Measuring employment volatility in South Africa using NIDS: 2008 – 2017

Rocco Zizzamia[1] and Vimal Ranchhod[2]

1. Introduction, motivation, literature

South Africa is a country with exceptionally high rates of unemployment, much of which is experienced as a chronic state. At the same time, a subset of the population experiences repeated transitions into and out of work, thus experiencing unemployment more as a transient than as a chronic state. Of the available data sources at present in South Africa, only a few have the longitudinal structure required to investigate the rates at which people find and lose employment.

There are many reasons that we would be interested in knowing what fraction of the labour force experiences chronic unemployment as compared to high levels of volatility between employment states. The first and most direct implication of employment transitions is that they would be correlated with both individual and household level welfare. For example, the strongest predictor of household level poverty is whether someone in the household has a formal sector job or not (Leibbrandt, Bhorat and Woolard, 2001). Schotte *et al* (2018) show that wage income dominates as a source of income for those households on the edge of falling into poverty and those just below the poverty line. The net effect of a high dependence on wages to sustain consumption, combined with the high rates of worker flows in South Africa, means that labour market transitions are responsible for a large share of poverty transitions. Just as poverty transitions appear to be determined largely by employment transitions, chronic poverty appears to be determined by structural exclusion from the labour market. Thus, the material well-being of most people in the economy is determined by overall labour market conditions as well as how successfully a particular individual can navigate this environment.

In addition to the negative effect of actually losing one's job, the *a priori* probability of losing one's employment will reduce one's welfare. This is implicitly captured by the concavity of most utility functions, where the degree of concavity represents the degree of risk aversion. A simple illustrative example of this effect is that most people would prefer to earn R100 with certainty than earn either R200 or R0 with a

[1] Research Officer, SALDRU, University of Cape Town.
[2] Professor, SALDRU, University of Cape Town.

50% probability of each outcome occurring. Psychologically, this uncertainty would be a cause of stress for individuals, and chronically high levels of stress are associated with a number of poor health outcomes (Witte, 1999; Cafiero & Vakis, 2006). Thus, it is not only current income that matters for welfare, "but also the risks a household faces, as well as its (in)ability to prevent, mitigate and cope with these" (Klasen and Waibel, 2012: 17).

A different mechanism by which volatility affects welfare is through the inability to plan appropriately within a dynamic context. Consider two people with the same expected net present value of their future income, but where one person's income is stable while the other person's income is volatile. The person with a volatile income stream will always need to keep a greater share of their income as precautionary savings, in anticipation of a proverbial 'rainy day'. This will work in effectively the same way as a tax on that individual, thus lowering their consumption in any given period. Alternatively, the greater volatility might result in a greater share of their income being spent on insurance, which again would reduce their utility. Other dynamic optimisation behaviours are also likely to be affected. Key amongst these are the ability to invest or buy durable assets such as property. For most people, buying a property involves securing credit from a mortgage provider, and either this becomes impossible, or the interest rates charged becomes higher, for people with unstable income streams.

Yet another pathway through which employment volatility affects welfare is via the impact of unemployment on subsequent employment and income. Here, there are at least two pathways through which one's subsequent employability are affected by contemporary unemployment. First, our human capital is likely to depreciate as we spend time out of the labour market. For example, one's level of expertise becomes relatively less up-to-date if one leaves the labour market for a prolonged period of time, relative to someone who had experienced sustained employment over the same time period. In addition to skills depreciation, the experience of unemployment can induce stress, depression, and other mental health issues that further reduce our levels of human capital (McKee-Ryan, 2005; Kingdon and Knight, 2004). This process will probably affect both the likelihood of obtaining subsequent employment, as well as adversely impacting on the wage offers of subsequent employment.

In addition to the effects of volatility on human capital, a second pathway that links future employment prospects with historical employment volatility involves the way that firms evaluate job applicants in a world with imperfect information. When firms cannot observe productivity directly, they use observable signals to differentiate between job applicants. A person with a chequered employment history is not necessarily an undesirable worker, but the probability that they will be problematic is higher than a similar worker with a stable employment track record. As such, firms will prefer to employ people who have already been consistently employed, thus generating some degree of unemployment persistence for individuals who have already experienced spells of unemployment in the past. Thus, the experience of volatile employment has potentially long term implications for both the productivity as well as the subsequent employment opportunities of an individual. Given all of the above, it may not be surprising that a recent study on public attitudes toward work in South Africa finds that job stability is ranked highest among several characteristics of what South Africans value most in a job (Mncwango, 2016).

The existence of a large fraction of people who experience chronic employment volatility also has societal implications. A large and stable middle class is often argued to be key for sustaining a democratic order; and for the existence of a stable middle class, the stability and quality of employment are essential (Schotte et al., 2018). Thus, sufficiently high levels of employment volatility can affect the overall socio-political environment in ways that have far reaching consequences for a country's development path.

While the study of employment volatility can easily be motivated by focussing on its welfare implications, understanding employment volatility also has potential policy implications. For one thing, both the value and sustainability of social security programs, such as unemployment insurance and a universal basic income grant, will be affected by the proportion of people who find and lose jobs at a high rate. In a context of high unemployment and employment volatility, it is also clear that a system of social insurance which is tied to work-based contributions will be inappropriate if the most vulnerable populations are to be included. Moreover, policy designed to alleviate unemployment may benefit from better targeting if one understands how dynamic patterns of unemployment vary in different environments and for different sub-populations. For example, long term chronic unemployment may require a structural intervention, while employment volatility may be better addressed by reducing frictions in the labour market.

Despite the importance of this topic, there is only a relatively small literature that focusses specifically on employment volatility in the country, and this is most likely due to data limitations. To investigate employment volatility, one needs longitudinal data on individuals, and the only nationally representative individual level panel study is the National Income Dynamics Study (NIDS). Nonetheless, a few earlier studies have measured the degree of churning in the South African labour market. To date, researchers have used various data sources to investigate employment volatility in South Africa. For example, Cichello et al. (2005) use the KIDS dataset to analyse poverty transitions amongst Indian and

African households in the Kwa-Zulu Natal province. Other researchers have used dwelling level matched panel data (Banerjee et al., 2008; and Ranchhod and Dinkelman, 2008), a Cape Town specific youth panel study (Pugatch, 2018), and administrative tax data (Kerr, 2017), in order to investigate various facets of employment dynamics. The general finding is that South Africa has high levels of labour market churning by international standards, although each of the studies is limited due to the way that the data are collected.

Since NIDS is nationally representative by design, and also tracks individuals who migrate to different households, this provides an ideal dataset to measure employment transitions over time. Cichello et al. (2014) use the first two waves of NIDS data to look at labour market churning. Ranchhod (2013) used the first three waves of NIDS to measure the levels of earnings volatility in the labour market. Essers (2016) looks at the first two waves of NIDS as well as matched Quarterly Labour Force Survey data to investigate labour market transitions between 2008 and 2011/12. In accordance with the previously discussed literature, the NIDS-based studies also find substantial amounts of employment volatility.

In this paper, we measure the degree of labour market transitions across the first five waves of NIDS, which spans the period from 2008 to 2017. Our contribution, relative to the existing body of knowledge on employment volatility in South Africa, is thus twofold. First, we are including more recent data and thus have more up-to-date estimates. As Essers (2016) notes, the first two waves of NIDS covered exactly the period following the global financial crisis of 2008. Thus, employment volatility at the time may have been exaggerated relative to long run trends. Second, and perhaps more importantly, the earlier NIDS studies had fewer waves of data with which to estimate volatility. This makes it impossible to separately identify people who experience a once-off labour market transition from people who experience repeated labour market volatility. With five waves of data we are able to measure the size of the latter group, and compare the size of this group relative to those who are in stable employment or in long term unemployment.

The remainder of this paper is structured as follows. In Section 2, we describe the methods that we use in this paper. Section 3 contains a description of the data and sample characteristics. Section 4 contains descriptive results that summarize employment transition patterns. In Section 5, we report regression results on estimating employment volatility by separately investigating the likelihood of finding or losing employment. Section 6 concludes.

2. Methods

This paper is concerned with measuring employment volatility for the South African population and differentiating these volatility patterns for various subsets of society. These groups are distinguished by household and individual level characteristics, which in many cases reflect structural inequalities that are associated with different intertemporal employment patterns. Our investigation is structured around variants of two empirical methodologies, both of which exploit the panel dimension of NIDS.

First, in Section 4, we use the full longitudinal scope of NIDS by constructing a series of employment transition "trees". These are essentially multi-period transition matrices, represented in stylised form in Figure 1. In these transition trees, the employment patterns for each individual is treated separately based on whether the individual in question was employed or not employed in the first period. Since this paper is motivated by an attempt to understand how different subgroups of society experience different probabilities of gaining and losing work, we construct several of these transition trees for different population sub-groups. For example, suppose that we are interested in understanding volatility differences by gender. In this case, we limit our sample to the population of working age males. Within this group, some individuals in Wave 1 are employed, while others are not. For the subset of this group that are employed, some fraction will still be employed in Waves 2, 3, 4 and 5, some will lose their employment in subsequent waves, while others will move into and out of employment between waves. All transition patterns are exhaustively represented in Figure 1. An identical categorization can be used to classify the subsequent employment status of the subset of working age males who were *not* employed in Wave 1. This exercise is repeated for the sub-sample of working age women. On this basis, a comparison of volatility patterns by gender can be undertaken.

The construction of these transition trees forms the basis of our descriptive analysis of employment volatility in Section 4. These trees can be conveniently summarised by distinguishing between individuals who remained employed in all periods, those who remained employed in four, three or two periods, and those who were never subsequently employed. This can give us a sense of how the likelihood of finding and losing employment differs for different subsets of the population, how important initial conditions (being employed or not employed in 2008) are in predicting employment in the longer run for different population groups, and how chronic versus transient unemployment is distributed across the population.

Figure 1: Stylised employment transition tree

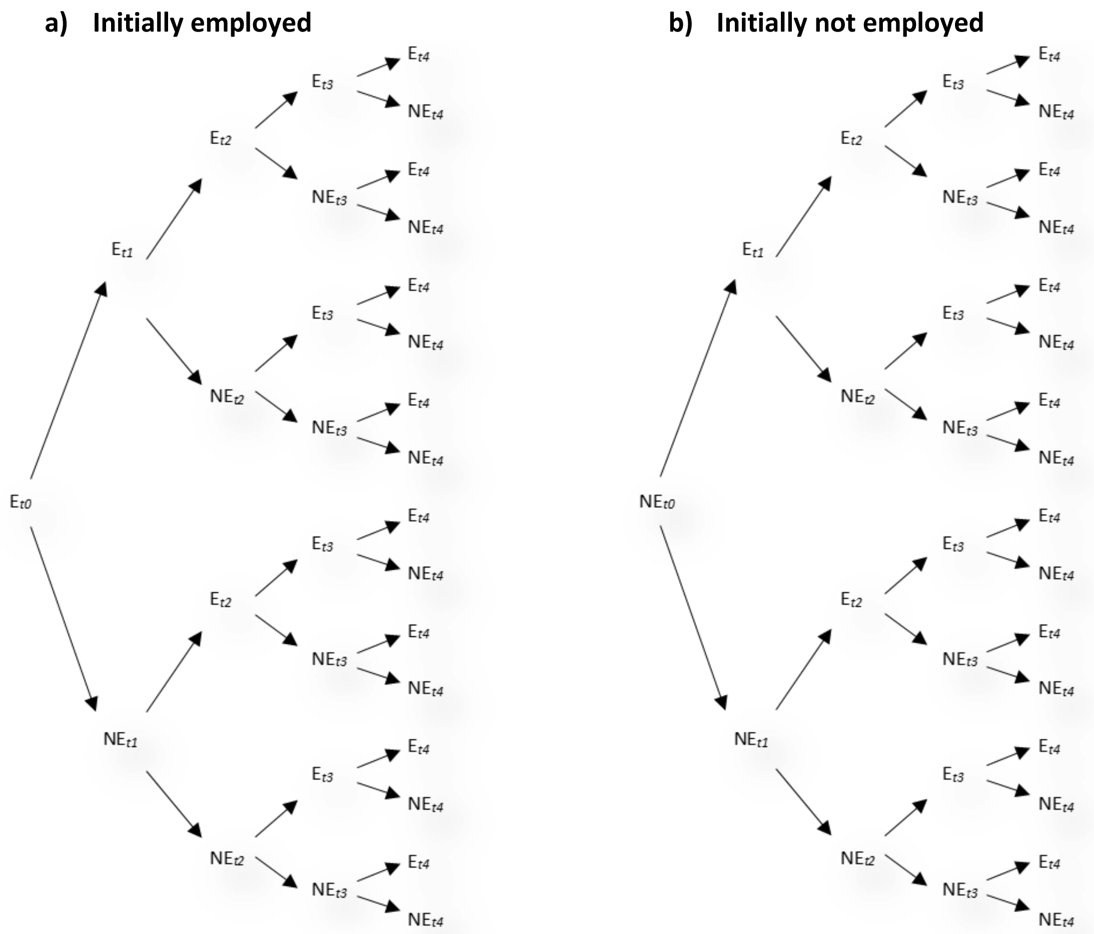

Notes: Author's representation. "E" signifies "employed" in period t, "NE" signifies "not employed" in period t.

The second core component of our analysis is to model the probability of losing (gaining) a job in time $t+1$, given household and individual level characteristics in time t. In doing so, we assume that individuals can be characterised as having a latent propensity to lose (gain) employment if they were initially employed (not employed). This relationship is modelled as:

$$empl^*_{it+1} = \boldsymbol{\beta}'\boldsymbol{X}_{it} + u_{it} \text{ and } empl_{it+1} = I(empl^*_{it+1} \leq 0.5) \qquad (1)$$

where $i = 1, \dots, n$ indexes individuals, $empl^*_{it+1}$ is a binary employment status outcome, $\boldsymbol{X}_{it}$is a vector of explanatory variables for each i, $\boldsymbol{\beta}$ is a vector of parameters to be estimated, and u_{it} is an error term which is assumed to follow the standard normal distribution ($u_{it}{\sim}N(0,1)$). $I(empl^*_{it+1} > 0.5)$ is a binary indicator function which takes on the value of 1 if the latent propensity is greater than or equal to 0.5, and zero otherwise.

We fit a probit model to NIDS data to regress employment status in $t+1$ on individual and household level characteristics in t, as described in (1) above. The outcome variable is the propensity to lose (gain) employment if initially employed (not employed).

Explanatory variables in the base specification include household-level variables and individual-level characteristics. The former includes household size, a geographic area variable (distinguishing farm areas, traditional areas and urban areas), main source of household income, number of employed household members, total number of grants in the household, province, and year fixed effects. Individual-level characteristics include a gender dummy, a race variable, educational attainment, a dummy for household head status, age and age squared.

3. Data

This paper uses NIDS panel data (SALDRU 2018a,b,c,d,e). NIDS is South Africa's only nationally representative household panel study, which began in 2008 with a sample of over 28,000 individuals in 7,300 households. There are currently five waves of data available spanning the nine years from 2008 to 2017, where each wave of data is spaced approximately two years apart.

In Section 4 we use the balanced panel of respondents to exploit the full longitudinal scope of the data, restricting our sample to observations which appear in all five waves. In Section 5, however, we pool data from pairs of consecutive waves ($t-1$ and t), such that the analysis of changes over time represent changes between 2008 to 2010/11, 2010/11 to 2012, 2012 to 2014/15 and 2014/15 to 2017 respectively, controlling for period-specific fixed effects. In the analysis in Section 5, observations are included as long as they appear in two subsequent waves.

To facilitate comparisons across time, all monetary figures are deflated using the Stats SA consumer price indices and are calibrated to March 2017.[3]

As with any longitudinal study, one is generally concerned with how much attrition there is over time. In Table 1~~Table 1~~ below, we show how panel attrition has affected our sample of prime aged adults when restricting our sample to the full balanced panel. Starting with Wave 1 in 2008, and including

Formatte

[3] To adjust for inflation, for each line the food component (equal to the FPL) is inflated by using the food specific Stats SA CPI and the non-food component (equal to the difference between the FPL and the UBPL) is inflated by using the non-food specific Stats SA CPI.

respondents from just the Adult sample, we have a maximum possible sample of 15,597. This number decreases to 6,701 if we exclude adults younger than 25 years or older than 50 years in 2008. Of these, we are able to track only 3,172 in all five waves, giving a net attrition rate of 47.34 percent for this sample.

Table 1: Sample composition

Restriction applied	N	% retention in balanced panel
Wave 1 adults, successfully interviewed and prime aged	6,701	
Adults, successfully interviewed in all waves and prime aged in Wave 1	3,172	47.34%

Notes:

a) 15,597 adults were interviewed in 2008, with 8,896 not being prime aged.

b) "Prime age" identifies those aged between 25 and 50 years in 2008, which means this restriction will apply to those aged 34 to 59 years in 2017.

We use panel weights to correct for the presence of this substantial attrition in NIDS. Fortunately, NIDS has released panel weights along with the data which adjust the original baseline survey design weight to ensure that the weighted distribution of households in terms of several relevant characteristics is the same for those who survive to Wave 5 as it is in the cross-sectional distribution in the survey's baseline wave.

Variables:

We make use of several variables in our analysis which are important to define. These fall into two categories: Outcome variables, and variables defining demographic groups of interest.

The key outcome variable is employment status, and more specifically, change in employment status. We define employment status as a dummy variable, indicating whether or not a person is employed in a given wave. It ought to be noted that several different forms of employment are captured by this variable, including regular employment, casual employment or self-employment. Note that we make the deliberate choice not to expand the number of labour market states to differentiate between the unemployed and the "not economically active". This decision is informed by both pragmatic considerations: Defining employment as a binary variable without distinguishing between the unemployed and the "not economically active" limits the number of branches in our transition tree substantially, and keeps the econometric modelling relatively simple, while still enabling us to investigate the relevant transitions between employment and states of economic inactivity.

In Section 4, we distinguish mobility patterns by defining several demographic groups. It ought to be noted that, apart from results presented for age categories, all other results are limited to respondents aged between 25 and 50 years in 2008. These individuals would thus be aged between 34 and 59 in 2017. The variables that define our demographic groups, and how they were coded are:

- **Race**: Since there are too few White and Indian respondents for meaningful racial comparisons across all four racial groups, we create a single "non-African" racial category which includes White and Coloured respondents. This is done to create a comparison group off which to contrast the labour market experiences of the African racial group, rather than because this somewhat artificial and heterogenous "non-African" group may reveal something meaningful on its own. Indians are dropped because of their very small sample size.
- **Age**: Age variables are defined in Wave 1 (2008) with "Youth" identifying those aged 16 to 24 in 2008, "Prime" identifying those aged 25 to 50 in 2008, and "Older" identifying those aged 51 to 64 in 2008. Thus, these categories are dynamic, with "Youth" identifying those aged 24 to 33 in 2017, "Prime" identifying those aged 34 to 59 in 2017, and "Older" identifying those aged 60 to 73 in 2017.
- **Gender**: We compare the volatility between female and male respondents.
- **Education**: We compare the volatility observed by different levels of educational attainment as measured in Wave 1. To implement this component of the analysis, we needed to classify different levels of education into a small number of categories. We decided to create three groups, one for people who have not completed secondary school, one for people who only have a matric qualification, and one for people who have any form of a post-matric qualification. This categorization was used as it seems to conform with different levels of signals of human capital in the labour market. In all likelihood, employers would also differentiate between a four year university degree and a 6 month diploma, but our group sizes become too small for such a comparison.
- **Geographic location**: We separated the sample into respondents residing in urban or rural areas in Wave 1, where "Rural" refers to communally-owned land under the jurisdiction of traditional leaders, defined as "traditional" land in the 2011 Census. Labour markets are likely to operate quite differently in urban and rural areas, and the job prospects and wage distributions are also likely to differ substantially between them. This led us to consider a comparison along these lines. The third geographical category defined in NIDS – commercial farming areas – is not considered in our analysis. This is because of

the relatively small share of respondents in these areas and because of the fact that the very particular labour market structures in farming areas limits comparability.

- **Poverty status**: Poverty is defined using the StatsSA upper-bound poverty line (UBPL) (StatsSA, 2015). The upper bound poverty line is calculated to indicate the expenditure level at which individuals can satisfy both their food and non-food needs. Expressed in March 2017 Rands, it is R1,503 per capita per month. Each respondent is classified as poor or non-poor based on whether their per capita household expenditure falls below or above this poverty line.

Sample composition and descriptive statistics

In Table 2~~Table 2~~ below we show how the composition of the balanced sample compares with each cross-section, in terms of the demographic groups that we consider. We note several interesting observations from Table 2~~Table 2~~. First, the sample size grows quite substantially with time in the cross-sections, due to new members joining the households of continuing sample members (CSMs). In relative terms, the sample size in the balanced panel is much smaller. In addition to not containing new household members, the balanced panel is also smaller due to attrition from the Wave 1 cross-section. Nonetheless, the total sample size remains large enough to provide sufficient statistical power for the analyses that we undertake in this paper.

Formatte

Formatte

Of the variables that we focus on, race, education and urban/rural splits remain fairly consistently distributed in the panel and across all five waves. There does appear to be a fairly generalised increase in educational attainment between 2008 and 2017, although since cell-proportions are unweighted this should be interpreted with caution. In addition, the sample becomes more African with time, and the balanced panel over-represents Africans, especially with respect to the wave 1 cross-section. This reflects that attrition is disproportionately a non-African phenomenon.

Age categories in the balanced panel are substantially different to each of the cross-sections. This is likely caused by an aging panel and South Africa’s youth-heavy demographic pyramid. Particularly interesting is the fact that women are disproportionately likely to remain in the balanced panel relative to men. To correct for the differences between the panel and the respective cross-sections, we apply the balanced panel weights when estimating our main results.

Table 2: Descriptive statistics – composition of sample for each wave vs for balanced panel

	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Average balanced panel, all waves

N	*15631*	*17626*	*18689*	*22740*	*23891*	*6859*
African	0.79	0.81	0.81	0.82	0.78	0.84
Less than matric	0.84	0.84	0.83	0.82	0.79	0.77
Matric	0.10	0.10	0.10	0.10	0.12	0.13
Tertiary	0.05	0.06	0.06	0.08	0.10	0.10
Youth (16-24)	0.25	0.25	0.26	0.26	0.26	0.34
Prime (25-50)	0.30	0.30	0.29	0.29	0.29	0.48
Older (51-64)	0.10	0.09	0.09	0.09	0.08	0.17
Female	0.54	0.54	0.54	0.54	0.54	0.68
Rural	0.44	0.44	0.44	0.43	0.40	0.45
Urban	0.48	0.48	0.49	0.51	0.54	0.48

Notes: All cell proportions are unweighted.

Descriptive statistics:

In **Error! Reference source not found.**~~Table 3~~ below, we present a comparison of the mean proportion employed and mean earnings in Wave 1 and Wave 5 of the cross section as well as the corresponding subset of the balanced panel in each wave.[4] All the means and proportions are weighted. In the cross-sections, we use the conventional design weights which are adjusted for non-response. In the balanced panel, we use panel weights provided with the NIDS data.

We obtain several insights from **Error! Reference source not found.**~~Table 3~~. First, whether considering the cross-sections or the panel sub-sample from each wave, there appears to have been a generalised improvement in labour market outcomes across all groups. The cross-sectional annualised real wage growth is positive (or negligibly negative) throughout. The percentage employed also seems to have gone up across the board, although improvements appear larger for the panel than the cross-section. This suggests that employed members of the panel are less affected by attrition, on average, than those with poorer labour market outcomes.

Second, while the Wave 1 panel summary statistics of percentage employed yield a more pessimistic picture than the cross-sections, this is reversed in Wave 5, where the panel yields more optimistic results than the cross-section. However, mean wages are lower in the balanced panel for both Wave 1 and Wave 5.

There are two ways in which these discrepancies could arise, and these are not mutually exclusive. First, there is the effect of attrition in the panel. If attrition is correlated with having better labour market outcomes, then respondents who survive into the balanced panel will have worse labour market

[4] Results for all five waves are available from the authors upon request.

experiences than the cross-section. Second, there is the composition effect of selective migration into households and the formation of new NIDS households. People who enter into NIDS households after Wave 1 are not continuing sample members, but are nonetheless included in the cross-sectional summary statistics of subsequent waves.

We can get a sense of the respective importance of these two processes. Consider first the difference between the Wave 1 panel and Wave 1 cross-section, which measures only the attrition effect, since household composition changes cannot affect these estimates. Here we find that outcomes in the balanced panel subsample are poorer than in the cross section, suggesting that those individuals with better labour market prospects are disproportionately likely to attrite from the sample. Next, consider subsequent divergences between the panel and cross-section in Wave 5, which reflects the combined effects of attrition as well as selective migration and household formation. Thus, by comparing the differences in Wave 5 relative to those in Wave 1, we can obtain some sense of the effect of selection in household composition on the comparability of the balanced panel relative to the cross-section.

Two observations from the comparison between the cross-section and panel for Wave 5 are striking: On the one hand, mean wages are substantially lower in the panel than in the cross-section. The fact that mean wages are lower suggests that workers with very high wages are more likely to attrite from the panel than workers with low wages, since the mean is especially sensitive to the effect of losing or including workers on the extreme high end of the earnings distribution. This is consistent with the comparison of the panel and cross-section in Wave 1.

On the other hand, the proportion employed is higher in the panel than the cross-section. Unlike mean wages, the proportion employed is not especially sensitive to the attrition of high wage earners. The greater proportion employed of the Wave 5 panel compared to the Wave 1 panel (where only attrition is driving the panel/cross-section discrepancies) suggests that household compositional changes are driving this effect – specifically, the in-migration of unemployed temporary sample members (TSMs) into CSM households. Another way of seeing this is to note that between Waves 1 and 5 the proportion employed in the cross-section and in the panel increase by 3.9 and 10.0 percentage points respectively. Thus, employment prospects improved relatively more for members of the balanced panel than for the people that they lived with. The improvement observed in the cross-section potentially reflects a real improvement in overall employment prospects, while at the same time changes in household composition appear to have offset this effect.

Table 3: Representativeness of panel - comparing % employment and mean earnings of panel vs cross-section

		Wave 1		**Wave 5**		**Real wage growth**
		Cross-section	*Balanced panel*	*Cross-section*	*Balanced panel*	*Cross-section*
Total	*% employed*	42.97%	41.07%	46.85%	51.10%	0
	Average wage	6414.807	5121.678	7531.313	5776.045	1.80%
African	*% employed*	40.03%	39.19%	45.34%	51.34%	
	Average wage	4348.151	4097.691	6016.716	5057.794	3.67%
Less than matric	*% employed*	36.51%	36.91%	37.75%	43.27%	
	Average wage	2987.62	2591.409	3809.061	3406.195	2.74%
Matric	*% employed*	49.48%	42.85%	51.63%	58.28%	
	Average wage	6474.078	5709.507	6922.326	5344.734	0.75%
Tertiary	*% employed*	71.07%	69.63%	75.83%	74.56%	
	Average wage	15353.55	13376.55	14880.66	11195.65	-0.35%
Youth (16-24)	*% employed*	21.74%	17.08%	61.64%	59.88%	
	Average wage	3203.056	2607.628	5697.736	4650.937	6.61%
Prime (25-50)	*% employed*	60.15%	55.36%	62.90%	59.86%	
	Average wage	6636.999	5020.439	6667.454	6338.84	0.05%
Older (50-64)	*% employed*	45.91%	47.91%	18.09%	18.43%	
	Average wage	8098.547	7055.517	6244.193	7457.287	-2.85%
Female	*% employed*	34.40%	34.97%	38.45%	42.98%	
	Average wage	4983.219	4120.05	6221.039	5080.835	2.50%
Male	*% employed*	53.92%	52.01%	55.96%	64.72%	
	Average wage	7498.217	6241.703	8498.258	6543.199	1.40%
Rural	*% employed*	28.80%	30.08%	31.63%	37.43%	
	Average wage	3369.246	2884.703	4544.245	4152.853	3.38%
Urban	*% employed*	48.97%	47.29%	52.86%	58.08%	
	Average wage	7576.119	6239.626	8513.149	6385.917	1.30%

Notes:
a) Cross sectional cell proportions weighted using post stratified weights, balanced panel cell proportions weighted using Wave 5 panel weights.
b) Age variables defined in Wave 1 (2008) with "Youth" identifying those aged 16 to 24 in 2008, "Prime" identifying those aged 25 to 50 in 2008, and "Older" identifying those aged 51 to 64 in 2008. Thus, these categories are dynamic, with "Youth" identifying those aged 24 to 33 in 2017, "Prime" identifying those aged 34 to 59 in 2017, and "Older" identifying those aged 60 to 73 in 2017.
c) Unemployed respondents are assigned a zero wage. Wage figures include wages from formal employment, casual work and self-employment. Monetary figures are expressed in March 2017 Rand values.
d) "Rural" refers to communally-owned land under the jurisdiction of traditional leaders, defined as "traditional" land in the 2011 Census.

Overall, these findings suggest that we need to be cautious about generalising from the panel to society at large. However, on the employment dimension, if anything, the balanced panel paints an overly-optimistic picture of job finding rates in the South African labour market, since we observe that CSMs are more likely than the general population to be employed. In this sense, we err on the side of caution by presenting somewhat conservative results on the likelihood of finding employment. Despite these data concerns, our assessment of the data quality is that the subsequent analyses will nonetheless provide us with reasonable estimates of the labour market volatility experienced by a representative group of South Africans over the period from 2008 to 2017.

4. Employment transition patterns

In this section we summarise results from the transition trees represented in Figure 1, constructed for different subsets of the South African labour force. Table 4~~Table 4~~ presents these results, summarised by generating six dynamic employment categories based off the number of NIDS waves in which an individual is observed to be employed: Employed in five waves ("Always employed"), four, three, two or one waves, or no waves ("Never employed"). Formatte

Considering the entire sample of 3,591 adults observed in all five NIDS waves with valid employment responses, we find that 29.7 percent were observed to be employed in all five waves, with an additional 16.0 percent employed in four out of five waves. 13.3 percent of this sample had never been observed to have been employed in this period, and an additional 13.9 had not been employed in four out of five periods. 27.0 percent of the sample is observed to have transitioned into and out of employment more frequently – being observed to have been employed in either two or three waves. There is thus evidence of substantial employment volatility in the South African working-age population, as well as high levels of chronic exclusion from employment. This aggregate perspective is useful primarily as a benchmark against which to compare the labour market patterns of the specific sub-populations that we now consider.

Table 4: Number of periods employed

Number of periods employed	**Always employed**	**4**	**3**	**2**	**1**	**Never employed**	*No. of obs.*
Total	29.74%	16.02%	14.42%	12.62%	13.91%	13.29%	*3595*
African	27.60%	16.73%	14.55%	13.32%	14.20%	13.59%	*2970*
Non-African	40.21%	12.47%	14.42%	10.39%	12.63%	9.87%	*594*
< Matric	21.56%	15.67%	15.47%	14.17%	16.48%	16.65%	*2728*
Matric	37.86%	16.93%	13.08%	12.49%	11.44%	8.18%	*517*
Tertiary	60.43%	16.83%	10.81%	5.23%	3.62%	3.08%	*345*
Youth (16-24)	5.74%	11.35%	14.95%	22.81%	24.04%	21.11%	*2089*

Prime (25-50)	29.74%	16.02%	14.42%	12.62%	13.91%	13.29%	*3591*
Older (51-64)	7.97%	7.72%	12.60%	13.16%	24.24%	34.32%	*1214*
Female	20.76%	13.97%	14.71%	14.41%	18.22%	17.93%	*2428*
Male	44.28%	19.37%	13.89%	9.83%	6.85%	5.78%	*1167*
Urban	36.77%	16.80%	15.10%	10.94%	10.33%	10.05%	*1666*
Rural	15.05%	13.14%	11.79%	16.70%	22.31%	21.02%	*1313*
Poor	17.69%	16.22%	16.39%	16.12%	16.40%	17.17%	2748
Non-poor	51.16%	15.64%	10.92%	6.39%	9.48%	6.41%	843

Notes:

a) All cell proportions are weighted using Wave 5 panel weights.

b) Except for results presented for different age categories, all other results apply only to adults aged between 25 and 50 years in 2008.

c) “Non-African” identifies Whites and Coloureds. The Indian sample is small, and has been omitted.

d) Age variables are defined as described in Table 3 above.

e) “Poor” and “Non-poor” categories are defined using the StatsSA Upper Bound Poverty Line (R1,136 in March 2017 Rands) and per capita household consumption.

f) Unless otherwise stated, all restrictions applied (Column 1) are defined using Wave 1 variables (2008).

Figure 2~~Figure 2~~, a) to f), graphically present comparisons of the labour market patterns by several of the group variables that we have introduced above. Figure 2~~Figure 2~~a compares the dynamic labour market patterns of the African and non-African (White and Coloured) sample. We find that the non-African sample is far more likely to be employed in all five periods than the African sample, at 40.2 compared to 27.6 percent. The non-African sample is also less likely to experience chronic exclusion from the labour market. In addition, employment volatility appears slightly more prevalent for the African compared to the non-African sample, suggested by the relatively higher proportion of those experiencing two or three spells of employment. It should be noted that, because of the small sample size of the White sample, we could not make meaningful comparisons between White and African samples, which would presumably yield a much more striking contrast compared to this comparison, in which Coloured participants are included among the non-African sample.

The differences in volatility patterns by gender are striking. Males are more than twice as likely to be employed in all five periods than women, and less than one third as likely to *not* be employed in all five periods. 63.7 percent of men were employed in four or five periods, compared to only 34.7 percent of women, while as many as 17.9 percent of women were not employed in all five periods compared to only 5.8 percent of men. This indicates that men who experience unemployment in the South African labour market are far more likely to experience it as a transient state than unemployed women. For over one third of women, exclusion from the labour market is experienced as a persistent state (4 or 5 periods).

Women are also more likely to be affected by employment volatility than men, suggesting that finding a job is seldom experienced as a route into stable employment.

Figure 2: Number of periods employed

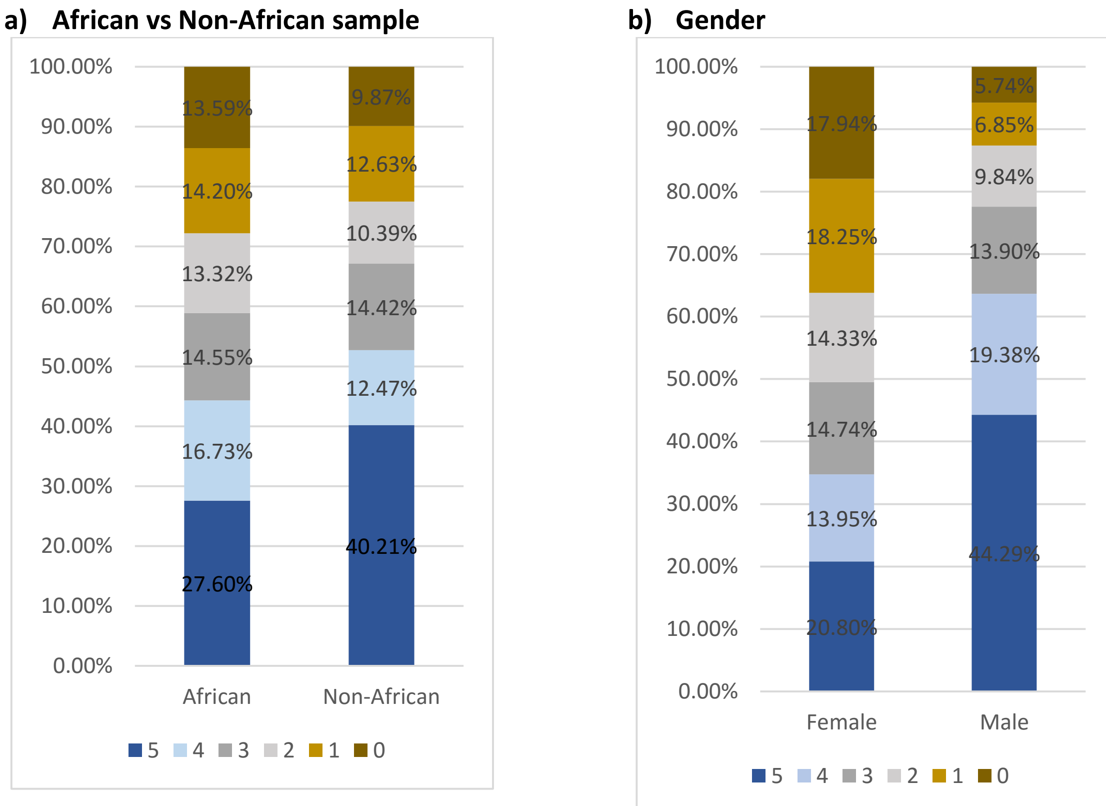

We also investigate how differences in educational attainment are related to labour market dynamics. Unsurprisingly, low educational attainment appears to be a strong predictor of chronic exclusion from (or weak attachment to) the labour market. Only about one in five of those with less than a matric were consistently employed, while one in three were not employed in all or all but one period. In contrast, those with post-secondary qualifications are far more likely to have been consistently employed, with 60.4 percent of the sample employed in five periods, and only 6.7 percent not employed in four or five periods. Of those with less than secondary schooling who did find employment, this was disproportionately likely to be experienced as a transient state, with 29.6 percent being employed in two or three periods, compared to only 16.0 percent of those with tertiary education experiencing these levels of employment transience. Those with completed secondary schooling are clearly better off in the labour market than those with incomplete secondary schooling and the population as a whole, but still fall short of those with tertiary qualifications. This suggests that there are large returns to tertiary education in terms of dynamic employment prospects.

Dynamic employment patterns display an interesting and complex association with age. In this part of the analysis we broaden the age range of respondents that we consider, where in other parts of this analysis we limit ourselves to those aged between 25 and 50 years old in 2008. This explains why the sample sizes for this part of the analysis are larger than elsewhere. Prime-aged workers appear most advantaged from a dynamic perspective. Young South Africans (16-24) are highly unlikely to be consistently employed, much more likely to experience employment as a transient state, and also more likely to experience unemployment as a chronic state. Older workers also compare unfavourably to prime-aged workers in the same respect. There are several potential explanations for these findings: First, it is likely that these differences are due partly to life-cycle dynamics: Younger workers are more likely to be at school or university, move more frequently between jobs, and are more likely to be able to tolerate spells of unemployment than older workers (Zizzamia, 2018). Older workers, on the other hand, may be moving out of stable employment into retirement, which may apply especially to more physically demanding labour. However, a second explanation may be linked more to historical investments in education and the more recent rapid entry of young people into the labour market. Older workers are much more likely to be poorly educated because of not having benefitted from the expansion of education in recent decades, placing them at a disadvantage relative to younger and prime aged workers when competing for jobs. Younger workers, on the other hand, are disadvantaged by the rapid increase in labour supply in recent years, especially for low-skill jobs. For many of these youth, unemployment is experienced as a trap – where an initial spell of unemployment creates negative signals to employers and leads to the depreciation of human capital built up through schooling, thereby rendering it even more difficult to find work in the future. In this case, the simple life-cycle narrative may provide an inadequate explanation for differences in the outcomes of younger, prime-aged and older workers in the South African labour market.

Finally, it is also worth noting that the unfavourable labour market experiences of younger and older workers reveal that a consequence of limiting our analysis elsewhere to initially prime-aged adults may lead to a substantially more optimistic picture than what would be the case if we included a broader age range in our analysis. Our results for other demographic groups should be interpreted with this in mind.

Differences in labour market dynamics along the rural/urban divide are similarly dramatic. As is well established in the literature, chronic unemployment is above all a rural phenomenon in South Africa, and this is also the primary determinant of the chronic poverty observed in rural areas (Seekings and Nattrass, 2005; Schotte et al., 2018). Our analysis of NIDS data confirms this, where we observe that 43.3 percent of rural adults were not employed in four or five periods, compared to 20.4 percent of urban adults. Only 15.1 percent of rural adults were consistently employed over the period, compared to 36.8 percent of urban adults. While rural adults are also more likely to experience volatile spells of employment (i.e. being employed in two or three spells out of five), the differences are not as large as those between the other groups discussed. This simply highlights that the primary urban/rural divide in terms of employment dynamics is the high concentration of chronic unemployment in rural areas relative to urban areas. These findings suggest that the employment challenges in urban areas are relevantly different to those in rural areas. As suggested by Schotte et al. (2018) and Zizzamia (2018), it seems that an important challenge in urban labour markets is to reduce the precariousness of employment and address labour market frictions, while in rural areas job creation as a means of overcoming chronic unemployment is a priority.

c) Education

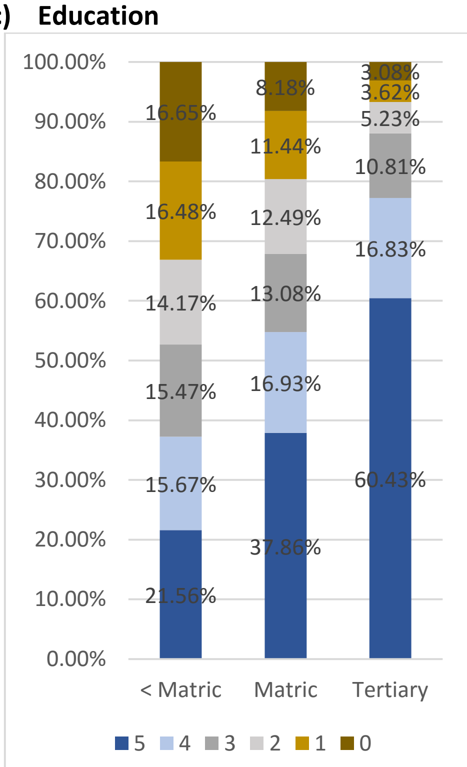

d) Age

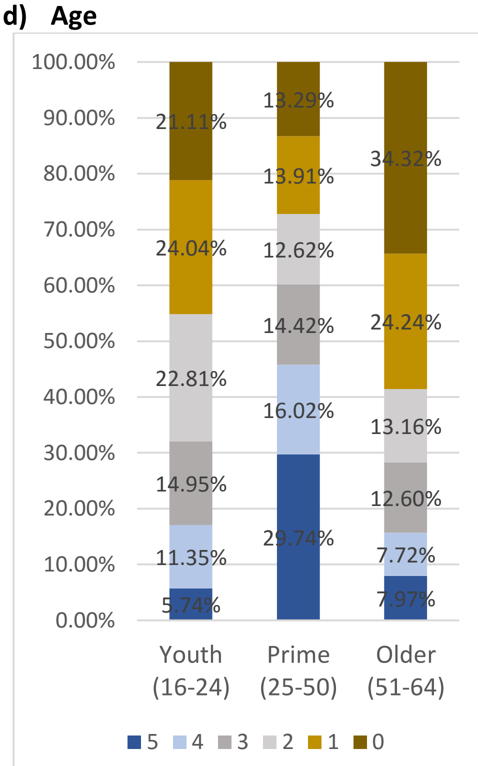

In recent research on poverty dynamics in South Africa, Finn and Leibbrandt (2017) and Schotte et al. (2018) have emphasised the importance of initial poverty status in determining poverty persistence. That is, the experience of poverty may itself, independently of other factors, increase the likelihood of remaining poor. One mechanism through which this may operate is if being poor compromises one's ability to find and keep work, either because of a difficulty in funding job search costs, or because of a psychological burden imposed by poverty (Stoop et al., forthcoming). While Table 4~~Table 4~~ provides a purely descriptive analysis, it is nevertheless revealing to observe how initial poverty status (as measured in 2008) is closely associated with employment patterns over the 2008 to 2017 period. The initially non-poor are more than twice as likely to have remained employed over five periods, while the initially poor are about three times as likely as the initially non-poor to remain unemployed in all five periods. Employment volatility among the initially poor is twice as high as it is for the initially non-poor, suggesting that the poor are being selected into more precarious forms of employment, or for other reasons are unable to sustain their attachment to the labour market, thereby further compromising their ability to use the labour market as a means to escape poverty.

e) Location

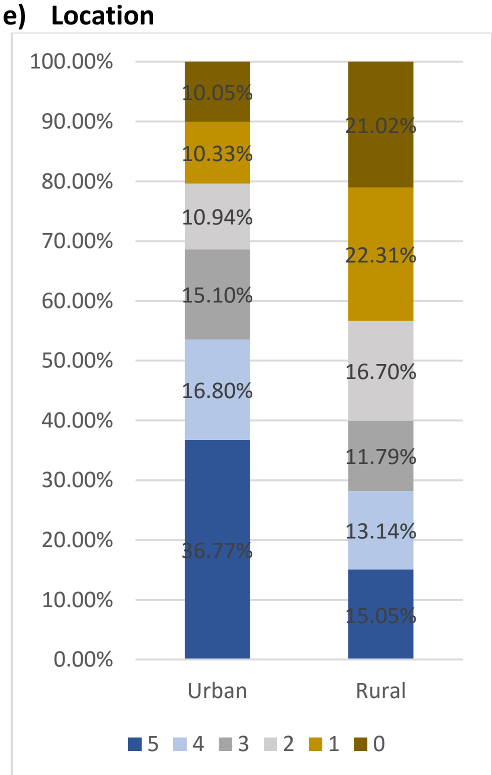

f) Initially poor vs intiially non-poor

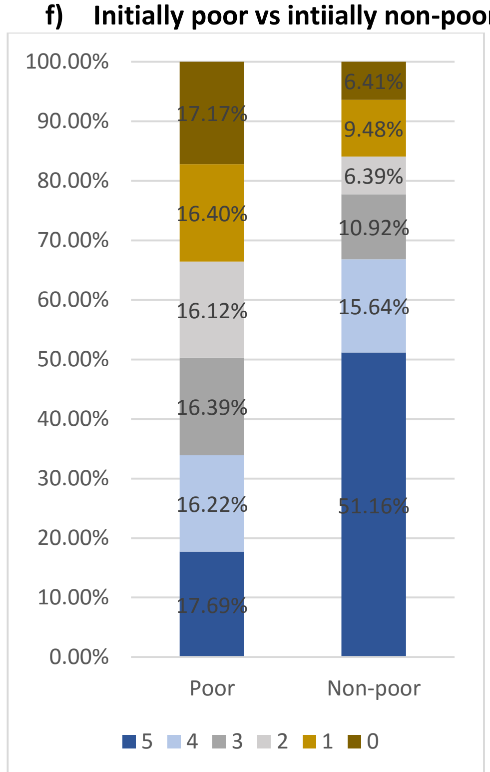

In Table 5~~Table 5~~ we investigate how prospects of employment differ by initial employment status for our differently defined demographic groups. Overall, we see that initial employment status is a strong predictor of remaining employed, with 70.4 percent of those initially employed remaining employed in at least three of the four subsequent periods. Of those initially not employed, only about one in four found employment in at least three of the four subsequent periods.

Across all groups, the patterns are as expected: Black and rural South Africans, those without matric, the youth and elderly, women and the initially poor, even when initially employed, are significantly less able to maintain their employment over the subsequent nine years, indicating high levels of employment precariousness for these groups. In addition, for those who were initially not employed in 2008, these same groups are much less likely to find stable employment in subsequent periods. These findings highlight how employment dynamics can deliver a "double whammy" blow to these groups: Not only is *maintaining* employment disproportionately difficult, but *finding* employment again once having lost it is too.

Table 5: % of sample~~-~~ employed in 3 or 4 subsequent periods, by initial employment status

	Employed in 2008	**Not employed in 2008**	No. of obs.
Total	70.42%	27.59%	3591
African	69.75%	28.40%	2970
Non-African	72.19%	23.33%	594
< Matric	62.20%	24.41%	2728
Matric	78.58%	35.71%	517
Tertiary	88.43%	49.39%	345
Youth (16-24)	56.85%	22.35%	2089
Prime (25-50)	70.42%	27.59%	3591
Older (51-64)	31.66%	5.21%	1214
Female	62.25%	25.04%	2428
Male	78.64%	36.49%	1167
Urban	74.23%	33.64%	1666
Rural	54.76%	17.71%	1313
Poor	54.76%	54.79%	2748
Non-poor	81.12%	53.25%	843

Notes:
a) All cell proportions are weighted using Wave 5 panel weights.
b) All variables defined as described in Table 4.

Table 6~~Table 6~~ approaches the same issue from a different perspective. In this table, we investigate how prospects of *exclusion* from the labour market differ by initial employment status for our various demographic groups. Again, we find that African and rural South Africans, those without matric, the youth and elderly, women and the initially poor are all much more likely to fall into sustained exclusion from the labour market, even when initially employed. This affords us another perspective through which to understand which demographic groups are most affected by precarious forms of work. We ought to be especially concerned about the high apparent precariousness experienced by young, female and rural workers.

*Table 6: % of sample **not** employed in 3 or 4 subsequent periods, by initial employment status*

	Employed in 2008	Not employed in 2008	No. of obs.
Total	16.55%	54.49%	3595
African	17.27%	53.51%	2970
Non-African	14.23%	58.25%	594
< Matric	21.70%	59.24%	2728
Matric	13.06%	40.12%	517
Tertiary	3.84%	28.44%	345
Youth (16-24)	24.00%	53.18%	2089
Prime (25-49)	16.55%	54.49%	3591
Older (50-64)	47.43%	87.59%	1214
Female	23.73%	58.81%	2428
Male	9.38%	39.39%	1167
Urban	12.72%	49.24%	1666
Rural	32.15%	63.72%	1313
Poor	22.33%	54.07%	4020
Non-poor	11.19%	46.24%	1052

Notes:
a) All cell proportions are weighted using Wave 5 panel weights.
b) All variables defined as described in Table 4.

5. Estimating labour market volatility

While we have a clear understanding of the demographic correlates of employment stability, volatility and chronic exclusion from the labour market, one of the challenges in the analysis thus far has been that these groups are not independent of each other. For instance, rural workers are more likely to be African and poorly educated, while White workers are more likely to be urban and highly educated. In the preceding descriptive analysis, the correlations between these characteristics makes isolating the effect of any one characteristic on employment patterns impossible. To address some of these concerns, in this section we estimate two simple probit models off the basis of several characteristics, including those used in the group-based analysis above. One probit model is used to predict the probability of gaining employment and the other is used to model the probability of losing employment. This allows us to obtain a measure of the correlation between a demographic characteristic and the likelihood of gaining/losing employment, under the assumption that the values of all the other variables are held constant. However, since identifying causation cleanly in a simple probability model is still fraught with challenges, the analysis which follows cannot be interpreted as anything more than a refined descriptive analysis.

The first model, which predicts job-loss (Table 7~~Table 7~~), is run on 17,421 observations under the base specification and successfully predicts job losses in 61,5 percent of cases.[5] In a second specification, we add controls for type of employment (employed, self-employed, casual worker, subsistence worker). Under this specification, the number of observations drop to 16,601, and job losses are correctly predicted for 65,7 percent of observations. In a third specification, controls for employment type are dropped and replaced by a dummy for whether an employee's contract is written or unwritten, a dummy for union membership, and a dummy for whether the employment agreement is permanent or of temporary/fixed-duration. Observations drop to 12,258 under this specification, while percent of job losses correctly predicted drops to 61.5 percent. In a final specification, we drop all employment controls and replace these by a single dummy for formal or informal employment in t. The model is run on 16,417 observations, with 61.4 percent of job losses correctly predicted. Because of the relatively larger sample size and superior predictive power, the second specification is preferred. The marginal effects of all four specifications are reported in Table 7~~Table 7~~.

All the variables that we include (except household size) are highly significant. This is unsurprising, since we have already established that these groups experience different employment prospects. People in urban households are between 8.8 and 22.9 percent less likely to experience a job loss than their rural counterparts, depending on the specification used. It is also worth noting that individuals in households in which grants or other sources of income are more important than labour income, or in which there are other wage earners, are more likely to lose employment.

Strikingly, even after controlling for other relevant factors, women are between 24.2 and 31.2 percent more likely than men to fall out of employment between periods. We also confirm the finding in the previous section that higher education is a strong predictor of employment stability, where we find that those with tertiary education are between 34.6 and 56.9 percent less likely to lose employment between periods. White workers are between 22.1 and 24.0 percent less likely to fall out of employment

[5] While the model correctly predicts employment status in t in 75.9 percent of cases, this is distorted upward by the relative ease in predicting job retentions vis-à-vis job losses. For this reason, we prefer to report the percentage of job *losses* correctly predicted since, while this is more conservative, it also is a better indicator of the predictive power of the model.

that black Africans. Again, consistent with the analysis in the previous section, there is evidence of a concave relationship between age and the likelihood of job loss, suggesting that prime-aged workers are better-off in the labour market. Finally, we also find evidence that employment type and the nature of the employment contract matter for employment security.

Table 7: Average marginal effects on probability of falling out of employment between t-1 and t

Specification:	I	II	III	IV
Household characteristics (t-1)				
Household size	-0.002	0.002	0	0
Geographic location (base: traditional)				
Urban	-0.229***	-0.136***	-0.088**	-0.197***
Farms	-0.494***	-0.322***	-0.266***	-0.443***
Primary source of household income (base: Labour)				
Grants	0.522***	0.273***	0.297***	0.416***
Other	0.617***	0.32***	0.398***	0.506***
No. of other employed household members	0.085***	0.045***	0.049***	0.082***
Number of state pensions/disability grants in household	0.034	0.041	0.021	0.035
Individual characteristics (t-1)				
Female	0.301***	0.312***	0.242***	0.274***
Education (base: Incomplete secondary)				
No schooling	0.215***	0.177***	0.063	0.133***
Incomplete primary	0.109***	0.089**	0.009	0.068*
Primary education	0.063	0.046	-0.042	0.019
Secondary education	-0.237***	-0.184***	-0.119**	-0.153***
Tertiary education	-0.569***	-0.485***	-0.346***	-0.408***
Race (base: Black African)				
Coloured	-0.12***	-0.087**	-0.102**	-0.096**
Indian/Asian	0.04	-0.015	-0.011	0.099
White	-0.24***	-0.307***	-0.27***	-0.221***
Age	-0.139***	-0.134***	-0.133***	-0.134***
Age squared	0.002***	0.002***	0.002***	0.002***
Individual is Household head/spouse of household head	-0.085***	-0.098***	-0.105**	-0.074**
Employment characteristics				
Employment type (base: Employee)				
Self-employed		0.62***		
Casual		0.617***		
Subsistence		1.332***		
Employment agreement				
Written contract			-0.055*	
Permanent contract			-0.319***	
Union member			-0.206***	
Formal sector worker				-0.482***

Year and province fixed effects	YES	YES	YES	YES
Constant	2.206	1.86	1.967***	2.26***
Log-likelihood	-8930.55	-8159.29	-5418.07	-8218.18
Model chi2	1789.365	2249.849	901.754	2041.516
Number of observations	17421	16601	12258	16417
Positive predictive value (% job losses correctly predicted)	61.48%	65.69%	61.54%	61.44%
Negative predictive value (% job losses retentions predicted)	77.37%	78.80%	81.83%	78.21%
Percentage correctly predicted	75.90%	77.28%	81.55%	76.31%

* p<0.10, **p<0.05, *** p<0.01

Notes:

a) Author's calculations using NIDS waves 1 to 5 pooled panel of wave-to-wave transitions.
b) Weights not applied.
c) All explanatory variables defined in *t-1.*
d) Sample upon which regression is run is restricted to those aged between 22 and 60 in *t-1.*
e) Standard errors are clustered at the individual level

In Table 8~~Table 8~~ we report marginal effects for the probit model predicting employment gains. Since this model is run on a sample of respondents who initially do not have employment, we clearly cannot control for employment type, as in the previous model. Instead, we vary the outcome variable, and estimate a first model predicting a gain of employment (without distinguishing whether this is formal or informal) and a second model predicating formal employment gains specifically. The first specification (Column 1 of Table 8~~Table 8~~) is run on 15,545 observations and successfully predicts job gains in 56,1 percent of cases. In the second model observations drop to 12,845, and job gains are correctly predicted for only 52,5 percent of observations. Both these models perform disappointingly in terms of predictive power. However, this is to be expected, since we do not distinguish between the not-employed who are searching for employment versus those who are not.

As in Table 7~~Table 7~~, all the variables that we include in the models in Table 8~~Table 8~~ (except household size) are highly significant, and the signs and magnitude of marginal effects are as expected, largely confirming the analysis in Section 4. A few findings are worth emphasising: Women are again shown to be at an enormous disadvantage from a dynamic perspective, being 37.7 percent less likely to find employment than men, controlling for other relevant factors. Those with matric or tertiary education are also far more likely to find (formal) employment than those with incomplete secondary education. Those in urban areas are 17.8 percent more likely to gain formal employment.

One unexpected result is the negative and significant coefficient on the White race variable – suggesting that white people are between 25.3 and 41.0 less likely to gain employment if initially not

Formatte
Formatte
Formatte
Formatte

employed. However, this result is likely being driven by whites who are not economically active by choice and the low unemployment rate among this demographic.

Table 8: Average marginal effects on probability of gaining employment or formal employment between t-1 and t

Outcome variable	Gain employment	Gain formal employment
Household characteristics (t-1)		
Household size	0.004	0.002
Geographic location (base: traditional)		
Urban	0.093***	0.178***
Farms	0.216***	0.223***
Primary source of household income (base: Labour)		
Grants	-0.065*	-0.112**
Other	-0.067	-0.003
No. of other employed household members	-0.048***	-0.038
Number of state pensions/disability grants in household	-0.119***	-0.131***
Individual characteristics (t-1)		
Female	-0.377***	-0.322***
Education (base: Incomplete secondary education)		
No schooling	-0.201***	-0.422***
Incomplete primary	-0.131***	-0.326**
Primary education	-0.134***	-0.163**
Secondary education	0.171***	0.384***
Tertiary education	0.354***	0.662***
Race (base: Black African)		
Coloured	0.074	0.062
Indian/Asian	-0.334**	-0.348*
White	-0.253**	-0.41**
Age	0.07***	0.054***
Age squared	-0.001***	-0.001***
Individual is Household head/spouse of household head	0.068**	0.048
Year and province fixed effects	YES	YES
Constant	2.014***	1.402***
Log-likelihood	-8687.16	-4344.72
Model chi2	937.809	930.418
Number of observations	15545	12845
Positive predictive value (% job losses correctly predicted)	56.08%	52.54%
Negative predictive value (% job losses retentions predicted)	73.00%	87.28%
Percentage correctly predicted	72.40%	87.12%

* p<0.10, **p<0.05, *** p<0.01

Notes:

a) Author's calculations using NIDS waves 1 to 5 pooled panel of wave-to-wave transitions.

b) Weights not applied.

c) All explanatory variables defined in *t-1*.
d) Sample upon which regression is run is restricted to those aged between 22 and 60 in *t-1*
e) Standard errors are clustered at the individual level.

6. Conclusion

In this paper, we make use of five waves of a nationally representative longitudinal dataset to explore the degree of employment volatility in the South African labour market. Over the period from 2008 to 2017, we find a substantial amount of volatility in the likelihood of finding or losing employment. Our primary finding is that, instead of a binary classification of labour market participants into employed or not employed groups, it would be more accurate to conceptualise a labour market with three groups: Stable employment, volatile employment and persistent unemployment. This middle group, the people who experience repeated spells of employment and non-employment, could not have been identified using cross-sectional data alone.

Overall, we estimate that this group makes up about 27% of the prime aged individuals in the economy. In addition, the experience of volatile employment is not uniformly distributed amongst variously defined demographic groups. Africans experience more employment volatility than non-Africans, women experience more employment volatility than men, while more highly educated individuals experience less employment volatility than people with lower levels of educational attainment. Youth, in particular, experience very high levels of employment volatility. The difference between urban and rural sub-populations is mostly in terms of the proportions in either stable employment or non-employment, with the proportion in the rural areas that experience chronic unemployment being more than double the corresponding proportion amongst urban dwellers. These bivariate findings were maintained in our multivariate regression analyses.

Our findings raise some important questions for both labour market policy as well as for further research. In terms of future research directions, there are several avenues to be explored. To what extent are the experiences of the different groups a function of individual characteristics, and to what extent do they reflect economic volatility in the labour market at an aggregate level? How much of the documented volatility is due to life cycle dynamics in one's labour force attachment? Do the people who experience a high degree of employment volatility tend to find undesirable forms of employment, which would then exacerbate the levels of volatility further? These are just some of the questions that researchers may find worth pursuing as part of a broader research agenda aimed at understanding the South African labour market.

From a policy perspective, one has to wonder whether a more nuanced approach to assist people may be warranted. The chronically unemployed may best be assisted with skills development policies, while the people who experience high levels of employment volatility could possibly benefit more from policies that reduce search frictions in the labour market. Of course, any successful policies will change both the size as well as the composition of the groups, and this too would need to be explicitly recognised in planning a dynamically consistent strategy to improve the lives of a large proportion of South Africans.

References:

Banerjee, A., Galiani, S., Levinsohn, J., McLaren, Z., & Woolard, I. (2008). Why has unemployment risen in the new South Africa? 1. *Economics of Transition*, *16*(4), 715-740

Cafiero, C., & Vakis, R. (2006). *Risk and Vulnerability Considerations in Poverty Analysis: Recent Advances and Future Directions, (*Social Protection Discussion Paper 0610). Washington D.C.: The World Bank.

Cichello, P. L., Fields, G. S., & Leibbrandt, M. (2005). Earnings and employment dynamics for Africans in post-apartheid South Africa: A panel study of KwaZulu-Natal. *Journal of African Economies*, 14(2), 143-190.

Paul Cichello, Murray Leibbrandt & Ingrid Woolard (2014) Winners and losers: South African labour-market dynamics between 2008 and 2010, *Development Southern Africa*, 31:1, 65-84, DOI: 10.1080/0376835X.2013.853612

Essers, D. (2016). South African labour market transitions since the global financial and economic crisis: Evidence from two longitudinal datasets. *Journal of African Economies*, 26(2), 192-222.

Finn, A., & Leibbrandt, M. (2017). *The dynamics of poverty in the first four waves of NIDS*. SALDRU Working Paper Number 174/ NIDS Discussion Paper 2016/1. Cape Town, South Africa: Southern Africa Labour and Development Research Unit (SALDRU), University of Cape Town.

Kerr, A. (2018), Job Flows, Worker Flows and Churning in South Africa*. South African Journal of Economics*, 86: 141-166.

Kingdon, G. G., & Knight, J. (2004). Unemployment in South Africa: The nature of the beast. *World development*, 32(3), 391-408.

Klasen, S., & Waibel, H. (Eds.) (2013). *Vulnerability to Poverty: Theory, measurement and determinants, with case studies from Thailand and Vietnam*. Palgrave Macmillan.

Leibbrandt, M., Bhorat, H. & Woolard, I. (2001), Understanding Contemporary Household Inequality in South Africa, Chapter 2 in Bhorat, H., Leibbrandt, M., Maziya, M., van der Berg, S. & Woolard, I. (2001). *Fighting poverty: Labour markets and inequality in South Africa*. UCT Press: Cape Town.

McKee-Ryan, F., Song, Z., Wanberg, C. R., & Kinicki, A. J. (2005). Psychological and physical well-being during unemployment: a meta-analytic study. *Journal of applied psychology*, *90*(1), 53.

Mncwango, B. 2016. *Public Attitudes to Work in South Africa*. LMIP Report 16. Pretoria: Labour Marker Intelligence Partnership.

Pugatch, T. (2018). Bumpy Rides: School-to-work Transitions in South Africa. *Labour*, 32(2), 205-242.

Ranchhod, V. (2013). *Earnings volatility in South Africa*. Cape Town: SALDRU, University of Cape Town. SALDRU Working Paper Number 121/ NIDS Discussion Paper 2013/3.

Ranchhod, V. and Dinkelman, T. (2008) *Labour Market Transitions in South Africa: What can we learn from matched Labour Force Survey data?* A Southern Africa Labour and Development Research Unit Working Paper Number 14. Cape Town: SALDRU, University of Cape Town

SALDRU (2018a). National Income Dynamics Study 2008, Wave 1 [dataset]. Version 5.3. Cape Town: SALDRU [producer],Cape Town: DataFirst [distributor].

SALDRU (2018b). National Income Dynamics Study 2010-2011, Wave 2 [dataset]. Version 2.3. Cape Town: SALDRU [producer], Cape Town: DataFirst [distributor].

SALDRU (2018c). National Income Dynamics Study 2012, Wave 3 [dataset]. Version 1.3. Cape Town: SALDRU [producer], Cape Town: DataFirst [distributor].

SALDRU (2018d). National Income Dynamics Study 2014-2015, Wave 4 [dataset]. Version 1.0. Cape Town: SALDRU [producer], Cape Town: DataFirst [distributor], Pretoria: Department of Planning Monitoring and Evaluation [commissioner].

SALDRU (2018d). National Income Dynamics Study 2014-2015, Wave 4 [dataset]. Version 1.0. Cape Town: SALDRU [producer], Cape Town: DataFirst [distributor], Pretoria: Department of Planning Monitoring and Evaluation [commissioner].

SALDRU (2018e). National Income Dynamics Study 2017, Wave 5 [dataset]. Version 1.0.0 Pretoria: Department of Planning, Monitoring, and Evaluation [funding agency]. Cape Town: Southern Africa Labour and Development Research Unit [implementer], Cape Town: DataFirst [distributor].

Schotte, S., Zizzamia, R., & Leibbrandt, M. (2018). A poverty dynamics approach to social stratification: The South African case. *World Development*, 110, 88-103.

Seekings, J., & Nattrass, N. (2005). *Class, race, and inequality in South Africa*. Yale University Press.

Stoop, N., Leibbrandt, M. & Zizzamia, R. (2018) *Exploring psychological well-being and poverty dynamics in South-Africa: evidence from NIDS waves 1-5.* Cape Town: SALDRU, UCT. Mimeo.

Witte, H. D. (1999). Job insecurity and psychological well-being: Review of the literature and exploration of some unresolved issues. *European Journal of work and Organizational psychology*, *8*(2), 155-177.

Zizzamia, R. (2018). *Is employment a panacea for poverty in South Africa? A mixed-methods investigation*. Cape Town: SALDRU, UCT. (SALDRU Working Paper Number 229).