.38	0.04	0.04	0.22	No	No
QLFS 2013:1	6.72	0.11	0.11	0.00	675.55	2.21	2.21	0.01	No	No
QLFS 2013:2	6.83	0.09	0.09	0.00	705.52	0.71	0.71	0.09	No	No
QLFS 2013:3	8.35	0.43	0.43	0.00	703.46	14.94	14.94	0.00	No	No
QLFS 2013:4	6.96	0.71	0.71	0.00	619.06	33.20	33.20	0.00	No	No
QLFS 2014:1	7.73	0.49	0.49	0.00	719.45	24.08	24.08	0.01	No	No
QLFS 2014:2	6.73	0.47	0.47	0.00	687.96	27.53	27.53	0.00	No	No
QLFS 2014:3	7.04	0.26	0.26	0.00	738.24	12.79	12.79	0.05	No	No
QLFS 2014:4	7.35	0.41	0.41	0.00	811.62	23.53	23.53	0.00	No	No
QLFS 2015:1	5.88	0.81	0.81	0.00	437.65	49.26	49.26	0.00	No	No
QLFS 2015:2	5.92	0.36	0.36	0.00	553.06	22.21	22.21	0.00	No	No
QLFS 2015:3	5.68	0.49	0.49	0.00	615.89	22.74	22.74	0.00	No	No
QLFS 2015:4	5.21	0.35	0.35	0.00	547.79	27.72	27.72	0.00	No	No

Note: Source: Authors own calculation based on Post-Apartheid Labour Market Surveys (PALMS) Note: Sample contains individuals aged between 15 and 65 years of age.

Table A2
Stochastic dominance ranking and tests (composition effect).

	d_1 max	d_2 max	d	$Pr(d \leq 0)$	s_1 max	s_2 max	s	$Pr(s \leq 0)$	First order	Second order
OHS 1995	5.16	0.03	0.03	0.01	565.90	0.00	0.00	0.73	No	Yes
OHS 1997	8.47	0.01	0.01	0.05	1050.32	0.00	0.00	0.69	No	Yes

(continued on next page)

Table A2 (continued)

	d_1 max	d_2 max	d	$Pr(d \leq 0)$	s_1 max	s_2 max	s	$Pr(s \leq 0)$	First order	Second order
OHS 1998	7.04	0.02	0.02	0.01	920.40	0.04	0.04	0.38	No	No
OHS 1999	7.78	0.06	0.06	0.00	993.73	0.00	0.00	0.51	No	Yes
LFS 00:1	5.41	0.06	0.06	0.00	724.62	0.00	0.00	0.51	No	Yes
LFS 00:2	3.07	0.63	0.63	0.00	240.90	19.50	19.50	0.10	No	No
LFS 01:1	6.29	0.23	0.23	0.00	873.05	2.64	2.64	0.26	No	No
LFS 01:2	7.06	0.04	0.04	0.01	791.55	0.36	0.36	0.28	No	No
LFS 02:1	7.82	0.03	0.03	0.08	919.68	0.00	0.00	0.36	No	Yes
LFS 02:2	6.57	0.41	0.41	0.01	798.27	1.65	1.65	0.34	No	No
LFS 03:1	6.96	0.01	0.01	0.10	1249.33	0.00	0.00	0.65	No	Yes
LFS 03:2	7.37	0.00	0.00	0.32	1105.29	0.00	0.00	0.65	Yes	Yes
LFS 04:1	8.84	0.00	0.00	0.19	1283.66	0.00	0.00	0.51	Yes	Yes
LFS 04:2	9.46	0.11	0.11	0.03	1115.43	0.00	0.00	0.44	No	Yes
LFS 05:1	9.04	0.00	0.00	0.31	1254.50	−0.01	−0.01	0.59	Yes	Yes
LFS 05:2	7.93	0.44	0.44	0.00	749.09	8.98	8.98	0.17	No	No
LFS 06:1	9.78	0.00	0.00	0.32	947.05	0.00	0.00	0.71	Yes	Yes
LFS 06:2	9.55	0.01	0.01	0.10	1432.54	0.12	0.12	0.31	No	No
LFS 07:1	8.39	0.03	0.03	0.07	1341.65	0.42	0.42	0.23	No	No
LFS 07:2	9.19	0.00	0.00	0.19	1156.80	0.02	0.02	0.36	Yes	No
QLFS 2010:1	1.74	0.68	0.68	0.00	115.88	0.27	0.27	0.19	No	No
QLFS 2010:2	2.47	0.51	0.51	0.00	158.39	1.80	1.80	0.22	No	No
QLFS 2010:3	1.34	0.84	0.84	0.00	35.92	6.64	6.64	0.12	No	No
QLFS 2010:4	1.12	1.75	1.12	0.00	41.92	59.48	41.92	0.14	No	No
QLFS 2011:1	1.91	0.47	0.47	0.00	130.96	0.31	0.31	0.28	No	No
QLFS 2011:2	0.59	2.19	0.59	0.00	20.85	182.10	20.85	0.02	No	No
QLFS 2011:3	1.23	1.44	1.23	0.00	59.34	0.45	0.45	0.13	No	No
QLFS 2011:4	1.20	1.17	1.17	0.00	41.05	32.48	32.48	0.09	No	No
QLFS 2012:1	0.84	1.83	0.84	0.00	40.58	43.53	40.58	0.17	No	No
QLFS 2012:2	1.05	1.90	1.05	0.00	23.39	86.97	23.39	0.09	No	No
QLFS 2012:3	1.05	0.98	0.98	0.00	81.97	0.00	0.00	0.25	No	No
QLFS 2012:4	2.12	0.17	0.17	0.00	158.74	5.13	5.13	0.16	No	No
QLFS 2013:1	2.04	0.38	0.38	0.00	162.64	0.00	0.00	0.32	No	No
QLFS 2013:2	2.34	0.25	0.25	0.00	163.60	0.68	0.68	0.25	No	No
QLFS 2013:3	3.31	0.11	0.11	0.00	196.89	2.97	2.97	0.19	No	No
QLFS 2013:4	1.89	0.07	0.07	0.00	233.58	0.29	0.29	0.28	No	No
QLFS 2014:1	2.32	0.04	0.04	0.00	192.97	1.38	1.38	0.22	No	No
QLFS 2014:2	2.28	0.05	0.05	0.00	273.82	0.69	0.69	0.23	No	No
QLFS 2014:3	2.37	0.12	0.12	0.00	178.11	3.29	3.29	0.10	No	No
QLFS 2014:4	3.41	0.18	0.18	0.00	281.71	1.10	1.10	0.22	No	No
QLFS 2015:1	1.85	0.26	0.26	0.00	130.96	12.43	12.43	0.07	No	No
QLFS 2015:2	1.19	0.77	0.77	0.00	54.31	56.15	54.31	0.05	No	No
QLFS 2015:3	1.02	0.55	0.55	0.00	57.87	39.22	39.22	0.09	No	No
QLFS 2015:4	1.07	0.62	0.62	0.00	71.95	38.06	38.06	0.14	No	No

Note: Source: Authors own calculation based on Post-Apartheid Labour Market Surveys (PALMS) Note: Sample contains individuals aged between 15 and 65 years of age.

References

- Abadie, A., 2002. Bootstrap tests for distributional treatment effects in instrumental variable models. *J. Am. Stat. Assoc.* 97 (457), 284–292.
- Bhorat, H., Goga, S., 2013. The gender wage gap in post-apartheid South Africa: a re-examination. *J. Afr. Econ.* 22 (2), 1–22.
- Blau, F., 2012. *Gender, Inequality, and Wages*. Oxford University Press, Oxford.
- Blau, F.D., Kahn, L.M., 1996. International differences in male wage inequality: institutions versus market forces. *J. Polit. Econ.* 104 (4), 791–837.
- Blau, F.D., Kahn, L.M., 2017. The gender wage gap: extent, trends, and explanations. *J. Econ. Lit.* 55 (3), 789–865.
- Branson, N., Wittenberg, M., 2014. Reweighting South African National household survey data to create a consistent series over time: a cross-entropy estimation approach. *South Afr. J. Econ.* 82 (1), 19–38.
- Carrillo, P., Gandelman, N., Robano, V., 2014. Sticky floors and glass ceilings in Latin America. *J. Econ. Inequal.* 12, 339–361.
- Casale, D., 2004. What has the feminisation of the labour market “bought” women in South Africa? Trends in labour force participation, employment and earnings, 1995–2001. *J. Interdiscip. Econ.* 15 (3–4), 251–275.
- Casale, D., Posel, D., Mosomi, J., 2021. Gender and work in South Africa. In: Oqubay, A., Tregenna, F., Valodia, I. (Eds.), *The Oxford Handbook of the South African Economy*. Oxford University Press, Oxford. <https://doi.org/10.1093/oxfordhb/9780192894199.013.32>.
- Chatterjee, A., Czajka, L., Gethin, A., 2022. Wealth inequality in South Africa, 1993–2017. *World Bank Econ. Rev.* 36 (1), 19–36.
- Chatterjee, A., Czajka, L. and Gethin, A. (2020). ‘Estimating the distribution of household wealth in South Africa’. WIDER Working Paper 45/2019. Helsinki: UNU-WIDER.
- Dong, Y., Shen, S., 2018. Testing for rank invariance or similarity in program evaluation. *Rev. Econ. Stat.* 100 (1), 78–85.
- Fortin, N.M., Bell, B., Böhm, M., 2017. Top earnings inequality and the gender pay gap: Canada, Sweden, and the United Kingdom. *Labour Econ.* 47, 107–123.
- Fortin, N.M., Lemieux, T., 2000. Are women’s wage gains men’s losses? A distributional test. *Am. Econ. Rev.* 90 (2), 456–460.
- Firpo, S., Fortin, N.M., Lemieux, T., 2009. Unconditional quantile regressions. *Econometrica* 77 (3), 953–973.
- Goldin, C., 2014. A grand gender convergence: Its last chapter. *Am. Econ. Rev.* 104 (4), 1091–1119.
- Granger, C.W., Maasoumi, E., Racine, J., 2004. A dependence metric for possibly nonlinear processes. *J. Time Ser. Anal.* 25 (5), 649–669.
- Hinks, T., 2002. Gender wage differentials and discrimination in the New South Africa. *Appl. Econ.* 34 (16), 2043–2052.

- Kerr, A., Lam, D., Wittenberg, M., 2017. Post-Apartheid Labour Market Series, Version 3.2. DataFirst, Cape Town.
- Kerr, A., Wittenberg, M., 2015. Sampling methodology and fieldwork changes in the october household surveys and labour force surveys. *Dev. South Afr.* 32 (5), 603–612.
- Kerr, A., Wittenberg, M., 2021. Union wage premia and wage inequality in South Africa. *Econ. Model.* 97, 255–271.
- Leibbrandt, M., Woolard, I., Finn, A., & Argent, J. (2010). Trends in South African income distribution and poverty since the fall of apartheid.
- Li, Q., Maasoumi, E., Racine, J.S., 2009. A nonparametric test for equality of distributions with mixed categorical and continuous data. *J. Econom.* 148 (2), 186–200.
- Linton, O., Maasoumi, E., Whang, Y.J., 2005. Consistent testing for stochastic dominance under general sampling schemes. *Rev. Econ. Stud.* 72 (3), 735–765.
- Maasoumi, E., Heshmati, A., 2000. Stochastic dominance amongst Swedish income distributions. *Econom. Rev.* 19 (3), 287–320.
- Maasoumi, E., Wang, L., 2019. The gender gap between earnings distributions. *J. Polit. Econ.* 127 (5), 2438–2504.
- Maasoumi, E., and L. Wang (2014). The gender gap between earnings distributions.' Available @ <https://economics.stanford.edu/sites/g/files/sbiybj9386/f/maaswang6.pdf>.
- Mandel, H., Shalev, M., 2009. How welfare states shape the gender pay gap: a theoretical and comparative analysis. *Social Forces* 87 (4), 1873–1911.
- Mosomi, J. (2019a). 'Distributional changes in the gender wage gap in the post-apartheid South African labour market'. WIDER Working Paper 2019/17, Helsinki: UNU-WIDER.
- Mosomi, J., 2019b. Distributional Changes in the Gender Wage Gap in the Post-Apartheid South African Labour Market. Ph.D. thesis. University of Cape Town.
- Mosomi, J., and Thornton, A. (2022). Physical proximity and occupational employment change by gender during the COVID-19 pandemic (No. 2022/90). WIDER Working Paper.
- Muller, C., 2009. Trends in the Gender Wage Gap and Gender Discrimination among Part-Time and Full-Time Workers in Post-Apartheid South Africa. University of KwaZulu-Natal, Durban.
- Ntuli, M., 2007. Exploring Gender Wage Discrimination in South Africa, 1995–2004: A Quantile Regression Approach. University of Michigan, Gerald R. Ford School of Public Policy, Michigan.
- Oyenubi, A., 2023. The COVID-19 pandemic and polarisation of income distribution in South Africa. *Econ. Labour Relat. Rev.* 34 (1), 157–178.
- Piketty, T., Saez, E., 2013. Top incomes and the great recession: recent evolutions and policy implications. *IMF Econ. Rev.* 61 (3), 456–478.
- Posel, D., Gelo, D., Casale, D., Oyenubi, A., 2023. Sorting the gender earnings gap: heterogeneity in the South African labor market. *Fem. Econ.* <https://doi.org/10.1080/13545701.2023.2279227>.
- Puhani, P., 2000. The Heckman correction for sample selection and its critique. *J. Econ. Surv.* 14 (1), 53–68.
- Roberts, G., Schöer, V., 2021. Gender-based segregation in education, jobs and earnings in South Africa. *World Dev. Perspect.* 23, 100348.
- Shepherd, D., 2008. Post-Apartheid Trends in Gender Discrimination in South Africa: Analysis through Decomposition Techniques. Stellenbosch University, Stellenbosch.
- Wittenberg, M., 2017. Wages and wage inequality in South Africa 1994–2011: part 1—Wage measurement and trends. *South Afr. J. Econ.* 85 (2), 279–297.
- World Bank (2022), 'Gini index (World Bank estimate)'. World Bank DataBank Development Indicators. URL: <https://data.worldbank.org/indicator/SI.POV.GINI>.