7

WAVE 2

National Income Dynamics Study (NIDS) – Coronavirus Rapid Mobile Survey (CRAM)

Mind the gap:
Analysing the effects of South Africa's national lockdown on gender wage inequality

Robert Hill - University of Cape Town

Tim Köhler - University of Cape Town

30 September 2020

N.i.D.S.
NATIONAL INCOME DYNAMICS STUDY

CRAM
CORONAVIRUS RAPID MOBILE SURVEY 2020

Mind the gap: Analysing the effects of South Africa's national lockdown on gender wage inequality

Robert Hill, University of Cape Town[1]

Tim Köhler, University of Cape Town[2]

30 September 2020

Abstract

The COVID-19 pandemic has had severe and potentially long-lasting impacts on the South African economy since the onset of the national lockdown in March 2020, however, these effects have not been equally felt by all. Research has shown that the effects of the COVID-19 pandemic have been disproportionately felt by vulnerable groups, and disemployment effects have been concentrated amongst women. This paper aims to investigate whether wage inequality has deepened among job retainers, particularly along gender lines. Using the first two waves of the NIDS-CRAM data, this paper presents results for the conditional and unconditional gender wage gap at the mean, showing that women earned approximately 29% less than men per hour in February 2020, which expanded to approximately 43% less in June 2020. Monthly figures are more severe, with the gender wage gap estimated at approximately 30% in February 2020 and 51.6% in June 2020. Given the heterogeneity of the gender wage gap over the wage distribution, we proceed to use Recentred Influence Function (RIF) regressions to estimate the gender wage gap along the distribution. Although there is evidence that the gender wage gap exists at almost all points of the distribution in both periods, it has not changed uniformly across the wage distribution: in fact, evidence suggests that the monthly gender wage gap has deepened significantly for individuals below the 40th percentile of the distribution, but not for those at the top. Hourly wage gaps show no evidence of deepening across the distribution, suggesting that changes in wage inequality are driven by decreased working hours amongst women relative to men, perhaps due to being unable to work effectively from home, or due to increased childcare burdens during lockdown. This finding is robust to sample selection corrections carried out via a DiNardo, Fortin and Lemieux (DLF) reweighting.

1 Junior Researcher, Development Policy Research Unit; Assistant Lecturer, School of Economics; and PhD Candidate, University of Cape Town. Email: robert.hill@uct.ac.za.

2 Junior Researcher and PhD Candidate, Development Policy Research Unit, School of Economics, University of Cape Town. Email: tim.kohler@uct.ac.za.

Executive Summary

Introduction

Wave 1 of the NIDS-CRAM project showed that women have been disproportionately affected by the national lockdown. Of the estimated three million fewer people employed in April relative to February 2020, two million were women. However, less is known about the impact the national lockdown has had on those women who have remained in employment. This paper investigates this hypothesis by using representative survey data collected prior to and during the lockdown to construct comparable estimates of the evolution of the unconditional and conditional gender wage gaps. Given the evidence of heterogeneity across the earnings distribution (Mosomi, 2018), we also estimate the gender wage gap across the entire distribution of earners both before and during the lockdown, while accounting for selection into labour market participation, and make use of new information collected in Wave 2 of the NIDS-CRAM.

The unconditional gender wage gap was large and evident both before and during lockdown, while the average gap widened during the period. The gender wage gap – before accounting for any confounding factors like age, marital status, and occupation for example – was evident in both February and June 2020, regardless of whether monthly or hourly real wages are used. On average, our point estimates suggest that this unconditional gap was higher in June relative to February, although the differences are not statistically significant. This was driven by higher average wages of both men and women, however men's wages increased at a marginally higher rate (11%) than women's (9%). The average man's real monthly wage increased from R9 500 to R10 600 while that of the average woman increased from R5 600 to R6 100. Similar patterns exist once we account for working hours. Both changes in average monthly and hourly wages may be indicative of higher earners being more likely to remain employed during the lockdown period, rather than individual's actual wages increasing.

The unconditional gender wage gap widened most amongst poorer earners and remained particularly large amongst White individuals, those living in urban areas, and those with a tertiary qualification. The unconditional wage gap seems to have widened amongst the poorest 20% of earners. In June 2020, the average woman in this group earned about 83% of the average man's wage for a given hour of work – down from approximately parity in February. On the other hand, the unconditional gap seems to have narrowed amongst the richest quintile of earners. The gap remains largest and unchanged amongst White individuals. In both February and June, for every Rand earned by the average White man for a given hour of work, the average woman earned half. In both periods, the gap amongst those living in urban areas exceeded that of those in non-urban areas. Turning to education, the gap remains highest amongst those with a tertiary qualification. For a given hour of work in June, the average woman within this group earned 56 cents for every Rand earned by the average man, up from just 45 cents in February.

Even after accounting for several confounding factors, the gender wage gap was 46% - 73% higher in June 2020 relative to February on average. Of course, differences in wages between men and women can be partially explained by factors other than gender itself. After controlling for age, marital status, education, and occupation to name a few, we find that women earned 29% lower hourly wages than men before the lockdown in February, which deepened to 43% in June - indicative of an increase in the conditional gender wage gap of about 46%, at least on average. When we consider monthly rather than hourly wages but control for working hours, the increase in the conditional gender wage gap is even larger (73%). Although the conditional gender wage gap is statistically significant in each period, the changes in the gaps over the period are not statistically significant. Despite this, the differences in the magnitudes of our point estimates are compelling and as such, we interpret these changes as indicative of a widening gender wage gap at the mean.

The monthly gender wage gap in the bottom third of the wage distribution increased up to five-fold from its February size by June 2020. Overall, the only statistically significant changes in the gender wage gap occurred below the 40th percentile of the wage distribution, while there was little to no significant change above the 60th percentile. This is indicative of inequality-deepening effects between genders amongst the most vulnerable individuals in the economy. These results are robust to specifications that include industry-level effects for June 2020, as well as a reweighting technique used to adjust for sample selection. When estimating the hourly gender wage gap, there is no statistically significant difference between February 2020 and June 2020. This lack of a statistically significant result shows that the hourly gender wage gap has not necessarily changed between February and June 2020. Even when only considering point estimates of the February-to-June gender wage gap ratio, the results are milder for hourly wage inequality than they are for monthly inequality.

Taken together, changes in the monthly and hourly gender wage gaps speak to women decreasing their working hours relatively more than men. This could be as the result of women being employed in jobs that are less amenable to working from home, or as a result of women bearing a greater share of the increased childcare responsibilities that arose during the lockdown due to school closures.

How might these new insights assist policymakers? The widening of the size of the conditional gender wage gap amongst poorer earners is of particular concern for policymakers, given that this result speaks to deepening inequality amongst an already vulnerable group. This observation may have to do with the disproportionate incidence of childcare responsibilities on women during the pandemic, particularly amongst those who are unable to afford or access childcare. Policy which seeks to mitigate the adverse implications of such widening wage inequality ought to consider providing targeted income support to such workers at the bottom of the distribution. Several policy options are available. In their analysis of the NIDS-CRAM Wave 2 data, Köhler and Bhorat (2020) find that the distribution of personal and household-level receipt of the special COVID-19 Social Relief of Distress (SRD) grant, despite not being means-tested, is relatively pro-poor. However, the authors note that because other grant recipients are not eligible to apply for the grant, and because nearly 85% of all grant recipients are women, most recipients (two-thirds) of the COVID-19 SRD grant are men. Therefore, despite it being progressively targeted, this grant in its current form may not be appropriate in this context. However, considering the Child Support Grant is also progressively targeted and that about 98% of recipients (caregivers) are women, the pandemic-induced expansion of the grant may be an optimal mechanism to provide such income support. Policymakers ought to consider further extending this expansion of social assistance beyond October 2020.

State-subsidised childcare at schools may also assist in remedying the deepening gender wage gap that has arisen during lockdown. With schools reopening, women already have more opportunity to engage in labour market activities and increase their hours worked. However, if the state could provide after-school care for children, this may afford women the opportunity to further increase working hours, and help remedy the inequality that arose during the lockdown period. Careful structuring of this care could also benefit students through academic assistance, food provision, and ultimately better educational attainment outcomes.

1. Introduction

Wage inequality and discrimination on the basis of gender have been the subject of many empirical studies over the past few decades. Inequality in South Africa is already high, and progressive policies have been implemented by the state to address this. However, the onset of the coronavirus (COVID-19) pandemic and the subsequent national lockdown at the start of 2020 has had potentially devastating effects on inequality. In this paper, focus falls on the gendered impacts of the COVID-19 lockdown. Gender is an important factor in determining the economic impact of the pandemic. International literature suggests that, unlike previous recessions where men have borne the brunt of the economic downturn, this 'pandemic recession' is likely to disproportionately and persistently impact women (Alon et al., 2020). This is already clearly the case in South Africa, where initial research has shown that of the estimated three million fewer employed people in April relative to February 2020 as a result of the pandemic, two in every three were women (Casale and Posel, 2020; Ranchhod and Daniels, 2020).

However, although research has been conducted on the employment effects of the COVID-19 pandemic in South Africa, less research has been conducted on whether there have been inequality-deepening effects for those individuals who have managed to remain employed during the national lockdown. This paper aims to investigate the impact that the lockdown has had on gender wage inequality in South Africa for those individuals who have remained in employment during the period. Given international evidence that women have been found to take on greater shares of responsibility in the home relative to men during this period of working from home (Alon et al., 2020; Collins et al., 2020), it is our hypothesis that, even amongst those women who have remained employed, they are likely to have been more adversely affected by the onset of the COVID-19 pandemic relative to their male counterparts.

We make use of a comparable econometric specification using the first two waves of the National Income Dynamics Study: Coronavirus Rapid Mobile Survey (NIDS-CRAM) data, conducted from May to June and July to August 2020 respectively, to estimate the evolution of the unconditional and conditional gender wage gaps in South Africa. We use these data as two independent cross-sections that are broadly representative of the population aged 18 years and older in 2017. With this, we are able to construct estimates of the gender wage gap for a pre-lockdown period and compare them to estimates from during the lockdown to determine whether there have been any inequality-deepening impacts of the lockdown on inter-gender wages. We begin by considering the conditional and unconditional gender wage gap at the mean of the earnings distribution. We show that the unconditional gender wage gap was large and evident both before and during lockdown, while our point estimates suggest the average gap widened during the period regardless of whether monthly or hourly wages are used, although these differences are not statistically significant. This was driven by higher average wages of both men and women, however men's wages increased at a marginally higher rate than those of women. We emphasise that changes in average monthly and hourly wages may be indicative of higher earners being more likely to remain employed during the lockdown period, rather than individual's actual wages increasing. We further show that the unconditional gap widened most amongst the poorest 20% of earners and remained particularly large amongst White individuals, those living in urban areas, and those with a tertiary qualification. After accounting for several confounding factors through Mincerian-style regressions, we show that the gender wage gap was 46% to 73% higher in June 2020 relative to February on average. Although this change in point estimates is compelling, these differences are, again, statistically insignificant.

In order to obtain a more nuanced understanding of the impact of the pandemic on gender wage inequality, we utilise Recentred Influence Function (RIF) regressions in order to estimate the conditional gender wage gap at various points along the wage distribution in February and June 2020. Our estimates indicate that the conditional gender wage gap is indeed heterogeneous across the wage distribution, and there is evidence of a widening monthly gender wage gap for the bottom of the distribution, while there is no evidence for any widening gap at the top of the distribution. Using a reweighting technique to account for sample selection bias in the June 2020 sample, we find that the results are robust to sample selection corrections. This is indicative of a trajectory of deepening gender inequality amongst an already vulnerable group of individuals. As a result, this evidence can be used to protect those who are suffering most as a result of the COVID-19 pandemic.

Coupled with the fact that there is no indication of any deepening hourly wage gap, one can conclude that the driving force behind monthly wage inequality may be an adjustment in hours worked that has disproportionately impacted women. Potential reasons for this change are that women may be employed in jobs less amenable to remote working practices, or that they have disproportionately taken up the burden of childcare relative to men during the lockdown period. Either of these explanations would be consistent with findings from the international literature on the impact of the pandemic on female economic outcomes (Alon et al., 2020; Collins et al., 2020).

The remainder of the paper is structured as follows: Section 2 provides an overview of research conducted on the gender wage gap, both in general and during the COVID-19 pandemic. Section 3 details the data used for this study as well as the creation of new weights used to correct for selection into bracket responses in the data. Section 4 briefly outlines our adopted methodology and discusses the DiNardo, Fortin and Lemieux (DFL) reweighting procedure undertaken to correct for sample selection. Section 5 presents several descriptive statistics, while Section 6 presents the results of the econometric models. Finally, Section 7 provides several concluding remarks as well as policy considerations.

2. Literature Review

2.1. The gender wage gap: global and local evidence

Gender wage inequality has been the focus of a large body of literature, both within South Africa and abroad. This research has been mostly unanimous in concluding that the gender wage gap, although narrowing, is still a persistent feature of the global labour market. According to Wiechselbaumer and Winter-Ebmer (2005), early estimates of the gender wage gap in the international labour market began at approximately 65% in the 1960s and narrowed to approximately 30% by the late 1990s. Furthermore, in South Africa specifically, gender inequality, and in particular, the gender wage gap, has continued to narrow in the post-apartheid period (Mosomi, 2019; Posel and Casale, 2019). Mosomi (2019) estimated the South African gender wage gap at the mean of the wage distribution to have narrowed from approximately 40% in 1993 to approximately 16% in 2014. The gender wage gap at the median of the distribution has also decreased, but not to the same extent. In 1993, the gender wage gap at the median of the distribution was approximately 35%, while in 2015, it had decreased to approximately 23% (Mosomi, 2018). These estimates, using survey data, are slightly lower than those which use administrative data, where the gender wage gap is estimated to be approximately 35% in the South African formal sector (Bezuidenhout et al., 2019). Estimates of the South African median gender wage gap are relatively comparable with international estimates for the same time period: In 2009, full-time female workers in the US earned approximately 80 cents per dollar earned by male workers, indicating a gender wage gap of approximately 20% (Hegewisch et al., 2010; Blau and Kahn, 2017). The German gender wage gap is also at a comparable level, having been estimated to be approximately 20% (Antonczyk et al., 2010).

However, estimates of the gender wage gap at the mean or median of the distribution, while informative, can obscure important variation in wage inequality across the wage distribution. For example, Bhorat and Goga (2013) find that the gender wage gap is most pronounced (approximately 63%) at the 10th percentile of the distribution, but decreases to only approximately 7.2% by the 90th percentile. Although the reported size of the gender wage gap at different points along the South African wage distribution differs, the over-riding conclusion of heterogeneity in wage inequality across the wage distribution has been consistent. Ntuli (2007) shows that the gender wage gap has not consistently narrowed across the distribution. Rather, the narrowing of the mean gender wage gap was driven by decreasing inequality at the top and bottom of the distribution. Findings by Mosomi (2018) clearly support this narrative, showing stagnating inequality at the middle of the distribution with decreasing inequality at the top and bottom. This narrowing of the gender wage gap at the bottom of the distribution is likely driven by a combination of increased human capital characteristics and upward pressure on wages as a result of minimum wage legislation, particularly in the female-dominated domestic workers sector (Mosomi, 2018, 2019).

This heterogeneity of the gender wage gap across the wage distribution is not only a South African phenomenon, however. In the United States, Blau and Kahn (2017) find that the gender wage gap declined substantially more slowly at the top of the distribution than at the middle or bottom. As a result, the United States has experienced a widening of the gender wage gap at the top of the wage distribution. The German labour market has shown similar trends, with evidence of a shrinking gender wage gap only present at the bottom of the wage distribution, while wage inequality at the top has increased over time (Antonczyk et al., 2010).

Given that evidence presented in the literature provides a strong argument for heterogenous wage inequality across the wage distribution, we opt for a distributional analysis in this paper. By analysing the gender wage gap across the entire distribution of wages, rather than simply at the mean, we will be able to better understand the interaction between wages and employment dynamics that have occurred in the South African economy as a result of the national lockdown. This will provide a more nuanced platform from which to engage in policy discussions, as impacts on individuals at either end of the distribution will be hidden by simply estimating an average effect.

Studies on the gender wage gap, both locally and internationally, have provided a number of socio-economic characteristics that impact wage inequality. For example, the race of a worker has been found to be highly significant in correctly estimating the gender wage gap. Hinks (2002) found that the gender wage gap at the mean of the distribution is found to be highest amongst White individuals at approximately 40%, whilst amongst Coloured individuals, the gap is only estimated to be approximately 5%. Similarly, the age of workers is found to be a significant driver of the gender wage gap. Wage inequality between men and women is substantially lower for younger cohorts (Mosomi, 2019). The gender wage gap increases steadily over the course of an individual's lifetime; however, this is potentially explained by labour market interruptions as a result of childbirth for women (Budlender, 2019), or that women are more likely to be employed in occupations that provide limited room for real wage growth (Mosomi, 2019).

Education is a further factor that acts to narrow the gender wage gap, especially given the complementarities that arise between education and skills-biased technical change. Specifically, Mwabu and Schultz (2000) argue that the returns to higher education in particular are higher for women. In recent years, women have realised greater increases in human capital than men, and there has been a pattern of skills-biased technical change underway in the South African economy (Mosomi, 2019). Combined, these factors are thought to explain why education has played a large role in the narrowing of the South African gender wage gap (Mosomi, 2019). Skills-biased technical change has not only narrowed the gender wage gap in South Africa, but all around the world. The mechanisation of occupations that have a focus on manual or routine tasks has primarily occurred in male-dominated occupations, thus placing downward pressure on male wages and narrowing the gender wage gap in the United States (Yamaguchi, 2018). Evidence from Germany supports these findings, showing that the returns to labour market skills have risen over time (Antonczyk et al., 2010). Coupled with the fact that labour market skills that receive lowest returns are predominantly held my men, this could partially explain the narrowing of the gender wage gap in parts of the developed world (Yamaguchi, 2018).

Occupational segregation is a persistent cause of gender wage inequality, with female-dominated occupations generally presenting a higher gender wage gap than male-dominated occupations (Hegewisch et al., 2010; Hinks, 2002). In fact, according to a predictive model proposed by Hegewisch et al. (2010), a high-skilled occupation in the United States that is 100% female would pay approximately 46% less than one that is 100% male.[3] A similar finding is true for female-dominated industries when compared to male-dominated industries. There is a general decrease

3 The predicted wages for men and women in these hypothetical occupations are $1555 and $840, respectively.

in the gender wage gap as the proportion of male employment in the industry increases (Landman and O'Clery, 2020; Hegewisch et al., 2010). This finding holds in the South African context, and it is hypothesised that the reason for this has to do with compliance with the Employment Equity Act (No. 55 of 1998). In particular, because of legislation that forces South African firms to representatively hire female employees, it is necessary to entice female workers to enter and remain in male-dominated industries. The easiest way to accomplish this is through higher wages. Through this mechanism, the gender wage gap in male-dominated industries is forced downwards and wage inequality decreases (Landman and O'Clery, 2020).

2.2. Gender wage inequality and COVID-19

Evidence from the local and global literature has shown that the gender wage gap can be influenced by a number of socio-economic characteristics and trends. The COVID-19 pandemic has had a large impact on both the local and global economy, and studies have shown the disproportionate impact it has had on women in South Africa (Casale and Posel, 2020). As a result, it is likely that gender-based wage inequality will also be affected. Given that at the time of writing, much of the world is still struggling with the COVID-19 pandemic, this area of research is rather sparsely populated, particularly for the developing world.

The COVID-19 pandemic has produced an economic crisis quite different to any other in recent history, and as such, the effects of the pandemic on economic outcomes is not clear-cut. For example, Alon et al. (2020) report that the Global Financial Crisis of 2007/2008 disproportionately impacted male labour market outcomes, while the COVID-19 pandemic has quite clearly had a more severe impact on female labour market outcomes. One channel through which this disproportionate effect on women has been felt is working hours. In the United States, women with young children have reduced their working hours between four and five times more than fathers, leading to the gender gap in working hours growing by between 20 and 50% (Collins et al., 2020). The effects of these reductions in working hours may feed through to future labour market inequality as employers may choose to reward longer working hours with higher pay and, as a result, increase male wages disproportionately over female wages once again (Collins et al., 2020; Alon et al., 2020).

It is possible that inequality in labour market outcomes has been exacerbated by an inability to work effectively from home. In the United States, it was found that only 28% of men and 22% of women were employed in so-called tele-commutable occupations and able to work from home (Alon et al., 2020). This discrepancy in working conditions may lead to disproportionate job or pay losses for women, as they cannot meet the same obligations as before the pandemic. A similar result in the United Kingdom showed that women made up a greater share of employment amongst those sectors that needed to shut down during COVID-19 lockdown, thus disproportionately impacting women's ability to work, and ultimately, their wages during the pandemic (Blundell et al., 2020).

In the South African context, it is clear that women are still feeling the brunt of the COVID-19 lockdown. South Africa's national lockdown was implemented from the end of March 2020. Of the estimated three million fewer people employed in April relative to February 2020, women accounted for approximately two in every three fewer people employed (Casale and Posel, 2020). Using pre-crisis data, only 13.8% of workers have been estimated to be able to work from home (Kerr and Thornton, 2020). Considering these individuals are concentrated at the top end of the wage distribution, it is likely that wage inequality in South Africa is likely to increase as a result of the lockdown. Furthermore, with South Africa's lockdown-related workplace restrictions considered amongst the most stringent in the world (Gustafsson, 2020), impacts on wage inequality are likely to be more severe in South Africa than other comparable countries.

In this paper, we aim to estimate the impact of the national lockdown on gender wage inequality in South Africa by estimating the unconditional and conditional gender wage gap across the wage distribution. The following section of the paper describes the data that is available for use, as well as some of the corrections and manipulations performed in order to ensure that our estimates of the impact are accurate.

3. Data

3.1. The National Income Dynamics Study: Coronavirus Rapid Mobile Survey

This paper uses data from the first two waves of the National Income Dynamics Study: Coronavirus Rapid Mobile Survey (NIDS-CRAM), conducted from 7 May to 27 June and 13 July to 13 August 2020, respectively. The NIDS-CRAM is a broadly representative, panel, individual-level, and individual-based survey of approximately 7 000 South African adults, which will be repeated over several months as South Africa's national lockdown progresses. Conducted as a collaborative research project by several South African universities, the aim of the survey is to provide frequent, representative data on key socioeconomic outcomes in South Africa during the COVID-19 pandemic and national lockdown. The survey forms part of a broader study – of which this paper is part – which aims to inform policymaking using rapid, reliable research in the context of the COVID-19 pandemic. The survey instrument includes a wide array of questions on income and employment, household welfare, and COVID-19-related knowledge and behaviour.

In order for the NIDS-CRAM to be representative while simultaneously adhering to public health protocol and lockdown regulations, the mobile phone numbers of existing sample participants needed to be obtained. In this light, the NIDS-CRAM sample frame consists of individuals resident in South Africa aged 18 years or older at the time of fieldwork in April 2020 who were surveyed in Wave 5 of the National Income Dynamics Study (NIDS) conducted in 2017.[4] The NIDS is a nationally representative, panel, face-to-face, individual-level, household-based survey conducted approximately every two years from 2008 to 2017 and has followed the same 28 000 South African individuals over five waves. The NIDS-CRAM sample is a sub-sample of the NIDS Wave 5 sample and was drawn using a stratified sampling design. For more information on the NIDS-CRAM sampling design, the interested reader is referred to Ingle et al. (2020). Considering attrition across the first two waves of the NIDS-CRAM panel, of the 7 073 individuals who were successfully interviewed in Wave 1, 80.2% (or 5 676 individuals) were successfully interviewed in Wave 2. The attrition rate between the two waves is approximately 19%. Because individuals who attrit tend to be systematically different to individuals who remain in the panel, panel weights are used here for the NIDS-CRAM Wave 2 data to correct for non-random attrition.[5]

In order to examine the effects of South Africa's national lockdown on the country's gender wage gap, we required adequate data prior to and during the lockdown. The only representative survey conducted during the lockdown available at the time of writing was the first two waves of the NIDS-CRAM. The NIDS-CRAM Wave 1 data include information on individuals' earnings in both February (pre-lockdown) and April 2020 (one month into lockdown), where Wave 2 includes information on individuals' earnings in June 2020 (three months into lockdown). As such, although the February earnings data in Wave 1 is retrospective in nature, we choose to use it for our analysis of the pre-lockdown gender wage gap. For our analysis of the gender wage gap during South Africa's lockdown, we choose to use earnings data for June 2020. Although we can additionally conduct the analysis for April 2020, we choose to leave this for future work where we evaluate how gender wage inequality evolved from February to April to June 2020. It should, however, be noted that the NIDS-CRAM Wave 1 data do not include several key variables required to analyse the conditional gender wage gap (such as an individual's marital status and usual employment industry). Not controlling for these variables will likely result in an overestimated gender wage gap. The solution to overcoming this problem underlies our use of the NIDS-CRAM Wave 2 data, where data on these two variables are included. Given the longitudinal nature of the NIDS-CRAM, we are able to use information from Wave 2 in place of missing information in Wave 1. This then assumes these variables were time-invariant across the two waves. Given the period between the two waves was one to two months,

4 Sample members in the NIDS could be Continuing Sample Members (CSMs) or Temporary Sample Members (TSMs). CSMs were interviewed in every wave of the NIDS, whereas TSMs were interviewed in a given wave only if they were a co-resident of a CSM.

5 The results of a multivariate probit model of the probability of being re-interviewed in Wave 2 conditional on a range of Wave 1 characteristics including race, sex, age categories, province, urban location, an indicator that household income was missing, an indicator that the individual was employed at the time of the Wave 1 interview, and NIDS Wave 5 household per capita income quartiles (not shown here) showed that attrition was statistically significantly higher among Coloureds and Indians, urban dwellers, the employed, those with missing household income and individuals in the top per capita income quintile in NIDS Wave 5.

we believe this is not an unreasonable assumption to make, at least as far as marital status goes.[6] Considering our analysis focuses on heterogeneity in wages conditional on employment, we restrict our within-wave samples to working-age adults (18-64 years) who were employed at the time of the relevant reference period of their earnings (that is, February and June 2020).[7] In essence then, our analysis can be regarded as pooled cross-sectional.[8]

A number of representative surveys conducted prior to the lockdown are available to examine the pre-lockdown gender wage gap; however, not all of them include many of the variables we require for this analysis. For instance, Statistics South Africa's (StatsSA) latest Quarterly Labour Force Survey (QLFS) was conducted in the first quarter of 2020 just prior to the onset of South Africa's national lockdown. However, the public release of the QLFS does not contain any earnings data – a key variable of interest. StatsSA release the earnings data collected in the QLFS in a separate, annual publication called the Labour Market Dynamics in South Africa, the latest of which that is publicly available is for 2018. However, it is not ideal to use this data to analyse the effects of the national lockdown on the gender wage gap. This is because between 2018 and June 2020, many other factors may have affected gender wage inequality in South Africa other than the lockdown (such as economic recessions) and as such, attributing any change in the gap to the lockdown would be spurious. Additionally, there are important differences in sampling and questionnaire design between the QLFS and NIDS-CRAM which may render any comparisons unreliable.

3.2. How representative is the NIDS-CRAM data? Caveats to consider

We believe our estimates come with important caveats due to unavoidable imprecision which render them still approximations. We discuss these caveats here. First, it is important to note that because the NIDS-CRAM sample is drawn from a representative sample of individuals in NIDS Wave 5 (conducted in 2017), the weighted estimates are not necessarily representative of the South African adult population in 2020. Although post-stratification weights make the NIDS Wave 5 sample representative of the South African population in 2017, the weighted NIDS-CRAM estimates are only representative of the outcomes in 2020 of those aged 15 years and older in 2017 who were followed up 3 years later – hence 'broadly' representative. Second, it is not unusual to observe disparities between estimates of the NIDS and other household surveys such as the QLFS. This is because NIDS was always designed to be a panel survey representative of the population in 2008, and as such, factors such as selective migration from the sample over time are not accounted for. Since the NIDS-CRAM sample is drawn from the NIDS Wave 5 sample, this characteristic will continue into NIDS-CRAM. We emphasise that the reader keeps these concerns in mind throughout the paper.

There remain, however, important advantages of the NIDS-CRAM data that make it incredibly valuable for understanding the current context in South Africa. First, because the survey is designed as a panel, it can provide a substantial amount of information about the dynamics of sampled individuals as the pandemic and lockdown unfolds. At the time of writing there is no comparable existing dataset which can be used to analyse these dynamics. Second, despite unavoidable comparability issues with the NIDS-CRAM and other surveys, internal validity and comparisons over time for the sample are not issues of concern. Many of the operational challenges experienced by the NIDS-CRAM survey will almost certainly be experienced by other surveys being conducted during this period.

3.3. Earnings in the NIDS-CRAM data: Adjusting for outliers and selection into bracket response

6 It should be noted that despite the addition of these important variables, we were not able to control for several others not available, such as union membership.

7 This lower age bound of 18 years, as opposed to the standard lower bound of the working-age population of 15 years, is used because younger individuals were not sampled in the NIDS-CRAM.

8 We considered restricting our sample to those employed in all time periods; however, doing so may result in biased estimates given that employment outcomes in one period may be endogenous to an individual's characteristic(s) in another period. This is important considering that employment loss between February and April 2020 was more prevalent within several groups such as women and individuals at the lower end of the earnings distribution in February 2020.

The questionnaire items used to collect earnings data in the NIDS-CRAM differ to those used in more detailed household surveys such as the NIDS. In the NIDS, several items exist which seek to gather information on an individual's labour market earnings and other sources of income (such as that from an individual's primary and secondary jobs, casual work, self-employment). On the other hand, in the NIDS-CRAM Waves 1 and 2 data, much fewer items exist. Specifically, individuals were asked what their last take-home pay and/or profit for February, April, and June 2020 was. For the latter two reference periods, the data varies by respondent's reported frequency of payment (daily, weekly, fortnightly, or monthly). We interpret these items collectively as an inclusive wage variable. It is however reassuring that in an assessment of the quality of the NIDS-CRAM Wave 1 data, Ardington (2020) shows that the weighted distribution of earnings is plausible given its similarity to distributions in other surveys (such as the NIDS Wave 5 and General Household Survey 2018). We convert all wages reported in monetary amounts to a monthly frequency.[9] For earnings in all three reference periods, individuals were asked to report an actual monetary (Rand) amount after taking deductions into account. If they were not willing, they were asked to report which bracket their income lies in. There are several concerns with the use of such raw income data. First, the data may be contaminated by implausible or unreliable values, i.e. outliers. Second, there may be selection into responding with bracket information, or into refusing to provide any income information. Simply ignoring bracket responses incorrectly ignores responses that may come from the top end of the income distribution. For instance, in an analysis of South African household survey data, Wittenberg (2017) shows that individuals who respond with brackets do tend to have higher incomes. Thus, any analysis which does not address these concerns beforehand may produce biased estimates. We adopt several statistical techniques to address these issues and adjust raw earnings in the NIDS-CRAM Waves 1 and 2 data.

3.1.1 Outliers

First, outlier values are identified and coded as missing by using the "extreme studentised regression residuals" approach as advised by Wittenberg (2017). This is done by estimating a Mincerian-style Ordinary Least Squares (OLS) regression of the logarithm of nominal monthly earnings on a vector of observable covariates and identifying outliers as those observations with absolute residuals in excess of four.[10] This adjustment resulted in just three February 2020 earnings values being coded as missing in the NIDS-CRAM Wave 1 data, and two June 2020 earnings values in the NIDS-CRAM Wave 2 data.

3.1.2 Selection into bracket response

Second, we address selection into responding with bracket information through a reweighting technique. A simpler approach would be to impute and assign the within-bracket means, medians, or midpoints to individuals who reported bracket responses. However, one of the major drawbacks of such imputation is that it produces artificial spikes in the data at the imputation values, which would affect the percentiles – an important aspect of our distributional analysis here. Instead, we construct bracket weights, similar to those constructed in the Post-Apartheid Labour Market Series (PALMS) version 3.3 dataset. These are calculated as the inverse of the probability of an actual monetary (Rand) response in a particular bracket in a particular wave (NIDS Wave 1 or Wave 2), multiplied by the sampling weight for each individual. This process weights up individuals whose reported incomes are in brackets where the proportion of actual monetary responses are lower, relative to brackets where such response is high.[11] To generate these weights, we first needed to generate consistent brackets, considering that in the NIDS-CRAM Wave 2 data the bracket categories differ by the frequency of payments (weekly, fortnightly, and monthly). More information pertaining to this harmonisation process is provided in *Tables A1* and *A2* in the Appendix.

9 To do so, we assume 4.33 weeks per month and use self-reported data on individuals' average work week lengths measured in days.

10 The covariates we use here include age, age squared, self-reported gender and race, area (traditional, farms, or urban), province of residence, highest level of education, marital status, usual occupation, usual industry, type of contract (written or verbal), and weekly hours worked.

11 For example, if we observe 95% of individuals within the bracket R1 000 – R2 000 gave actual Rand responses, then these individuals will get revised weights equal the sampling weight divided by 0.95. On the other hand, individuals within the bracket R20 000 – R30 000 where 35% gave actual Rand responses will get revised weights equal the sampling weight divided by 0.35. The latter will be weighted up relative to individuals in the lower bracket.

The outcome of our reweighting process is summarised by the unweighted and weighted (with sampling and bracket weights) distribution of wages in *Figure A1* in the Appendix. In both waves, the weighted (sampling weight or bracket weight) distributions are notably shifted to the right relative to the unweighted. The sampling weighted February 2020 distribution (mean of R7 285.94) is substantially different from the bracket weighted distribution (mean of R7 607.05). Similarly, the sampling weighted June 2020 distribution (mean of R8 157.08) is substantially different from the bracket weighted distribution (mean of R8 779.92). These observed differences between the sampling and bracket weighted distributions are attributable to the varied likelihoods of responding with an actual monetary (Rand) amount across the distribution (see *Table A2*). Unless indicated otherwise, all estimates for all periods are weighted using these computed bracket weights while accounting for the NIDS-CRAM complex survey design. Lastly, it is important to note that his reweighting approach does not do anything about those who refuse to answer or who otherwise have missing data; it only corrects for bracket responses.[12]

After these adjustments, our final sample consists of 2 590 employed, working-age individuals with non-missing monthly wage data in February 2020 (78.1% of the working-age employed sample of 3 316 individuals) and 1 738 in June 2020 (78.75% of the working-age employed sample of 2 207 individuals). All earnings data were inflated to July 2020 Rands ($1 was approximately R16.50 at the time of writing). In our analysis, we focus on both real hourly and real monthly wages seeing as a considerable proportion of the NIDS-CRAM sample report zero working hours, despite being employed.[13]

4. Method

This section very briefly discusses the method used for estimating the unconditional and conditional gender wage gaps in this paper. Although there are a number of methods available, the choice of method was informed by a combination of the best practice in the available literature and practicality of implementation given the size of the sample in the first two waves of the NIDS-CRAM.

First, we estimate the unconditional and conditional gender wage gaps in both February and June 2020 at the mean through Mincerian-style regressions. That is, we employ Ordinary Least Squares (OLS) to regress the natural logarithm of real hourly (or monthly) wages on a vector of observable covariates, including a binary indicator for women. This allows us to estimate the evolution of the conditional gender wage gap at the mean; that is, the percentage difference between the real hourly or monthly wages of men and women on average, while accounting for variation in wages induced by variation in other characteristics. Our estimate of interest is, of course, the coefficient on the binary indicator for women.

After estimating the conditional gender wage gaps at the mean of the wage distributions in February and June 2020, we seek to analyse the gap across the entire distribution in both periods. The econometric method utilised in this paper for this purpose is that of Recentred Influence Function (RIF) regressions, as proposed by Firpo et al. (2009). The RIF regression method essentially allows for the marginal effect of a change in an explanatory variable on the dependent variable to be estimated at each of a number of specified quantiles of the unconditional distribution of the dependent variable (Firpo et al., 2009). In other words, the coefficients from a RIF regression at the τ^{th} quantile can be interpreted as the marginal effect of a change in x_i on y at quantile τ. Estimation of a RIF regression relies heavily on the influence function, defined as $IF(Y; v, F_Y)$, where Y is the dependent variable of interest; v is the distributional statistic of interest in the influence function

12 In this paper, we chose not to impute wage values for such observations because furloughed workers are very imprecisely defined, and there is a considerable amount of them in the data due to the national lockdown. Consequently, imputations would be biased for furloughed workers in particular, and therefore biased overall for all income imputations. We thank Reza Daniels (UCT) for this valuable input. This work may be pursued in the future, subject to data availability in future waves of the NIDS-CRAM.

13 These individuals may be regarded as "furloughed"; that is, they are temporarily absent from employment due to lockdown restrictions. This is an important group to consider in any labour market analysis during a national lockdown. As expressed by Ranchhod and Daniels (2020), previous conventions and definitions of employment may not necessarily be appropriate in the context of the COVID-19 pandemic given the nature of temporary absences and furloughed workers.

in this case, the quantile; and F_Y is the unconditional distribution of Y. To produce a recentred influence function, one simply adds the influence function to the distributional statistic of interest. In other words, given that the functional form of the quantile influence function is known, the recentred influence function for the τ^{th} quantile of the distribution, q_τ, is defined as follows:

$$RIF(Y,\ q_\tau,\ F_Y) = q_\tau + \frac{\tau - \mathbb{1}[Y \leq q_\tau]}{f_Y(q_\tau)} \tag{1}$$

The regression estimation simply uses this newly defined RIF of Y_i, estimated at quantile q_τ as the dependent variable in an OLS regression. This leads to a regression model of the following form to be estimated:

$$RIF(Y,\ q_\tau,\ F_Y) = \alpha_\tau + \beta_\tau female_i + \gamma_\tau X_i + \varepsilon_i \tag{2}$$

In the above model, the dependent variable is either the log of monthly wages or the log of hourly wages; the matrix of individual-level covariates, X_i, includes variables such as race, marital status, home language, occupation and education level, amongst others. The coefficient β_τ is the point estimate of the gender wage gap at the τ^{th} quantile, which is the estimate of primary interest to this study.

One particular concern regarding the estimation of the gender wage gap is concerns around endogeneity of estimates due to the selection of individuals into labour force participation. Mwabu and Schultz (2000) find that women are significantly less likely to participate in the labour market than men, which introduces selection bias into the estimation of the gender wage gap. A number of studies have attempted to correct for this bias by estimating a two-stage Heckman selection model and controlling for the inverse Mills ratio in their subsequent regression estimates (Ntuli, 2007; Hinks, 2002; Mwabu and Schultz, 2000). However, in all studies, the selection term remained insignificant, indicating that controlling for sample selection did not substantially improve the estimates produced.

Even though studies have found selection effects to be insignificant, we suspect that this is unlikely to be the case here. As a result of the national lockdown, many individuals lost their jobs, however, these job-losers were not a random sample of the employed; rather, those individuals who lost their jobs were disproportionately concentrated amongst the more vulnerable and lower-earning groups in South Africa (Ranchhod and Daniels, 2020; Casale and Posel, 2020). As a result, it is important to account for the changes in the characteristics of the employed population between February 2020 and June 2020.

In the absence of a valid instrument to control for selection in a Heckman two-stage model, we opt to make use of the DiNardo, Fortin and Lemieux (hereafter DFL) reweighting technique to create a hypothetical distribution for June 2020 wage earners that matches the distribution of characteristics in the February 2020 wage-earner population (DiNardo et al., 1996). This technique has been used previously with the NIDS-CRAM data to investigate poverty incidence by Jain et al. (2020). The DFL reweighting procedure essentially entails adjusting sample weights for June 2020 by a factor θ, defined as follows:

$$\theta = \frac{Pr(T = Feb \mid X) Pr(T = June)}{Pr(T = June \mid X) Pr(T = Feb)} \tag{3}$$

These components are relatively simple to estimate from the data: the unconditional probabilities are simply the probability of an observation in the pooled sample being from February 2020 or June 2020,[14] while the conditional probabilities are estimated from a binary choice model with a dependent variable equal to 1 if the observation is from February 2020, and 0 if it is from June 2020. The covariates in X capture characteristics of individuals that we expect may differ between wage earners in the two periods, such as race, gender, occupation, child cohabitation status, and others.

14 Or, equivalently, the probability of being in the NIDS-CRAM Wave 1 and Wave 2 samples, respectively.

Figure A2 in the Appendix plots the wage distributions for February 2020, June 2020, and the reweighted June 2020 sample that has the same characteristics as the February 2020 sample. The hypothetical June distribution lies noticeably to the right of the real June distribution, indicating that there has been a selection effect at play to arrive at the June 2020 sample. As a result, the use of the DFL reweighting technique to control for selection is justified in this case. In essence, the DFL reweighting procedure is equivalent to an inverse probability weighting (IPW) procedure, which is commonly used to weight regression analysis in the programme evaluation literature (Elder et al., 2015). To this end, we rerun the June 2020 regressions as specified in Equation (2), above, but using the adjusted DFL weights as regression weights.

In the following section, we present a brief overview of the structure of the data through several descriptive statistics. This will provide the background for the econometric estimation of the gender wage gap undertaken in Section 6.

5. Descriptive statistics

In both February and June 2020, the unconditional gender wage gap is statistically significant and evident, regardless of if monthly or hourly real wages are observed. When we consider changes in monthly wages over the period, our point estimates suggest that the unconditional gap has increased at the mean, however these differences are not statistically significant. On the other hand, once we account for working hours, our point estimates suggest that the unconditional gap seems to have decreased, given that the average woman's hourly wage increased at a greater rate than the average man. The increase in the latter was not statistically significant; however, that of the average woman's was. *Figure 1* presents the average real monthly and hourly wages for men and women in February and June 2020. The average man's real monthly earnings increased (R9 548 to R10 614) while that of the average woman also increased (R5 639 to R6 134). However, despite both increasing, our point estimates still suggest a widening of the unconditional gap given that the former rate of increase (11%) exceeds the latter (9%), even if marginally. Once we account for working hours, the wages of both men and women have also risen on average, from R69 to R86 for men and R40 to R61 for women. Both changes in average monthly and hourly wages may be indicative of higher earners being more likely to remain employed during the lockdown period, and not individual's actual wages increasing.

Figure 1: Absolute mean real monthly and hourly wage by gender in February and June 2020

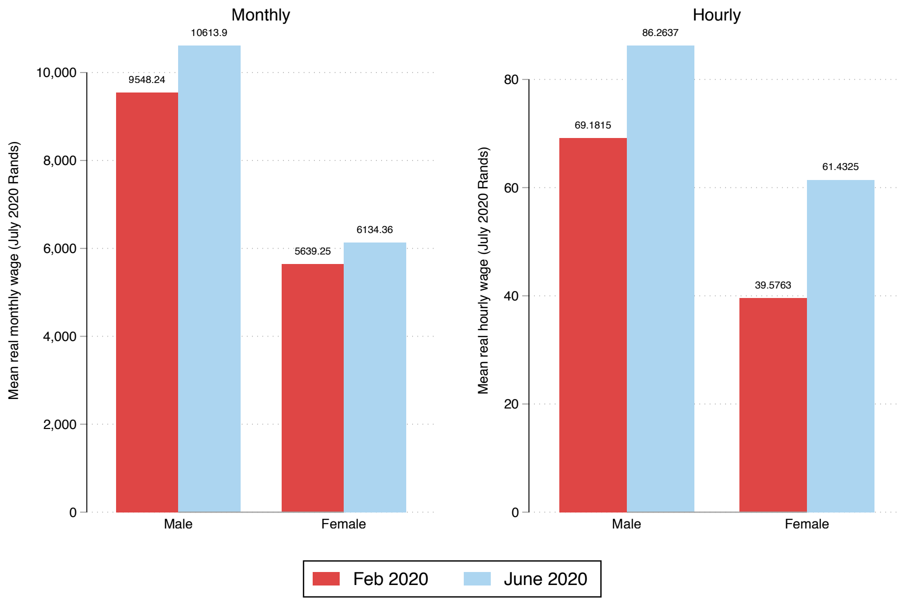

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes: [1] Within-wave samples restricted to employed individuals aged 18-64 years. [2] Estimates are weighted using computed bracket weights and account for complex survey design. [3] Wages inflated to July 2020 Rands.

To examine these changes in more detail, *Figure 2* presents the distributions of real monthly and hourly wages for men and women in February and June 2020. The shifts in distributions are substantial for both men and women. Considering the evolution of real monthly wages, it appears that the increase in wages for men was driven by a reduction in the number of poorer earners towards the bottom of the distribution and an increase in the number of richer earners from the middle towards the top. This is indicative of higher earners being more likely to remain employed during the lockdown period. For women on the other hand, the increase in wages seems to be driven also by a reduction in the number of poorer earners but also an increase in the number of earners in the middle of the distribution. Once we account for working hours, we see similar patterns for men. Male earners at the 90th percentile earned R161 per hour in February but R208 per hour in June. For women, most of the hourly wage distribution has shifted to the right, with the most notable changes at the top end. In February and June 2020 respectively, female earners at the 75th percentile earned R42 and R71, at the 90th percentile R108 and R129, and the 99th percentile R217 and R476. Again, this likely reflects a selection issue; that is, higher-earning men and women were far more likely than their lower-earning counterparts to remain employed during the national lockdown.

Figure 2: Distribution of log real monthly and hourly wages by gender in February and June 2020

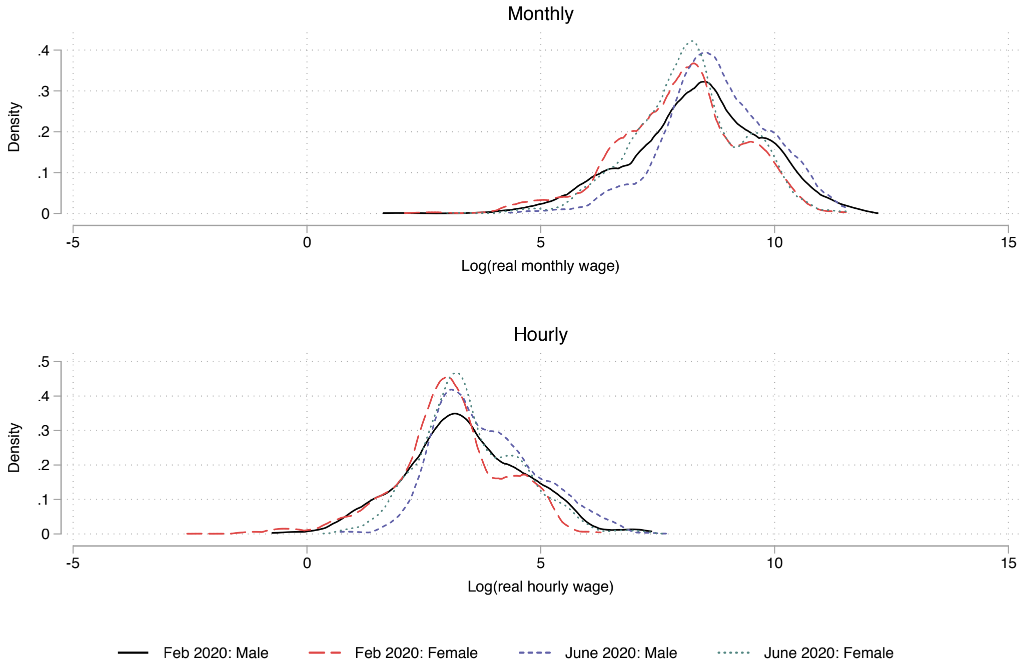

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes: [1] Within-wave samples restricted to employed individuals aged 18-64 years. [2] Estimates are weighted using computed bracket weights. [3] Wages inflated to July 2020 Rands.

It is worth investigating how the unconditional gap has changed within and between varied demographic groups. *Table 1* presents the mean real hourly wages for men and women in February and June 2020 as well as the computed inter-group unconditional wage gaps; that is, the ratio of the average women's real hourly wage relative to that of men's in a given month. We investigate these gaps between several demographic groups, including self-reported racial population group, age group, highest level of education, and usual occupation to name a few. The table suggests that the overall unconditional gender wage gap seems to have narrowed between February and June 2020, although the gap is still evident and significant in both periods. This however excludes individuals whose reported weekly working hours in June 2020 were zero, despite earning a positive wage. Overall, for every Rand earned by the average man for a given hour of work, the average woman earned 57 cents in February, but 71 cents in June. This varies, however, once we consider variation across the wage distribution. The unconditional wage gap seems to have widened amongst the poorest 20% of earners. In June, the average woman in this group earned about 83% of the average man's wage for a given hour of work – down from approximately parity in February. Amongst the richest quintile of earners, however, the unconditional gap seems to have narrowed. Within this group, the average woman earned 61 cents for every Rand earned by the average man for a given hour of work in February, up to 90 cents in June. Across the middle of the distribution, the gap is statistically unchanged.

Table 1: Mean real hourly wage across select covariates by gender, February and June 2020

	NIDS-CRAM Wave 1			NIDS-CRAM Wave 2		
	February 2020			June 2020		
	Male	**Female**	**Ratio**	**Male**	**Female**	**Ratio**
Overall	69.18	39.58	0.57	86.26	61.43	0.71
Real hourly wage quintile						
1	4.73	4.86	1.03	10.82	9.01	0.83
2	13.74	14.31	1.04	20.25	20.49	1.01
3	24.17	24.44	1.01	32.70	31.94	0.98
4	47.82	47.78	1.00	65.43	69.59	1.06
5	221.60	134.67	0.61	257.66	232.45	0.90
Race						
Black African	53.30	28.79	0.54	72.01	47.67	0.66
Coloured	36.35	45.01	1.24	58.99	60.79	1.03
Indian/Asian	54.37	34.18	0.63	54.32	282.90	5.21
White	195.68	98.27	0.50	199.28	99.58	0.50
Age group						
18-34	44.28	33.02	0.75	63.02	50.58	0.80
35-49	66.95	42.68	0.64	82.98	60.08	0.72
50-64	142.50	49.27	0.35	150.87	90.78	0.60
Area						
Urban	73.97	41.13	0.56	97.81	63.40	0.65
Non-urban	39.35	29.14	0.74	57.80	57.89	1.00
Education						
Up to primary	19.21	16.97	0.88	52.30	60.44	1.16
Up to secondary	29.90	20.30	0.68	59.17	45.80	0.77
Complete secondary	43.43	27.20	0.63	56.44	37.49	0.66
Tertiary	148.49	66.13	0.45	153.91	86.00	0.56
Occupation						
Managers	183.88	64.05	0.35	204.94	143.12	0.70
Professionals	172.54	82.91	0.48	186.37	93.85	0.50
Technicians and associate professionals	91.47	71.64	0.78	74.76	28.50	0.38
Clerical support workers	30.92	37.34	1.21	55.47	107.21	1.93
Service and sales workers	38.51	19.14	0.50	36.58	29.60	0.81

Skilled agricultural, forestry and fish	20.27	12.54	0.62	20.49	21.80	1.06
Craft and related trades workers	27.50	18.95	0.69	102.98	50.09	0.49
Plant and machine operators	50.95	20.61	0.40	49.08	45.73	0.93
Elementary occupations	22.04	17.54	0.80	45.81	28.07	0.61
Industry						
Private households	.	.	.	59.16	27.05	0.46
Agriculture, hunting, forestry and fish	.	.	.	24.19	24.09	1.00
Mining and quarrying	.	.	.	128.49	62.24	0.48
Manufacturing	.	.	.	61.81	50.94	0.82
Electricity, gas and water supply	.	.	.	47.75	75.70	1.59
Construction	.	.	.	76.30	38.23	0.50
Trade	.	.	.	47.86	67.55	1.41
Transport, storage and communication	.	.	.	54.30	194.19	3.58
Financial services	.	.	.	134.42	57.95	0.43
Community, social and personal services	.	.	.	116.40	79.76	0.69

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes:
1. Within-wave samples restricted to employed individuals aged 18-64 years.
2. Estimates are weighted using computed bracket weights after accounting for complex survey designs.
3. All hourly wages inflated to July 2020 Rands.

We observe significant heterogeneity in the evolution of the unconditional gender wage gap across self-reported racial population groups. Amongst Black African individuals, the gap is still substantial in both February and June, however, it seems to have narrowed, from 0.54 to 0.66. On the other hand, amongst Coloured individuals, in February women earned R1.24 for every Rand earned by men for a given hour of work on average. This reduced to approximately parity by June. Due to the small sample size, we prefer not to infer anything about the unconditional gender wage gap amongst Indian/Asian individuals. Lastly, the gap remains largest and unchanged amongst White individuals. In both February and June, for every Rand earned by the average White man for a given hour of work, the average women earned half. Moving onto varied age groups, our point estimates suggest that the gap has narrowed slightly but is still evident for all groups, especially one in particular: those aged between 50 and 64 years. Amongst this group, the gap narrowed from 0.35 in February to 0.60 in June: a 71% reduction.

We now consider the unconditional gender wage gaps across the remaining demographic groups included in *Table 1*. In both February and June, the gap amongst those living in urban areas exceeds that amongst those in non-urban areas. However, for individuals in both areas, the gaps seem to have narrowed, from 0.56 to 0.65 amongst individuals in urban areas and 0.74 to parity amongst individuals in non-urban areas. Considering varied levels of education, the gap seems to have narrowed for all groups, but remains highest amongst individuals who have a tertiary qualification. Notably, the average woman within this group earned 56 cents for every Rand earned by the average man for a given hour of work in June, up from just 45 cents in February. The evolution of the gap across individuals' usual occupations varied considerably. For instance, Managers saw their gap reduce from 0.35 to 0.70, whereas Technicians and associate professionals saw their gap widen from 0.78 to 0.38. Turning to industries, we are unable to observe the change in the gap considering this data does not exist in the NIDS-CRAM Wave 1 data.[15] However, we do observe that in June, the gap was particularly large amongst individuals working in the Private households (0.46), Mining and quarrying (0.48), Construction (0.50), Community, social, and personal (CSP) services (0.69), and most notably those working in Financial services (0.43).

As expressed in the beginning of the analysis in this section, it is important to interrogate changes in the unconditional gender wage gap across the entire wage distribution, as opposed to just at the mean. In this light, *Table 2* presents several inter- and intra-gender wage inequality statistics in February and June 2020. It is clear that the real hourly wages of both men and women were higher in June relative to February across the distribution, with the exception of male earners at the 99th percentile. The extent of these changes, however, varied. For instance, where the average hourly wage increased by 25% for men, it increased by 55% for women. The increase in mean wages amongst men seems to be driven by relatively greater increases from the middle towards the bottom of the distribution. Male earners at the 25th percentile saw 68% higher wages in June relative to February, as opposed to 30% higher wages for those at the 90th percentile. Amongst women on the other hand, the increase in mean wages was driven by relatively greater increases at the top of the distribution, where earners at the 99th percentile saw 120% higher wages in June relative to February.

15 Although given the longitudinal nature of the NIDS-CRAM we are able to impute a given respondent's industry in Wave 1 by using their reported industry in Wave 2, we prefer not to do so, considering that there may be some individuals who transitioned between industries during the national lockdown.

Table 2: Inter and intra-gender hourly wage inequality statistics in February and June 2020

	NIDS-CRAM Wave 1		NIDS-CRAM Wave 2	
	February 2020		June 2020	
	Male	Female	Male	Female
Real hourly wage (July 2020 Rands)				
p25	12.05	11.92	20.24	16.36
Mean	69.18	39.58	86.26	61.43
Median	26.78	21.09	37.04	26.79
p75	68.28	42.08	83.33	71.43
p90	160.66	108.45	208.33	129.25
p99	819.38	216.89	595.24	476.19
Wage inequality				
Gini coefficient	0.66	0.56	0.60	0.61
Theil's T index: GE(1)	0.95	0.57	0.71	0.79
Atkinson index (e=1)	0.58	0.46	0.48	0.48
90:10 ratio	33.33	24.00	15.75	15.51
50:10 ratio	5.56	4.67	2.80	3.21
90:50 ratio	6.00	5.14	5.62	4.83
Bottom 50% wage share (%)	9.12	13.92	11.60	12.30
Middle 40% wage share (%)	16.50	22.04	19.04	18.62
Top 10% wage share (%)	53.15	41.42	47.07	48.93
Top 1% wage share (%)	18.07	8.30	10.02	16.28

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes:
1. Within-wave samples restricted to employed individuals aged 18-64 years.
2. Estimates are weighted using computed bracket weights, and all real hourly wages are additionally estimated after accounting for the complex survey design.
3. All hourly wages inflated to July 2020 Rands.

Intra-gender wage inequality can also be informative. Three different inequality statistics – the Gini coefficient, Theil's T index, and Atkinson's index – all exhibit analogous patterns: from February to June 2020, wage inequality decreased amongst male earners, but increased amongst female earners. Interestingly, in February wages amongst men were more unequally distributed relative to women, however this pattern seems to have reversed in June, albeit to a small extent. Changes in wage ratios and shares allow us to examine which parts of the distribution drove the above observed changes in intra-gender wage inequality. We estimate that amongst men in February, earners at the 90th percentile earned 33 times more than those at the 10th percentile, and 6 times more than the median male earner. This latter statistic was similar in June, whereas the aforementioned 90:10 ratio decreased by 53%. Female earners also saw a large decrease in their 90:10 ratio, albeit by a smaller magnitude (35%). These patterns are reflected by the observed changes in wage shares. The shares of the bottom 50% and middle 40% of earners have not changed significantly. On the other hand, the top 10% of male earners accounted for 53% of all wages earned in February, decreasing to 47% in June 2020. The opposite occurred amongst female earners, where the richest 10% of earners accounted for 41% of all wages earned in February, increasing to nearly half (49%) in June.

6. Model results

Our analysis up to this point has only included examining variation in real monthly and hourly wages between men and women across the distribution in an unconditional environment. Of course, the observed variation in inter-gender wages can be explained by factors other than gender itself. We now turn to examining changes in the gender wage gap while controlling for possible confounding variables – i.e. a conditional environment. First, we conduct this analysis at the mean of the distribution; that is, we estimate a Mincerian-style OLS regression for February 2020 and June 2020 separately. As discussed in Section 4, we regress the natural logarithm of real hourly (or monthly) wages on a vector of observable covariates, including a binary indicator for women. This allows us to estimate the evolution of the conditional gender wage gap at the mean; that is, the percentage difference between the real hourly or monthly wages of men and women on average, while accounting for variation in wages induced by variation in other characteristics. Here, we specifically control for age, race, highest level of education, main occupation, type of employment contract, marital status, geographic area and province of residence, home language, the number of household members younger than 18 years, and a dummy variable indicating if a respondent lives with at least one child who is younger than seven years old.[16] When we regress using real monthly wages, we additionally control for weekly hours worked. Our estimate of interest is, of course, the coefficient on the binary indicator for women. The succinct results of these regressions are presented in *Figure 3* which shows the estimated unconditional and conditional estimates of our female binary indicator in both February and June 2020. The complete results are available in *Table A4* in the appendix.

Figure 3: Unconditional and conditional Mincerian regression coefficient plots of the gender wage gap in February and June 2020

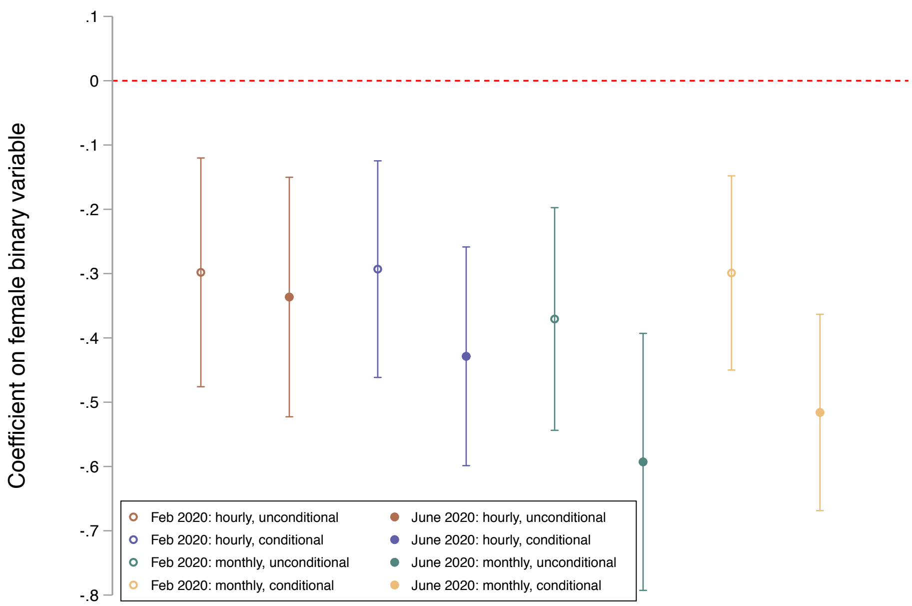

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes: [1] Within-wave samples restricted to employed individuals aged 18-64 years. [2] Estimates are weighted using computed bracket weights and account for complex survey designs. [3] Wages inflated to July 2020 Rands. [4] Conditional estimates obtained from regressing the natural logarithm of real hourly (monthly) wage on sex, age, age squared, race, highest level of education, main occupation, type of employment contract, marital status, geographic area and province of residence, home language, number of household members younger than 18 years, and a dummy variable for living with at least one child younger than 7 years (and weekly hours worked). Unconditional estimates obtained from regressing the natural logarithm of real hourly (monthly) wage on sex. [5] 95% confidence intervals presented as capped spikes.

Overall, our estimates suggest that even after controlling for several individual-level characteristics of men and women, both the unconditional and conditional monthly and hourly gender wage gaps were higher in June 2020 relative to February 2020, at least on average. Without controlling for any confounders, the average woman earned about 30% lower wages than a man for a given hour of work in February 2020 and 34% less in June 2020 – indicative of an increase in the unconditional

16 We are only able to control for variables included in the NIDS-CRAM. As such, we cannot control for other sources of variation, such as union membership for example.

gap of about 13%. This increase is even larger after we control for the aforementioned vector of covariates: the estimated conditional gender hourly wage gap at the mean is 29% in February 2020 and 43% in June 2020. That is, women earned 29% less than men of the same observable characteristics in February and 43% less in June – indicative of an increase in the conditional gender wage gap of about 46%. When we consider monthly rather than hourly wages, the increase in both the unconditional and conditional gender wage gaps are even larger. The unconditional gap increases from 37% in February to 59% in June 2020, while the conditional gap increases from 30% to 52% (representing a nearly 73% increase). Although these estimates of the conditional gender wage gap in both February and June 2020 are statistically significantly different from zero regardless of whether hourly or monthly wages are used, the changes in the gaps over the period are not statistically significant. This may be due to our relatively small samples. Despite this, the differences in the magnitudes of our point estimates are compelling. As such, we interpret these changes as indicative of a widening gender wage gap at the mean.

Next, we explore the evolution of the conditional gender wage gap across the entire earnings distribution, as opposed to just at the mean. One particular concern with the results presented in the remainder of this section is that in order for the wage regressions to be comparable between the two periods, we cannot control for industry. Research has shown strong results supporting the fact that an individual's industry of employment can have an impact on the size of the gender wage gap (Landman and O'Clery, 2020). In particular, the gender wage gap is found to be lower in male-dominated industries, likely due to male-dominated industries needing to diversify their workforce and recruit and retain female employees (Landman and O'Clery, 2020). In order to check the extent to which industry impacts our results, we ran the June 2020 regressions including industry dummies (as this data was available in the NIDS-CRAM Wave 2 data). The resulting estimates did not differ from those presented below by much, and as a result, we are confident that the results below do not suffer greatly from having industry left out of the model specification.

Figure 4 plots the conditional gender wage gap in monthly wages across the earnings distribution for February 2020 and June 2020, along with the relevant 90% confidence intervals for the estimates.[17] The left-hand panel of *Figure 4* shows the estimates of the gender wage gap from the actual distribution of June 2020 wages, while the right-hand panel plots the gender wage gap from the reweighted June 2020 sample to account for sample selection.

Figure 4: Estimates of the conditional gender wage gap in real monthly wages across the wage distribution, February 2020, June 2020, and June 2020 (reweighted)

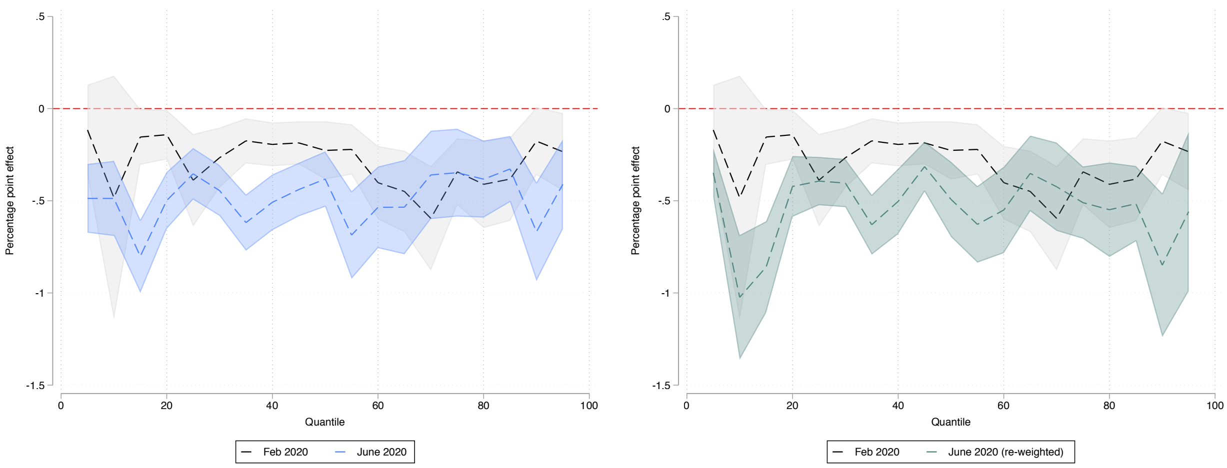

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes: [1] Left-hand panel shows February 2020 estimated gender wage gap compared to June 2020 estimated gender wage gap. Right-hand panel shows February 2020 estimated gender wage gap compared to reweighted June 2020 gender wage gap. [2] Variables controlled for in regressions include age, age squared, race, highest level of education, main occupation, area of residence, province, home language, marital status, presence of a written contract, number of cohabiting children under age 18, and weekly hours worked. [3] Estimates are weighted using computed bracket weights, or DFL reweighted bracket weights. [4] Wages inflated to July 2020 Rands. [5] Shaded areas represent 90% Confidence Intervals.

17 Figures presenting 95% confidence intervals are presented in the Appendix.

For the most part, the confidence intervals for the February 2020 and June 2020 estimates (whether reweighted or not) overlap, providing little evidence of a unilateral deepening of the monthly gender wage gap. However, in both panels of *Figure 4*, the 90% confidence intervals overlap less often at the bottom of the distribution than at the top. This indicates that there is a more severe and statistically significant deepening of the monthly gender wage gap at the bottom of the distribution than at the top. Comparing the two panels, it is also clear that the monthly gender wage gap has deepened more severely for those at the bottom of the distribution in the right-hand panel than in the left-hand panel. This indicates that by not correcting for sample selection bias in the June 2020 sample, one may underestimate the gender wage gap. This is sensible, as low-wage earners at the bottom of the distribution were more likely to lose their jobs during the lockdown (Ranchhod and Daniels, 2020), meaning that the bottom of the uncorrected earnings distribution in June 2020 would have been made up of individuals from higher up in the wage distribution.

Even though there is not consistent statistical significance of a deepening monthly gender wage gap, point estimates of the February-to-June gender wage gap ratio point towards a trend of deepening wage inequality. *Table 3* reports the point estimates of the gender wage gap as extracted from the RIF regressions at each quantile for February 2020, June 2020 and the reweighted June 2020 sample. Interestingly, in February 2020, there is little statistical evidence of a gender wage gap up to the 20th percentile of the wage distribution as the coefficient estimates are either insignificant or only significant at the 10% level. However, in both samples for June 2020, women clearly earn statistically significantly lower monthly wages than men at every quantile of the earnings distribution, all else constant. As a sense check, we confirm that our estimates at the median of the distribution accord with reported gender wage gap estimates from Mosomi (2018) and Bezuidenhout et al. (2019).These researchers estimate a gender wage gap of approximately 35% at this point in the distribution, and while the point estimates we present are slightly different, all 95% confidence intervals overlap with an estimated 35% gap, indicating that our estimates are more or less in line with the literature.

Table 3: Distribution of conditional gender wage gap estimates in real monthly wages, February 2020, June 2020, and June 2020 (reweighted)

	Gender wage gap estimates			Ratios	
Quantile	**February 2020**	**June 2020**	**June 2020 (reweighted)**	**Feb 2020:June 2020**	**Feb 2020: June 2020 (reweighted)**
5	-0.116	-0.487***	-0.349***	0.238	0.332
10	-0.483	-0.487***	-1.024***	0.992	0.472
15	-0.154*	-0.801***	-0.859***	0.192***	0.179***
20	-0.142*	-0.498***	-0.422***	0.284**	0.335
25	-0.388**	-0.354***	-0.393***	1.099	0.988
30	-0.266***	-0.444***	-0.404***	0.600	0.660
35	-0.175**	-0.618***	-0.630***	0.283***	0.278***
40	-0.194***	-0.508***	-0.503***	0.382*	0.386*
45	-0.186***	-0.439***	-0.315***	0.424	0.591
50	-0.227**	-0.382***	-0.492***	0.593	0.461
55	-0.221***	-0.686***	-0.628***	0.322**	0.352*
60	-0.402***	-0.536***	-0.550***	0.750	0.731
65	-0.449***	-0.535***	-0.352***	0.840	1.278
70	-0.595***	-0.360**	-0.424***	1.655	1.405
75	-0.343***	-0.347**	-0.510***	0.986	0.672
80	-0.411***	-0.383***	-0.548***	1.073	0.749
85	-0.382***	-0.328***	-0.516***	1.167	0.741
90	-0.177	-0.667***	-0.849***	0.265*	0.208*
95	-0.234*	-0.410***	-0.559**	0.570	0.419

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes:
1. Within-wave samples restricted to employed individuals aged 18-64 years.
2. Point estimates are the coefficient on the female dummy in the relevant RIF regression at a given quantile.
3. Estimates are weighted using computed bracket weights, or DFL reweighted bracket weight
4. All monthly wages inflated to July 2020 Rands.
5. * $p<0.1$; ** $p<0.05$; *** $p<0.01$
6. Variables controlled for in regressions include age, age squared, race, highest level of education, main occupation, area of residence, province, home language, marital status, presence of a written contract, number of cohabiting children under age 18, and weekly hours worked.

Where the evolution of the conditional gender wage gap across the distribution is concerned, the results are generally insignificant, indicating that the gap has not statistically significantly changed from February 2020 to June 2020. This being said, however, there is some statistical significance in the changes below the 40th percentile. Over this portion of the distribution, the February-to-June gender wage gap ratios show that the gender wage gap has deepened to up to five times its February 2020 level (and these are often significant changes). The point estimates of the ratios above the 60th percentile are much milder, and indicate that although there may have been a deepening of the gender wage gap for these individuals, it was substantially less severe than for those at the bottom of the distribution. Although broadly insignificant, these results are still interesting: if there is truly a deepening of the gender wage gap amongst low earners, then it means that already-vulnerable workers are being subjected to further inequality and welfare losses. Conversely, those individuals at the top of the wage distribution are experiencing much less severe effects of the lockdown on wage inequality, with some individuals around the 70th percentile even seeing a narrowing of the gender wage gap.

If this evidence is considered indicative of trends in wage inequality that may promulgate into the future, then there is cause for concern that individuals at the lower end of the wage distribution are being disproportionately disadvantaged relative to those at the top of the wage distribution. This pattern correlates with the incidence of ability to work from home, as estimated by Kerr and Thornton (2020). Evidence from Collins et al. (2020) cites a disproportionate incidence of childcare responsibilities during the pandemic falling to women. It is possible that part of the changing wage dynamic could be explained by women who have had to adjust their working hours in order to shoulder the more substantive childcare burdens placed on them due to school closures during the national lockdown. In order to investigate this, we turn our attention now to the trajectory of the hourly wage gap between February 2020 and June 2020.

Figure 5: Estimates of the conditional gender wage gap in real hourly wages across the wage distribution, February 2020, June 2020, and June 2020 (reweighted)

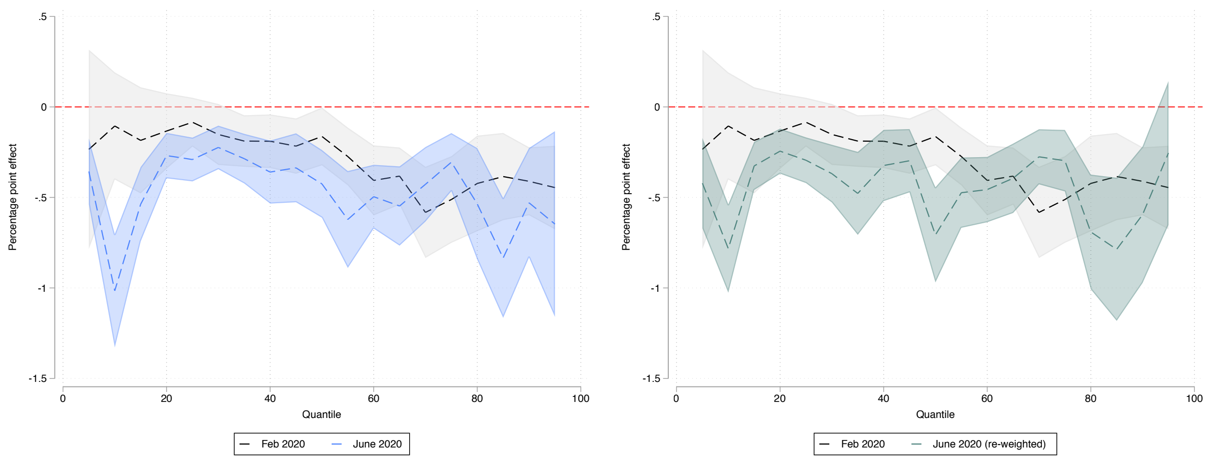

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes: [1] Left-hand panel shows February 2020 estimated gender wage gap compared to June 2020 estimated gender wage gap. Right-hand panel shows February 2020 estimated gender wage gap compared to reweighted June 2020 gender wage gap. [2] Variables controlled for in regressions include age, age squared, race, highest level of education, main occupation, area of residence, province, home language, marital status, presence of a written contract, number of cohabiting children under age 18, and weekly hours worked. [3] Estimates are weighted using computed bracket weights, or DFL reweighted bracket weights. [4] Wages inflated to July 2020 Rands. [5] Shaded areas represent 90% Confidence Intervals.

Figure 5 replicates the results from *Figure 4*, but with hourly wage as the dependent variable instead of monthly wage. The most striking result observed in *Figure 5* is that there are almost no points where the estimated hourly gender wage gap in February is different from the hourly gender wage gap estimated in June 2020, whether from the original sample or the DFL-reweighted sample. This result is replicated in *Table 4*, where the significance of the February-to-June hourly gender wage gap ratio is generally insignificant across the distribution. Even when focus is turned to the magnitude of the point estimates, one can see that the hourly gender wage gap has deepened much less severely than the monthly gender wage gap.

Table 4: Distribution of conditional gender wage gap estimates in real hourly wages, February 2020, June 2020, and June 2020 (reweighted)

	Gender wage gap estimates			Ratios	
Quantile	**February 2020**	**June 2020**	**June 2020 (reweighted)**	**Feb 2020:June 2020**	**Feb 2020: June 2020 (reweighted)**
5	-0.234	-0.357***	-0.421***	0.657	0.556
10	-0.105	-1.015***	-0.782***	0.104**	0.135**
15	-0.185	-0.538***	-0.326***	0.344	0.567
20	-0.133	-0.269***	-0.244***	0.496	0.546
25	-0.084	-0.291***	-0.295***	0.290	0.286
30	-0.152	-0.224***	-0.368***	0.680	0.413
35	-0.189**	-0.286***	-0.478***	0.660	0.395
40	-0.190**	-0.360***	-0.324***	0.527	0.586
45	-0.216**	-0.336***	-0.297***	0.643	0.728
50	-0.163*	-0.425***	-0.707***	0.384	0.231**
55	-0.274***	-0.622***	-0.474***	0.440	0.578
60	-0.406***	-0.496***	-0.457***	0.819	0.889
65	-0.382***	-0.547***	-0.394***	0.698	0.969
70	-0.583***	-0.427***	-0.276***	1.363	2.114
75	-0.512***	-0.305***	-0.297***	1.677	1.721
80	-0.423***	-0.536***	-0.691***	0.789	0.612
85	-0.385***	-0.834***	-0.787***	0.461	0.489
90	-0.411***	-0.529***	-0.595***	0.776	0.690
95	-0.446***	-0.647**	-0.255	0.689	1.751

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes:
1. Within-wave samples restricted to employed individuals aged 18-64 years.
2. Point estimates are the coefficient on the female dummy in the relevant RIF regression at a given quantile.
3. Estimates are weighted using computed bracket weights, or DFL reweighted bracket weights.
4. All monthly wages inflated to July 2020 Rands.
5. * p<0.1; ** p<0.05; *** p<0.01
6. Variables controlled for in regressions include age, age squared, race, highest level of education, main occupation, area of residence, province, home language, marital status, presence of a written contract, number of cohabiting children under age 18, and weekly hours worked.

Taken together with the results presented in *Figure 4* and *Table 3*, these results tell an interesting story: It seems that although there has been no significant change in the hourly wage gap between men and women between February 2020 and June 2020, there is evidence of a deepening monthly gender wage gap at the bottom of the wage distribution. This suggests that a change in hours worked may be the driving force behind the deepening gender wage gap at the monthly level. Given evidence from Collins et al. (2020), this may be as a result of an increased childcare burden that disproportionately affected women's ability to work effectively during this period of working from home. With the majority of schools closed between March 2020 and June 2020 due to the national lockdown, parents would have had to take responsibility for childcare, and this responsibility is found to disproportionately fall on mothers, thus negatively impacting the number of hours they are able to work at their jobs. Thus, even if hourly wage inequality between men and women remained unchanged (as it has done between February 2020 and June 2020, according to our estimates), then monthly wage inequality would have increased due to adjustments in working patterns along the intensive margin. Alon et al. (2020) suggest that women may also be less likely to be employed in jobs that are easily adapted to be tele-commutable. In other words, women may not be able to work from home as easily as men. If this is true, then women would likely decrease their hours worked, also potentially explaining the patterns in the deepening gender wage gap we see above.

Overall, the model results indicate that although the statistical significance is sparse, there is still some indication of a widening monthly gender wage gap between February 2020 and June 2020. This finding is relatively robust when reweighting to account for selection concerns. However, it is likely that even with insignificant results, the trajectories created by these point estimates are indicative of a general pattern that we should take note of. There seems to be a generally significant deepening of the conditional gender wage gap in the bottom third of the wage distribution, while the gap has not deepened as severely in the top of the distribution. Coupled with the fact that there is little evidence to support a deepening hourly gender wage gap, it is possible that the deepening monthly gender wage gap is attributable to women decreasing their hours worked more than men. According to the literature, one reason for this may be because of increased childcare responsibilities (Collins et al., 2020). In the following section of the paper, we bring together our findings from the rest of the paper and provide suggestions for how policy might intervene to avert increased inequality as a result of the national lockdown before it becomes a more serious issue.

7. Conclusion and policy considerations

Unlike previous recessions where it has been observed that men have borne the brunt of the economic downturn, the COVID-19 ‘pandemic recession’ is likely to disproportionately and persistently impact women. In the context of South Africa, initial research has already shown that of the estimated three million fewer employed people in April relative to February 2020 as a result of the pandemic-induced national lockdown, two in every three were women. However, less is known about the implications of the pandemic on those women who managed to remain in employment during the lockdown period. In this light, we use newly available, representative survey data to analyse the evolution of gender wage inequality in South Africa prior to and during the national lockdown. We do so by constructing estimates of the unconditional and conditional gender wage gaps through Mincerian-esque regressions and Recentered Influence Function (RIF) regressions for a pre-lockdown period and compare them to similar estimates from during the lockdown to determine whether there have been any inequality-deepening impacts of the lockdown on intergender wages. Additionally, we analyse variation in gender wage inequality across the entire wage distribution, given the evidence of distributional heterogeneity in the South African literature.

Concerns regarding sample selection were raised as a result of systematic differences between job-losers and job-retainers. It is known that more vulnerable individuals were more likely to lose their jobs as a result of the national lockdown (Ranchhod and Daniels, 2020), and as a result, the sample of wage earners in February 2020 and June 2020 may not be comparable. Using the DiNardo, Fortin and Lemieux (DFL) reweighting technique, we are able to create a hypothetical earnings distribution that would have arisen if the June 2020 sample had the same characteristics as the February 2020 sample. Running all RIF regressions on the original as well as the reweighted sample served to confirm that our results are robust to sample selection concerns.

Our main results begin by showing that the unconditional average gender wage gap in South Africa was large and evident both before and during lockdown, while the average gap widened from February to June 2020, regardless of whether monthly or hourly wages are used. Even after accounting for several confounding factors, our estimates suggest that the conditional gender wage gap was 46% to 73% higher in June 2020 relative to February on average. Although the conditional gender wage gap is statistically significant in each period, the changes in the gaps over the period are not statistically significant. Despite this, the differences in the magnitudes of our point estimates are compelling and as such, we interpret these changes as indicative of a widening gender wage gap at the mean.

When we further investigate this gap across the entire wage distribution, it is clear that the gap exists and varies considerably. Although estimates of the deepening of the monthly gender wage gap were generally insignificant for the majority of the distribution, there were some significant changes below the 40th percentile which were not reflected at the top of the distribution. At the bottom of the distribution, the monthly gender wage gap in June 2020 is estimated to be up to 5 times larger than its value in February 2020. This result is robust to sample selection corrections. An increase in the size of the conditional gender wage gap amongst the poorest third of earners is of particular concern for policymakers, given that this result speaks to deepening inequality amongst an already vulnerable group. When analysing how the hourly gender wage gap deepened, however, almost all estimates were insignificant, irrespective of their position along the wage distribution. The fact that hourly wage inequality did not change much, while monthly wage inequality did, suggests that there has been an adjustment of working hours amongst women that outweighs the change in working hours amongst men. Reasons for this could include women being less able to work from home (Alon et al., 2020), or that women have had to carry the brunt of increased childcare responsibilities during lockdown (Collins et al., 2020).

Policy which seeks to mitigate the adverse implications of such widening wage inequality ought to consider providing targeted income support to such workers at the bottom of the distribution. Several policy options are available. In their analysis of the NIDS-CRAM Wave 2 data, Köhler and Bhorat (2020) find that the distribution of the application for and personal and household-level receipt of the special COVID-19 Social Relief of Distress (SRD) grant, despite not being means-tested, has been relatively pro-poor. However, the authors note that because other grant recipients are not eligible to apply for the grant, and because nearly 85% of all grant recipients are women, most recipients (two-thirds) of the COVID-19 SRD grant are men. Therefore, despite it being progressively targeted, this grant in its current form may not be appropriate in this context. However, considering the Child Support Grant is also progressively targeted and that about 98% of recipients (caregivers) are women, the pandemic-induced expansion of the grant may be an optimal mechanism to provide such income support. Policymakers ought to consider further extending this expansion of social assistance beyond October 2020, as this would provide low-earning women with much-needed monetary support to combat deepening levels of inter-gender inequality as a result of the national lockdown.

A second option open to policymakers is that of state-subsidised childcare. With schools reopening for students to return to their studies, female caregivers will have more time to engage in labour market activities, thus decreasing the monthly gender wage gap once again. However, in order to provide women with an opportunity to engage in the labour market as fully as possible, the state could provide after-school care for students at public schools. If carefully rolled out, this type of intervention could assist in a number of spheres: not only would it free up time for women to adjust their working hours upwards once again, but if this after-school care were to provide food and academic support to students, then it may actively bolster academic outcomes for learners in the future.

REFERENCE LIST

Alon, T. M., Doepke, M., Olmstead-Rumsey, J., and Tertilt, M. (2020). The impact of COVID-19 on gender equality. Technical report, National Bureau of Economic Research.

Antonczyk, D., Fitzenberger, B., and Sommerfeld, K. (2010). Rising wage inequality, the decline of collective bargaining, and the gender wage gap. Labour economics, 17(5):835–847.

Ardington, C. (2020). NIDS-CRAM Wave 1 Data Quality. National Income Dynamics Study Coronavirus Rapid Mobile Survey (NIDS-CRAM) Technical Document, Wave 1.

Bezuidenhout, C., van Rensburg, C. J., Matthee, M., and Stolzenburg, V. (2019). Trading firms and the gender wage gap: evidence from South Africa. Agenda, 33(4):79–90.

Bhorat, H. and Goga, S. (2013). The gender wage gap in post-apartheid South Africa: A reexamination. Journal of African Economies, 22(5):827–848.

Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. Journal of Economic Literature, 55(3):789–865.

Blundell, R., Dias, M. C., Joyce, R., and Xu, X. (2020). COVID-19 and inequalities. Fiscal Studies.

Budlender, D. (2019). Changes in the care burden over the transition to adulthood. SALDRU Working Paper Number 234, Version 1/ NIDS Discussion Paper 2019/1.

Casale, D. and Posel, D. (2020). Gender and the early effects of the Covid-19 crisis in the paid and unpaid economies in South Africa. National Income Dynamics Study Coronavirus Rapid Mobile Survey (NIDS-CRAM) Working Paper, Wave 1(4).

Collins, C., Landivar, L. C., Ruppanner, L., and Scarborough, W. J. (2020). COVID-19 and the gender gap in work hours. Gender, Work & Organization.

DiNardo, J., Fortin, N. M., and Lemieux, T. (1996). Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. Econometrica, 64:1001–1044.

Elder, T. E., Goddeeris, J. H., and Haider, S. J. (2015). Isolating the roles of individual covariates in reweighting estimation. Journal of applied econometrics, 30(7):1169–1191.

Firpo, S., Fortin, N. M., and Lemieux, T. (2009). Unconditional quantile regressions. Econometrica, 77(3):953–973.

Gustafsson, M. (2020). How does South Africa's COVID-19 response compare globally? A preliminary analysis using the new oxcgrt dataset. Stellenbosch Economic Working Papers, WP07/2020. Stellenbosch: University of Stellenbosch.

Hegewisch, A., Liepmann, H., Hayes, J., and Hartmann, H. (2010). Separate and not equal? Gender segregation in the labor market and the gender wage gap. IWPR Briefing Paper, 377:1–16.

Hinks, T. (2002). Gender wage differentials and discrimination in the new South Africa. Applied Economics, 34(16):2043–2052.

Ingle, K., Brophy, T., and Daniels, R. (2020). National Income Dynamics Study – Coronavirus Rapid Mobile Survey (NIDS-CRAM) Panel User Manual. Wave 2 Release September 2020. Version 1.

Jain, R., Budlender, J., Zizzamia, R., and Ihsaan, B. (2020). The labour market and poverty impacts of Covid-19 in South Africa. National Income Dynamics Study Coronavirus Rapid Mobile Survey (NIDS-CRAM) Working Paper, Wave 1(5).

Kerr, A. and Thornton, A. (2020). Essential workers, working from home and job loss vulnerability in South Africa. A DataFirst Technical Paper 41. Cape Town: DataFirst, University of Cape Town.

Köhler, T. and Bhorat, H. (2020). Social assistance during South Africa's national lockdown: Examining the COVID-19 grant, changes to the Child Support Grant, and post-October policy options. National Income Dynamics Study Coronavirus Rapid Mobile Survey (NIDS-CRAM) Working Paper, Wave 2.

Landman, M. S. and O'Clery, N. (2020). The impact of the employment equity act on female inter-industry labour mobility and the gender wage gap in South Africa. World Institute for Development Economic Research (UNU-WIDER), WIDER Working Paper 2020/52.

Mosomi, J. (2019). An empirical analysis of trends in female labour force participation and the gender wage gap in South Africa. Agenda, 33(4):29–43.

Mosomi, J. N. (2018). Distributional changes in the gender wage gap in the post-apartheid South African labour market. PhD thesis, Faculty of Commerce. University of Cape Town.

Mwabu, G. and Schultz, T. P. (2000). Wage premiums for education and location of South African workers, by gender and race. Economic Development and Cultural Change, 48(2):307–334.

National Income Dynamics Study Coronavirus Rapid Mobile Survey (NIDS-CRAM) 2020 Wave 2 [dataset] (2020). Version 1.1.0. Cape Town: Allan Gray Orbis Foundation [funding agency]. Cape Town: Southern Africa Labour and Development Research Unit [implementer].

Ntuli, M. (2007). Exploring gender wage "discrimination" in South Africa, 1995-2004: A quantile regression approach.

Posel, D. and Casale, D. (2019). Gender and the economy in post-apartheid South Africa: Changes and challenges. Agenda, 33(4):3–10.

Ranchhod, V. and Daniels, R. (2020). Labour market dynamics in South Africa in the time of Covid-19: Evidence from Wave 1 of the NIDS-CRAM survey. National Income Dynamics Study Coronavirus Rapid Mobile Survey (NIDS-CRAM) Working Paper, Wave 1(9).

Republic of South Africa (1998). Employment Equity Act, No. 55 of 1998. Government Gazette, 400(19370).

Weichselbaumer, D. and Winter-Ebmer, R. (2005). A meta-analysis of the international gender wage gap. Journal of Economic Surveys, 19(3):479–511.

Wittenberg, M. (2017). Wages and Wage Inequality in South Africa 1994–2011: Part 1 – Wage Measurement and Trends. South Africa Journal of Economics, 85(2).

Yamaguchi, S. (2018). Changes in returns to task-specific skills and gender wage gap. Journal of Human Resources, 53(1):32–70.

Appendix

Table A1: Harmonisation of NIDS-CRAM Wave 2 income brackets into one consistent bracket

Original NIDS-CRAM bracket (Rands per month)	Formula for conversion	Final monthly equivalent bracket
Weekly bracket	**Lower or upper bound/7 × 30**	
0		0
1-700		<= 3 000
701-1 400		3 001-6 000
1 401-2 800		6 001-12 000
2 801-5 500		12 001-24 000
Fortnightly bracket	**Lower or upper bound/14 × 30**	
1 401-2 800		3 001-6 000
2 801-5 500		6 001-12 000
5 501-11 000		12 001-24 000
> 11 000		> 24 000

Source: NIDS-CRAM Wave 2. Authors' own calculations.
Note: Only relevant brackets with non-missing data shown here.

Table A2: Probability of reporting an actual Rand amount by final income bracket and wave

	Probability of reporting a Rand amount	
Final monthly earnings bracket	**NIDS-CRAM W1: Feb 2020**	**NIDS-CRAM W2: June 2020**
1. Zero/nothing	0.639	0.945
2. Less than R3000	0.853	0.879
3. Between R3001 and R6000	0.831	0.932
4. Between R6001 and R12000	0.865	0.871
5. Between R12001 and R24000	0.735	0.762
6. More than R24001	0.756	0.771

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.

Figure A1: Distribution of log nominal monthly wages in February and June 2020: unweighted and weighted using the sampling weights and computed bracket weights

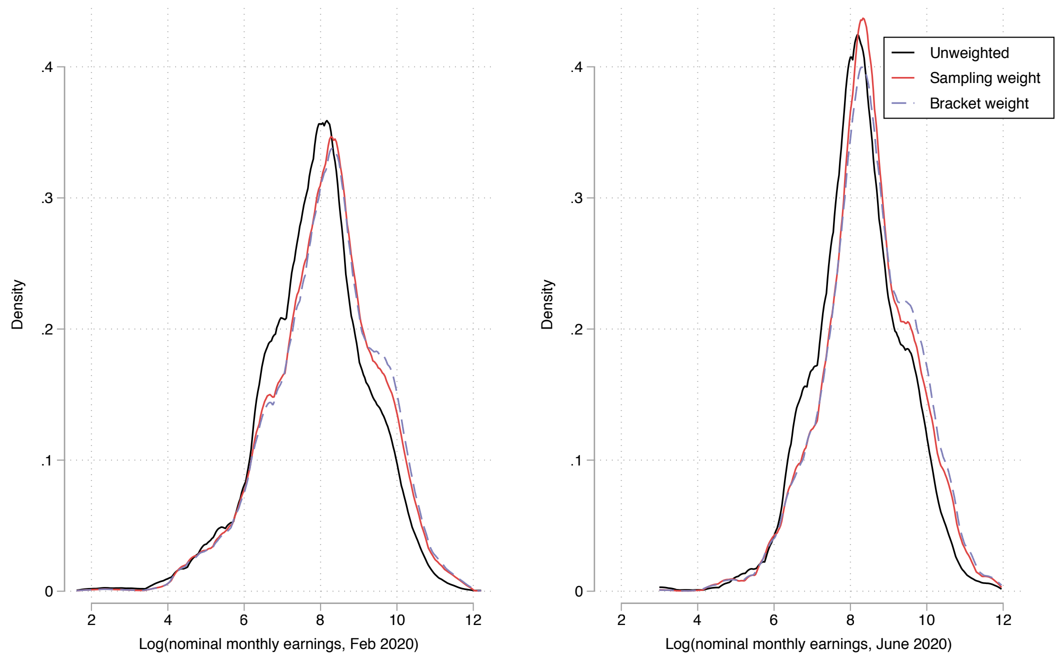

Authors' own calculations.
Source: NIDS-CRAM Waves 1 and 2.

Figure A2: February 2020, June 2020 and hypothetical June 2020 wage distributions after DFL reweighting

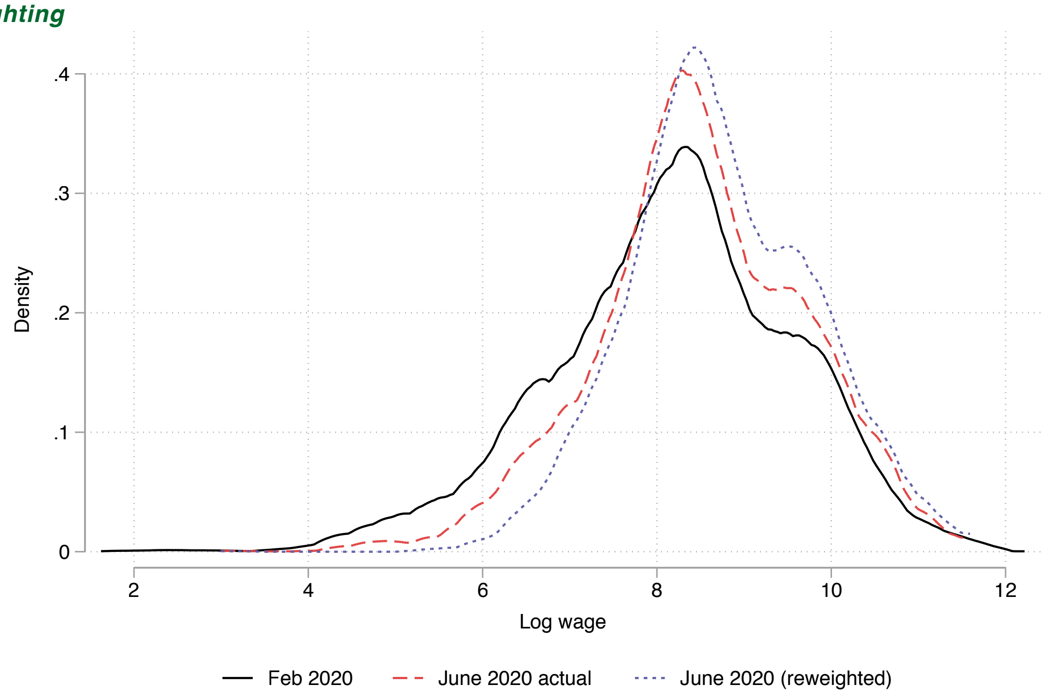

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes:
1. Variables controlled for in regressions include age, age squared, race, highest level of education, main occupation, area of residence, province, home language, marital status, presence of a written contract, number of cohabiting children under age 18, and weekly hours worked.
2. Hypothetical distribution represents the distribution of earnings that wage earners in June 2020 would have obtained if their characteristics were the same as February 2020 wage earners.
3. Wages inflated to July 2020 Rands.

Table A3: Mincerian OLS regression estimates for February and June 2020 to accompany Figure 4

Period	Feb	Feb	June	June	Feb	Feb	June	June
Dep var:	Log(real hourly wage)				Log(real monthly wage)			
Female	-0.298***	-0.293***	-0.336***	-0.429***	-0.371***	-0.299***	-0.593***	-0.516***
Age		0.138***		0.053**		0.147***		0.038*
Age squared		-0.001***		-0.000		-0.002***		-0.000
Coloured		-0.030		-0.542**		0.028		-0.288*
Asian/Indian		-0.135		-0.234		0.055		-0.232
White		0.479**		-0.177		0.516**		0.041
Urban		0.143		-0.101		0.157		0.033
Farm		-0.204		-0.206*		0.036		-0.099
Up to Secondary		0.154		0.406**		0.204		0.495***
Matric		0.385**		0.563***		0.436**		0.725***
Tertiary		0.856***		0.997***		0.919***		1.113***
Eastern Cape		0.141		-0.286		0.092		-0.228
Northern Cape		0.541**		-0.179		0.306**		-0.220
Free State		0.173		-0.457*		0.283*		-0.216
KwaZulu-Natal		0.115		-0.122		0.117		-0.118
North West		0.322*		-0.241		0.286		-0.148
Gauteng		0.294**		-0.028		0.331**		-0.043
Mpumalanga		0.283*		-0.178		0.279*		-0.133
Limpopo		0.332*		-0.081		0.259		-0.097
Managers		0.262		0.378		0.384*		0.274
Professionals		0.493*		0.164		0.553**		0.093
Technicians		0.322		-0.831***		0.365		-0.789***
Clerks		-0.229		-0.199		-0.100		-0.366**
Service workers		-0.390*		-0.752***		-0.231		-0.618***
Skilled agri		-0.381		-0.760***		-0.018		-0.711***
Craft		-0.352		-0.091		-0.328		-0.215
Plant operators		-0.186		-0.288		-0.014		-0.226
Elementary		-0.586**		-0.598***		-0.527**		-0.596***
Married		0.118		0.078		0.175**		0.207***
IsiXhosa		0.499		0.852***		0.659**		0.952***
IsiZulu		0.354		0.688***		0.541*		0.886***
Sepedi		0.110		0.606**		0.224		0.810**

Sesotho		0.373		0.943***		0.374		0.846***
Setswana		0.160		0.902***		0.391		1.008***
siSwati		0.384		0.641**		0.543		0.757**
Tshivenda		0.456		0.820**		0.795**		0.779**
Xitsonga		-0.013		0.595**		0.243		0.676**
Afrikaans		0.358		1.208***		0.447		0.988***
English		0.391		1.045***		0.462		0.944***
Written contract		0.306***		0.122		0.441***		0.232***
Number of residents < 4 18 years		-0.027		-0.051*		-0.004		-0.042*
Co-resides with child < 7 years		-0.019		0.028		-0.100		0.035
Working hours						0.089***		0.010***
Constant	3.368***	-0.658	3.822***	1.668***	8.344***	3.071***	8.749***	6.159***
Observations	2296	963	1168	983	2474	966	1577	983
R2	0.014	0.551	0.024	0.447	0.018	0.575	0.061	0.527

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes:
1. All estimates are weighted using computed bracket weights after accounting for complex survey designs.
2. * p<0.1, ** p<0.05, *** p<0.01.
3. Standard errors are suppressed.
4. Base groups for categorical variables: African, Traditional area, up to primary, Western Cape, Armed forces, and IsiNdebele.

Figure A3: Estimates of the conditional gender wage gap in real monthly wages across the wage distribution, February 2020, June 2020, and June 2020 (reweighted)

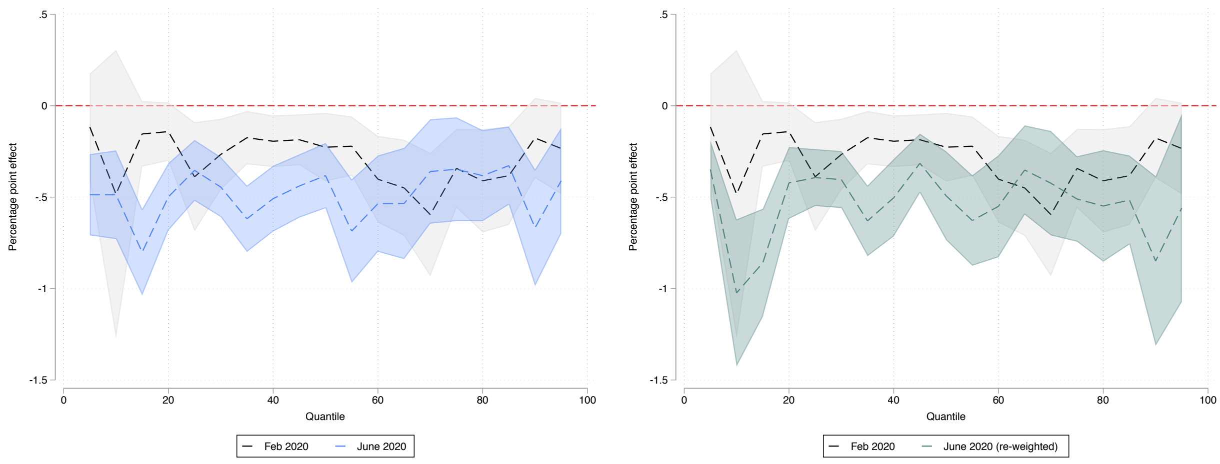

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes:
1. Left-hand panel shows February 2020 estimated gender wage gap compared to June 2020 estimated gender wage gap. Right-hand panel shows February 2020 estimated gender wage gap compared to reweighted June 2020 gender wage gap.
2. Variables controlled for in regressions include age, age squared, race, highest level of education, main occupation, area of residence, province, home language, marital status, presence of a written contract, and number of cohabiting children under age 18.
3. Estimates are weighted using computed bracket weights, or DFL reweighted bracket weights.
4. Wages inflated to July 2020 Rands.
5. Shaded areas represent 95% Confidence Intervals.

Figure A4: Estimates of the conditional gender wage gap in real hourly wages across the wage distribution, February 2020, June 2020, and June 2020 (reweighted)

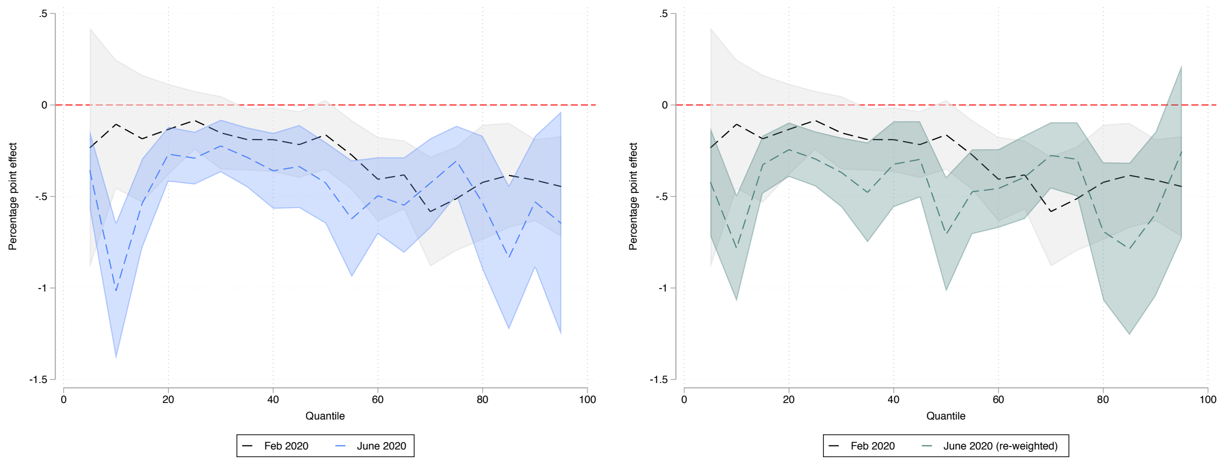

Source: NIDS-CRAM Waves 1 and 2. Authors' own calculations.
Notes:
1. Left-hand panel shows February 2020 estimated gender wage gap compared to June 2020 estimated gender wage gap. Right-hand panel shows February 2020 estimated gender wage gap compared to reweighted June 2020 gender wage gap.
2. Variables controlled for in regressions include age, age squared, race, highest level of education, main occupation, area of residence, province, home language, marital status, presence of a written contract, and number of cohabiting children under age 18.
3. Estimates are weighted using computed bracket weights, or DFL reweighted bracket weights.
4. Wages inflated to July 2020 Rands.
5. Shaded areas represent 95% Confidence Intervals.

For further information please see **cramsurvey.org**