

Fifteen Years On: Household Incomes in South Africa *

Murray Leibbrandt
University of Cape Town

James Levinsohn
Yale University and NBER

Abstract: This paper uses national household survey data to examine changes in real per capita incomes in South Africa between 1993 and 2008; the start and the end of the first fifteen years of post-apartheid South Africa. These data show an increase in average per capita real incomes across the distribution. Over this period growth has been shared, albeit unequally, across almost the entire spectrum of incomes. However, kernel density estimations make clear that these real income changes are not dramatic and inequality has increased. We conduct a series of semi-parametric decompositions in order to understand the role of endowments and changes in the returns to these endowments in driving these observed changes in the income distribution. This analysis highlights the positive role played by changes in endowments such as access to education and social services over the period. If these endowment changes were all that changed in South Africa over the post-apartheid period, we would have seen a pervasive rightward shift of the distribution of per capita real incomes. In the rest of the paper we explore why this did not happen.

*We thank the NBER Africa Project for funding. We are also grateful to Matthew Welch for outstanding assistance with the data.

1 Introduction

Measuring South African economic growth since the fall of Apartheid is a tricky business. One can simply measure GDP per capita and there the picture is a bright one. Real GDP per capita since the democratic elections in 1994 has risen an average of close to 1.5 percent per year. Individuals, though, can't really spend GDP when they go to the store. Rather, they spend *their* incomes. One can instead measure individual incomes, but this measure too is problematic. Examining the distribution of individual incomes will typically not speak to the welfare of the roughly 40 percent of South Africans age 18 and younger. In this paper, we measure incomes at the household level (adjusting for household size.) This measure encompasses all household members, even those not participating in the labor market, while still capturing a measure of economic welfare at the level of individuals. Our reasoning is that to the extent that real household per capita incomes increase, households are generally economically better off in a narrow but well-defined sense. With real household per capita income as our metric, we measure economic growth in South Africa from 1993 to 2008.

Our approach is a very microeconomic one. We rely on two nationally representative surveys of individuals. Our data, though, do not make up a panel, as such longitudinal data simply do not exist over the time period under consideration. Rather, we have used nationally representative household surveys from 1993 and 2008 and meticulously matched definitions of incomes so that we are confident that the temporal comparisons are valid. Because we rely on micro-data, we are able to both measure the changes in incomes and investigate what explains these changes. Additionally, we are able to examine changes throughout the entire distribution rather than focusing simply on a mean or median. We do so using relatively new nonparametric techniques augmented by more traditional para-

metric estimates.

Whether the news is good or bad surely depends on one's prior and the previous evidence is sufficiently diverse that it's hard to know just what constitutes a happy story. As noted above, the national income accounts tell a story of success. While the macroeconomy has shown robust growth over most of the past 15 years, it has been a period of relatively little job growth and unemployment has increased dramatically. Depending on the measure used, unemployment has increased from around 15 percent to well over 30 percent. Exactly how the growth in GDP together with the rise in unemployment has impacted households is something of an open question. In an earlier paper, we document that the first five years after the new government (from 1995 to 2000) saw real individual incomes decline almost forty percent. See Leibbrandt, Levinsohn and McCrary (2010c). Hoogeveen and Ozler (2004) also found the the first five years after transition were especially tough on the poor as poverty increased and household expenditures at the lower end of the distribution fell in real terms. With the recent release of a new nationally representative income survey, the dismal but provisional picture painted by these earlier studies merits revisiting.¹

In the next section, we describe our data. Section 3 describes the changes in real incomes from 1993 to 2008. Section 4 investigates what underlies these changes, while Section 5 concludes.

¹Using much of the same data that we employ in this paper, Leibbrandt, Woolard, Finn and Argent (2010b) examined the changes in inequality and poverty from 1993 to 2008. They found that inequality had increased while aggregate poverty had declined slightly. The authors did find some hopeful trends in indicators of non-monetary well-being (e.g. access to piped water, electricity, and formal housing.)

2 The Data

2.1 The 1993 Data

We benchmark incomes at transition using the LSMS household survey conducted by the World Bank in 1993. This survey is well-vetted and has been used by many researchers; including Case and Deaton (1998), Duflo (2003), and Thomas (1996). The survey was nationally representative and included about 44,000 individuals comprising just over 8800 households. One reason for this survey's widespread use is that it serves as a benchmark for what South Africa looked like on the eve of transition. Also, this survey has not been subject to some of the criticisms leveled at a plausible substitute survey, the 1995 Income and Expenditure Survey. We have elected to simply bypass that issue by using the 1993 LSMS survey.

Rather than using the widely available and easily downloaded merged version of the 1993 data, we have gone back to the original source data. We have done so because we want to be confident that our comparisons to 2008 are valid. This means making sure that every component of income is comparably defined in each of the two surveys— something we have taken great care to do. For this reason, we do not include imputed housing in our measurement of imputed income.

Especially for poorer households, the value of housing can represent a substantial fraction of real income. Most households do not report the value of the flow of housing they receive from their residence when they own it. It is of course possible to impute the value of housing and indeed one of us was responsible for this task for the current National Income Dynamics Survey. If we could be confident that the housing imputation used in 2008 (which we designed) could be applied to the 1993 data to construct a housing value that would

then be comparable to that used in 2008, we would do so. Because we are not able to do this, we strip housing out of our income measures for 2003 and 2008. We note, though, that if we include housing using the probably non-comparable definitions from 1993 and 2008, we find larger increases in household per-capita income.

2.2 The 2008 data

The most recent nationally representative income data for South Africa come from the first wave of the National Income Dynamic Survey (NIDS.) This survey, like the 1993 survey, is publicly available, free, and readily downloadable.² We use Wave 1 of the NIDS. These data were collected in 2008 and comprise the initial wave of what will be a national panel study. As was the case with the 1993 data, we use the original source data and then construct aggregates so as to ensure comparability with the 1993 data.

The 2008 NIDS includes data on 28,225 individuals comprising 7305 households. As noted above, we exclude the value of housing from our definition of income. The data include detailed expenditure data as well as income data. We focus in this paper on the latter.

2.3 Why log household per-capita income?

Throughout this paper, our analysis is focused on what happened to incomes at the household level. We have made this decision for a couple of reasons. First, we are trying to capture what happened to economic welfare at a national level using micro data. The obvious alternative to a household-level analysis is an individual-level analysis. An individual-level analysis has some advantages. It allows the researcher to investigate issues of income re-

²See <http://www.datafirst.uct.ac.za/home/index.php?/Metadata-and-Data-Downloads> to obtain the source data.

ciency and to better understand how the labor market works (or not). The drawbacks to the individual-level analysis are that it excludes children (who comprise almost half the population) and, if one elects to work with log incomes as is often done, the analysis excludes all those adults who did not receive income. The reciprocity issue can be addressed with careful econometric analysis, but the exclusion of children is part-and-parcel of any analysis of individual incomes. Because we want to better understand economic welfare at the national level, we elect the household-level approach and hence include children.

The household-level approach has the advantage of making moot most issues around reciprocity. Almost all households report at least some income, be it from remittances, grants, labor market earnings, or informal activities. We have elected to work with per-capita household incomes so as to adjust for household size. This has the obvious advantage that it corrects for household size, but it is a somewhat blunt way of dealing with changes in household composition over the course of the 15 years between the surveys. We have not employed equivalence scales and, in the analysis below, simply treat all household members equally. Table 1 reports the number of households (after applying frequency weights) in South Africa in 1993 and 2008. From 1993 to 2008, there was an increase of about 5.2 million households in South Africa with about 4.4 million of those self-identifying as “African.”³

Over this same period, household composition changed. Table 2 gives average household size by year and by population group. For all groups, the mean household size declined with the most marked decline being for African households. In only 15 years, mean household size declined from 5.30 to 3.68. This compositional shift underscores the importance of normalizing household incomes by some measure of household size when undertaking any

³We use the population group names, African, Coloured, Asian, and White to maintain consistence with the existing literature despite the fact that all South Africans are, in another sense, African.

inter-temporal comparisons.

Having decided to work with household per-capita incomes, we elect to conduct most of our analysis looking at log incomes. This has the advantage of decreasing the influence of outliers (and there are a handful especially in 2008) and of making our graphical analyses more practical.

Both surveys provide sampling weights and all of our analysis employs those weights.

3 Household Incomes from 1993 to 2008

Table 3 reports descriptive statistics on the distribution of per-capita household real incomes (hereafter “incomes” for the sake of expositional ease) in South Africa in 1993 and 2008. Focusing first on the top line of the table, the mean per capita income in 1993 was R10,741.⁴ In 2008, the comparable figure was R24,409. On the surface, this appears an impressive increase in real incomes at the household level. Not surprisingly given the inequality documented by other researchers, these mean figures hide huge heterogeneity in household welfare— both within and across population groups. The average African income increased from R6,018 in 1993 to R9,718 in 2008 and for Coloured households, the increase was from R7498 to R25,269— an almost four-fold increase. For Whites, the increase was similarly dramatic, from R29,372 to R110,195.

The standard deviation of incomes is reported in line 2 of the table and, for the population overall, this increased about ten-fold. The within-race inequality documented in Leibbrandt et al. (2010b) is evident in our data as well. For all population groups, the ratio of the mean income to its standard deviation increased from 1995 to 2008.

⁴In 2000, the exchange rate fluctuated mostly in the range of 6.5 to 7.5 Rand per US dollar.

The bottom panel of Table 3 reports percentiles of the distribution of income both overall and by population group. The median incomes (50th percentile) show increases that are substantially more modest than those of mean incomes. While mean income for all South Africans rose about 130 percent from 1993 to 2008, the median income rose just 15 percent over the same period— from R4444 to R5096. Especially for Whites, the increases are being driven by a small number of very large incomes.

Another “small numbers” issue with the data concerns zero incomes that are reported in 1993. It is somewhat hard to believe that these incomes are truly zero, especially for the White households which are more likely to report missing values. The 1993 data do, in principle, though, account for the difference between zero incomes and missing or non-reported incomes.

We elect to treat the data as it stands. In Table 3, we have not deleted the huge incomes reported nor have we set zero incomes to missing. In most of the analysis that follows, though, we work with log incomes, and this addresses each of these issues in different ways. The zero incomes are dropped and, especially given the large fraction of them that belong to White households, this strikes us as reasonable. More importantly, the huge outliers have diminished influence on means when working with log incomes. Hence, by working with log incomes, we report statistics that are both more interpretable in percentage terms and more robust to the handful of outliers.

Table 4 reports the means and distributions of log per-capita household real incomes (hereafter “log incomes”.) In 1993, the mean log income was 8.44 and it had grown to 8.58 by 2008— a 14 percent increase in real incomes. This figure is quite close to the 15 percent increase in median incomes found in Table 3. It is smaller than the 25 percent increase implied by the one and a half percent average growth rate compounded over 15 years as

indicated by the macro-data. In terms of orders of magnitude, though, the micro- and macro-data convey very similar messages.

Again, the growth in mean log incomes was not equal across the population groups. African households experiences a 26 percent point increase, while for Coloured households the figure was 8 percentage points. Asian households saw a 5 percent increase and Whites' a 28 percent increase. For all groups, the standard deviation of log incomes increased over this period.

The bottom panel of Table 4 reports percentiles. All population groups experienced increases in the median log income. Examination of the entire distribution for the overall population shows increases at each reported percentile except the first. For African households, the first percentile is the only one to report a decline in log real income— all other reported percentiles increased. For Coloured households, the gains were less pervasive. Only the top half of the reported percentiles saw increases in real incomes. The same was true for Asian households although this group is much smaller. Like African households, White households saw increases throughout the distribution except for the bottom percentile.

The overall picture painted by Table 4 is one of modest but pervasive increases in real incomes over the fifteen years since the fall of Apartheid. An important exception to this is the bottom half of the distribution of Coloured households. Anecdotes that the Coloured population has been left behind relative to the larger African population group are supported by the nationally representative data in Table 4. On the whole, though, log incomes have increased. As is to be expected given the inequality in South Africa, the increase in log incomes is but a fraction of the increase in (level) incomes.

Figures 1 and 2 display the cumulative density functions of log incomes for all South

Africans and for African households respectively. As indicated by Table 4, Figure 1 shows more modest gains, but it is still the case that in most (but not every) parts of the distribution, log real incomes were higher in 2008. Figure 2, for African households only, shows a more distinct pattern of increased log incomes throughout the distribution.

Figures 3 and 4 give the kernel density estimates of the income distributions for all South Africans and for Africans-only respectively. These are presented for two reasons. First, they better highlight relative gains of different segments of the income distribution. We illustrate this point immediately below. Second, the probability density functions (as opposed to the cumulative density functions) will serve as the basis for our investigation of what might explain the differences between the 1993 and 2008 distributions. We rely on methods developed in DiNardo, Fortin and Lemieux (1996) and Leibbrandt et al. (2010c) and those methods are based on probability density functions.

We discuss the relationship between the density functions using the African-only examples given in Figures 2 and 4. This is because the cumulative density function in Figure 2 is easier to read than that in Figure 1. The logic is the same for the density functions for the entire population (Figures 1 and 3.) Figure 2 showed that the cumulative density function for 2008 lay to the right of that for 1993 indicating gains in real income throughout the distribution. Figure 4 highlights the fact that those gains were greater for the bottom and top third of the distribution than they were for the middle third. There is a section of the 2008 distribution, from log incomes of about 7 to log incomes of about 10 for which the 2008 distribution lies mostly to the left of that for 1993 in Figure 4. Put another way, while real incomes were higher for African households throughout almost the entire distribution of income, the larger gains went to the bottom and top third of the distribution.

Having documented the changes in incomes from 1993 to 2008, we now turn to an analysis

of what explains these changes.

4 What Drives the Changes in Household Incomes?

We investigate three possible explanations for what might account for the shift in the density functions given in Figures 2 and 4. The first candidate is that endowments have changed, the second that returns to those endowments changed, and the third that the Child Support Grant explains at least the shift for the bottom half of the income distribution. Each are discussed in turn.

4.1 Does a change in endowments explain the shift in the distribution of incomes?

To investigate the role that changes in endowments might have played in shifting the distribution of log real incomes, we apply the approach of (DiNardo et al. 1996) (hereafter simply DFL.) This is a nonparametric approach and as such has both advantages and disadvantages. A key advantage is the ability to examine how a counterfactual impacts the entire distribution of income and to do so in a way that does not impose strong parametric assumptions (as, for example, is the case in Blinder (1973) and Oaxaca (1973).) A disadvantage is that the standard sort of hypothesis tests typically applied in parametric settings are not applicable to the nonparametric approach.⁵

We begin by setting notation.⁶

⁵It is possible, though, to investigate the impact of a change in only one endowment as opposed to all of them.

⁶The description of how the endowments counterfactual distribution is estimated draws from Leibbrandt et al. (2010c).

The density functions for household income in periods t and t' may be written as

$$f(y|T = t) = \int g(y|x, T = t)h(x|T = t)dx \quad (1)$$

and

$$f(y|T = t') = \int g(y|x, T = t')h(x|T = t')dx \quad (2)$$

respectively, where T is a random variable describing the year from which a given household in the *pooled* dataset of observations from both survey years is drawn, $g(y|x, T = t)$ is the density of household income evaluated at y , given that the observable attributes of the household, X , are equal to x and that the survey year is t , and $h(x|T = t)$ is the density of attributes evaluated at x , given that the survey year is t . It is perhaps helpful to think of $g(y|x, T = t)$ as the function that “translates” observable attributes into income. Were this a traditional parametric regression of household income on household endowments for a given year t , the density of household income, $f(y|T = t)$, would be analogous to the dependent variable, income; $h(x|T = t)$ would be analogous to the endowments data; and $g(y|x, T = t)$ would be analogous to the returns to those endowments.

We are interested in how the density of household (log) income changes if attributes and/or returns to those attributes changed. In this case, we are interested in how the distribution of income in period t would differ, were the endowments as they were in period t' . That is, what if households’ endowments were those that obtained in 2008 (t') instead of the actual 1993 (t) endowments? We denote this counter-factual by $f_h^{t \rightarrow t'}$; it may be written symbolically as

$$f_h^{t \rightarrow t'}(y) \equiv \int g(y|x, T = t)h(x|T = t')dx. \quad (3)$$

Notationally, the subscript “ h ” indicates that it is the density of attributes, or $h(x|T = t)$,

that is being changed from an actual to a counter-factual density. The superscript, “ $t \rightarrow t'$ ” indicates that in this counter-factual, we are going to start with *data* from period t and use statistical techniques, in particular a re-weighting scheme, to transform the actual density of attributes from the $h(x|T = t)$ that reigned in period t to the counterfactual density $h(x|T = t')$ that reigned in period t' .

The key insight from DFL is that the counter-factual in (3) is easy to implement by simply re-weighting the data. The re-weighting idea of DFL is based on the simple recognition that Bayes’ Axiom implies

$$\frac{h(x|T = t')}{h(x|T = t)} = \frac{P(T = t'|X = x)}{1 - P(T = t'|X = x)} \bigg/ \frac{P(T = t')}{1 - P(T = t')} \equiv \tau_h^{t \rightarrow t'}(x) \quad (4)$$

In words, $\tau_h^{t \rightarrow t'}(x)$ is just the ratio of the conditional odds to the unconditional odds. This is the weighting function needed to conduct the endowments counter-factual of (3). To see this, rewrite the object of interest $f_h^{t \rightarrow t'}(y)$ as

$$\begin{aligned} f_h^{t \rightarrow t'}(y) &= \int g(y|x, T = t)h(x|T = t')dx = \int g(y|x, T = t)h(x|T = t) \frac{h(x|T = t')}{h(x|T = t)} dx \\ &= \int g(y|x, T = t)h(x|T = t)\tau_h^{t \rightarrow t'}(x)dx \end{aligned} \quad (5)$$

which differs from (1) only by the weight $\tau_h^{t \rightarrow t'}(x)$. Consequently, we estimate the weighting function $\tau_h^{t \rightarrow t'}(x)$ and then compute the counter-factual (3) using a re-weighted density estimate of incomes. A recipe-style description of exactly how this is done is given in (Leibbrandt et al. 2010c).

In order to estimate the counterfactual density, it is necessary to estimate the numerator of (4) using a simple logit regression. This is a regression in which the dependent variable

is an indicator for whether the year is 1993 or 2008 and the dependent variables are the household endowments. The results of this regression are given in Table 5. Although the sole purpose of this regression is to estimate the conditional probabilities that enter the numerator of the DFL weight, the results are interesting in their own right.

The dependent variable is coded so that it is 1 if the year is 2008 and 0 if the year is 1993. The results show that conditional on other regressors, household size shrank while the fraction of households that were African increased the most followed by Coloured followed by Asian with White households as the excluded group. All of these are consistent with the simple correlations in the data. Other results (again conditional on other regressors) indicate that the number of adults with formal jobs declined, the number of adults in the household declined, the highest education level of the household rose, the likelihood that a household member was eligible for a State Old Age Pension fell, the number of children eligible for a Child Support Grant rose and the fraction of households that were Metro or Urban rose (relative to those that were Rural.) Except for the number of adults in the household, all of these variables are quite statistically significant.⁷

The estimated counterfactual is given in Figure 5. This figure is for all households. The results for only African households are quite similar. It is clear from the figure that the endowments counterfactual does not change the upper tail of the 1993 distribution at all. Thus, the actual improvement in incomes in the top tail by 2008 is not driven by changes in endowments. However, the counterfactual simulation changes the shape of the rest of the 1993 distribution fairly dramatically. The bottom two thirds of 1993 distribution shifts to the right. This implies that for all but the top-end of the 1993 distribution, incomes would have been greatly improved with 2008 endowments.

⁷There is a pretty good argument that the number of formally employed adults should not be included as a regressor and we have replicated all results without this regressor. Results are essentially identical.

This strong positive result is interesting because, superficially at least, the logit results shown in Table 5 show a mixed bag of positive and negative (conditional) endowment changes. Higher levels of urbanization, higher levels of education and smaller household sizes are potential positives. However, the declining population share of White South Africans, the lower numbers of employed members per household and the lower number of members eligible for the old age pension are negative endowment changes. The fact that the counterfactual distribution shifts well past the actual 2008 distribution implies that, in reality, some other factors offset the impact of these improved 2008 endowments. Actual income changes in the bottom tail were much smaller than simulated and improvements in the middle of the distribution did not happen at all. The change in the returns to these endowments is one such factor that could either accentuate or counter-balance the endowments effect and we now turn to this issue.

4.2 Does a change in returns explain the shift in the distribution of incomes?

An alternative explanation is that the returns to a household's endowments have changed from 1993 to 2008. Just as it was possible to simulate what the entire distribution of incomes would have been if returns were constant but endowments changed, one can simulate what the distribution of household incomes would be if endowments were constant but returns were those that obtained in 2008. We do just this using the methodology developed in Leibbrandt et al. (2010c).

We label this counter-factual by $f_g^{t \rightarrow t'}$ and note that it may be written symbolically as

$$f_g^{t \rightarrow t'}(y) \equiv \int g(y|x, T = t')h(x|T = t)dx \quad (6)$$

We again use Bayes' Axiom to derive an appropriate weight

$$\frac{g(y|x, T = t')}{g(y|x, T = t)} = \frac{P(T = t'|X = x, Y = y)}{1 - P(T = t'|X = x, Y = y)} \bigg/ \frac{P(T = t'|X = x)}{1 - P(T = t'|X = x)} \equiv \tau_g^{t \rightarrow t'}(x, y) \quad (7)$$

and note that the counter-factual distribution may be rewritten as:

$$\begin{aligned} f_g^{t \rightarrow t'}(y) &= \int g(y|x, T = t')h(x|T = t)dx = \int g(y|x, T = t)h(x|T = t)\frac{g(y|x, T = t')}{g(y|x, T = t)}dx \\ &= \int g(y|x, T = t)h(x|T = t)\tau_g^{t \rightarrow t'}(x, y)dx \end{aligned} \quad (8)$$

In practice, estimation of the weight given in (7) requires estimating the same logit as used in the endowments counterfactual and an additional logit regression in which household income is included both as a regressor itself and also interacted with all the included household attributes.

In previous work ((Leibbrandt et al. 2010c)), the returns counterfactual showed that returns to endowments played a major role in explaining the change in the distribution of individual incomes between 1995 and 2000. However, the counterfactual distribution at the household level that is shown in Figure 6 makes it clear that simulating a change in returns for the 1993 distribution had very little impact on the distribution. As with the endowment simulation, there is no change to the upper tail of the 1993 distribution. Thus the actual improvement in incomes at the top end in 2008 is explained by neither endowments nor returns to endowments.⁸ The counterfactual shifts the lower tail of the 1993 distribution to the right, but not as significantly as the actual rightward shift in the density between 1993 and 2008. Nonetheless, this lowering of mass in the density at the bottom is accommodated

⁸This finding is probably due to a violation of the common support assumption underlying the nonparametric approaches.

by some improvements in the middle of the distribution.

4.3 Does the Child Support Grant explain the shift in the distribution of incomes?

Over the entire post-apartheid period the State Old Age Pension has formed the central plank of an extensive social security system. See Case and Deaton (1998) for an early analysis of this. Over eighty percent of the elderly receive this pension. However, as this pension has been in place over the entire period at roughly constant real values, it is unlikely to have been responsible for major changes in the distribution of income. A new grant, the Child Support Grant (CSG) was introduced in April 1998. It initially provided R100 for every child in the household younger than seven years of age. Over time, it became both more generous (the grant is rose to R240 per child) and more pervasive as the means test was relaxed and the age below which a child qualified was raised to 15 in January 2008. By April 2009, 9.1 million children were benefiting from Child Support Grants.⁹ In short, in the period between 1993 and 2008, the Child Support Grant became a significant income source for poorer households. In this section, we investigate the role the CSG might have played in explaining the difference in the distribution of real household incomes.

The returns explanation begins to address this issue because we have included the number of children who would qualify for the CSG as a household “endowment” or attribute. The CSG acts to increase the return to this household attribute. It is not possible, though, to estimate the counterfactual density that would obtain if the return to only one attribute changed.¹⁰

In order to investigate the impact of the CSG alone, we have simply computed what

⁹This figure is from Treasury (2010).

¹⁰The reason for this is explained in detail in Leibbrandt et al. (2010c).

household incomes would have been but-for the CSG by subtracting this source of income from 2008 household incomes. The results are reported in Figure 7. This figure gives the level of income for incomes below the median 2008 household income (including the CSG.)

Figure 7 shows that the CSG has played an important role in increasing incomes for poorer households. By comparing the actual 2008 density from that which would obtain but for the CSG, it is clear that while the CSG has benefited all income levels below the median, the benefit is larger the poorer the household. This is evidenced by the fact that the gap between the actual and but-for-the-CSG incomes is larger the poorer the household. Indeed, without the CSG, there would have been about three times as many households reporting zero incomes.¹¹ For most income levels in Figure 7, the but-for-the-CSG density lies below the 1993 density. This suggests that the CSG more than explains the income gains by households below the median income level. We conclude that the CSG has played an important role in explaining why incomes increased for the bottom half of households.

5 Conclusions: Elements of success but is it sustainable going forward?

This paper is based on national household surveys conducted in 1993 and 2008. These years mark the start and the end of the first fifteen years of post-Apartheid South Africa. The data are constructed so as to insure that that the two years are comparable. What does this comparison show? The data show an increase in average per capita real incomes. For the

¹¹Figure 7 is presented in levels rather than in logs so as to make this point. The issue of zero incomes is brushed aside when working with log incomes.

most part, this increase is evident across the distribution. This means that growth has been shared, albeit unequally, across almost the entire spectrum of incomes. This is especially true for the African group that makes up close to eighty per cent of the population. We cite evidence from other researchers that this income improvement was accompanied by strong improvements in access to important services such as water, housing and electricity. Thus, there are elements of genuine success.

However, as the kernel density estimations that we present make clear, these real income changes are not dramatic. The increases are modest and the densities hint at the fact inequality has increased. Our research and that of others confirm that the very high levels of inequality that apartheid bequeathed the incoming government in 1994 have increased even further. Also, rising unemployment makes it clear that the labor market has been a problem rather than part of the solution over the last fifteen years.

We conduct a series of semi-parametric decompositions in order to see if we can better understand the source of the shifts in the distribution of incomes. These decompositions look at the role of changes in endowments and changes in the returns to these endowments in driving the observed changes in the income distribution between 1993 and 2008. This analysis proves to be very useful in highlighting the positive role played by changes in endowments over the period. Indeed the resulting endowments counterfactual indicates that, if these endowment changes were all that changed in South Africa over the post-apartheid period, we would have seen a pervasive rightward shift of the distribution of per capita real incomes. This contrasts sharply with the actual shifts in the densities; which show clear improvements only at the bottom and the top of the densities. This is an important finding as it highlights the fact that the strong spending by the state on education and services, led to measurable improvements in levels of education and access to essential services but these improved endowments did not translate into generalized increases in real

incomes. Therefore, something dampened the translation between improved endowments and improved real incomes. Our semi-parametric analysis of returns indicates that, at the household level, this dampening was not due to a pernicious change in returns to endowments. *Ceteris paribus*, the change in returns makes a small positive contribution to the bottom and middle sections of the distribution. Unfortunately, the semi-parametric analysis is not able to assess the impact of changes in returns to each separate endowment. This is a pity as the evidence coming from the analysis of individual earnings in the labor market (e.g. Banerjee, Galiani, Levinsohn, McLaren and Woolard (2008)) is that there has been a skills twist in the returns to education in South Africa that has lowered the returns to education for all but the highest levels of schooling. This includes the incomplete secondary school years where the greatest gains have been made in post-apartheid South Africa.

From the advent of the post-apartheid period, South Africa has always had an extensive social welfare system based on a large state old age pension. This pension persisted through the post-apartheid years but has not been extended significantly. There has been one major extension to the welfare system; from 1998 onwards a child support grant was implemented with very high take up in the middle 2000s. In our semi-parametric framework this would change the returns to the endowment of the number of young children in the household. While we cannot isolate the impact of this change within the semi-parametric framework, we run a simple with CSG/without CSG simulation that shows just how important CSG income is to the lower part of the distribution of per-capita real incomes.

This is suggestive of the fact that it is the system of social grants in general and the new support coming from the child support grant in particular that counterbalances a strongly negative set of changes coming from the labor market. The strong social spending on social services, education and health have a potentially positive role to play. However,

our evidence suggests that they are yet to generate broad-based income returns. The net effect of all of these changes is a positive increase in real incomes over the post-apartheid period.

Figure 8 and table 5 taken together reiterate the point that this increase in real incomes is the net outcome of some strongly positive and some strongly negative forces. Figure 8 presents social expenditures over the post-apartheid period and extrapolates these expenditures into the next few years. It retells the remarkable story of the expansion of the social grants and also the large (by international standards) expenditures on education and health. As shown by the debt service figures, one of the accomplishments of South African government policy over the period has been that these expenditure expansions were accomplished while bringing down the daunting public debt that the apartheid state handed over to the new South African democracy.

It is exactly this combination of cash transfers and the expansion of education that is credited with the reduction of inequality in Brazil and Mexico since 2000. However, as we have shown, inequality has risen not fallen in South Africa. The key difference between the Latin American and the South African experiences seems to be that social grants and improved levels of education accompanied and contributed towards strong employment creation in Latin America whereas this employment creation has not happened in South Africa. Table 6 shows this quite vividly. It can be seen that even in 1993 high unemployment rates were the marker of those in the lowest deciles. By 2008 unemployment rates rise across all deciles and they rise particularly sharply in the bottom half of the distribution. Taken in isolation this table does not accord with a society generating positive, inclusive economic growth and social stability.

It is this balance that makes it hard to be unequivocally positive. The post-apartheid

state has clearly been pro-active. However, other than through the generation of rising tax revenues, this appears to have failed to generate virtuous interactions with the real economy. Indeed the global financial crisis has sharpened these dilemmas. It can be seen in figure 8 that the debt service is starting to rise again. This is a reflection of the fact that the growing budget deficits are being generated in order to finance the states expenditure programs.

While real spending on social grants has been protected, it has not continued to grow. To some extent this is due to the tighter financial conditions. However, this is also due to a growing recognition that these grants cannot be expanded indefinitely. Woolard and Leibbrandt (2010) review a large corpus showing that these unconditional transfers result in many virtuous behavioral effects. These studies certainly justify the state's program over the last fifteen years to expand these grants to where they are now; one of the largest programs in the world. However, with the state old age pension being larger than the median per capita income and with this pension and the child support grant making up the dominant share of income for those in the lowest deciles, there are also grounds for worrying about the further expansion of these grants. For one thing, the grants are specifically targeted at the elderly, the disabled and children and rely on a set of indirect behavioral responses to connect to the labour market. Policies that directly address the labour market have to be the first priority.

References

- Banerjee, Abhijit, Sebastian Galiani, Jim Levinsohn, Zoe McLaren, and Ingrid Woolard**, "Why has unemployment risen in the New South Africa?," *The Economics of Transition*, 2008, 16 (4), 715–740.
- Blinder, Alan**, "Wage Discrimination: Reduced Form and Structural Estimates," *Journal of Human Resources*, 1973, 8 (4), 436–455.

- Case, Anne and Angus Deaton**, “Large Cash Transfers to the Elderly in South Africa,” *Economic Journal*, 1998, *180*, 1330–1361.
- DiNardo, John, Nicole Fortin, and Thomas Lemieux**, “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semi-Parametric Approach,” *Econometrica*, 1996, *64* (5), 1001–1044.
- Duflo, Esther**, “Grandmothers and Granddaughters: Old Age Pensions and Intra-household allocation in South Africa,” *World Bank Economic Review*, 2003, *17* (1), 1–25.
- Hoogeveen, Johannes and Berk Ozler**, “Not Separate, Not Equal: Poverty and Inequality in Post-Apartheid South Africa,” February 2004. Unpublished Draft.
- Leibbrandt, M., I. Woolard, H. McEwen, and C. Koep**, “Better Employment to Reduce Inequality Further in South Africa,” 2010. Chapter 5 in OECD 2010. Tackling Inequalities in Brazil, India, China and South Africa: the Role of Labour Market and Social Policies.
- Leibbrandt, Murray, Ingrid Woolard, Arden Finn, and Jonathan Argent**, “Trends in South African income distribution and poverty since the fall of Apartheid,” 2010. OECD Social, Employment and Migration Working Papers No. 101.
- , **James Levinsohn, and Justin McCrary**, “Incomes in South Africa after the fall of Apartheid,” *Journal of Globalization and Development*, 2010, *1*, Article 2.
- Oaxaca, R.**, “Male-female wage differentials in urban labor markets,” *International Economic Review*, 1973, *14*, 693–709.
- Thomas, Duncan**, “Education Across the Generations in South Africa,” *American Economic Review*, 1996, *86* (2), 330–334.
- Treasury, National**, “2010 Budget Review,” 2010. Government Printer, Pretoria.
- Woolard, Ingrid and Murray Leibbrandt**, “The Evolution and Impact of Unconditional Cash Transfers in South Africa,” 2010. Southern Africa Labour and Development Research Unit Working Paper 51, University of Cape Town.

Table 1: Households in South Africa

Race	1993	2008
African	6,057,916	10,436,201
Coloured	658,717	1,146,969
Asian	228,238	334,613
White	1,551,149	1,717,498
Missing	0	87,637
Total	8,496,020	13722918

¹ Unit of observation is the self-reported household.

² Rates are calculated using sample weights.

Table 2: Mean Household Size

Race	1993	2008
African	5.30 (3.56)	3.68 (2.69)
Coloured	4.90 (2.29)	3.76 (2.16)
Indian	4.50 (1.80)	3.72 (2.22)
White	3.25 (1.57)	2.62 (1.23)
Total	4.87 (3.25)	3.55 (2.52)

¹ Standard Errors in Parentheses.

Table 3: Household Incomes in South Africa: 1993 and 2008

	1993					2008				
	All	African	Coloured	Asian	White	All	African	Coloured	Asian	White
Mean	10741	6018	7498	12555	29372	24409	9718	25269	27423	110195
Std. Dev.	17994	7903	8599	12423	31483	172733	23869	145855	43051	453800
Percentiles										
1st	0	0	0	0	0	62	87	62	80	306
5th	230	328	0	0	0	510	510	764	80	4000
10th	644	621	966	1191	1159	834	764	1276	1558	6803
25th	1581	1288	2439	4680	9606	1911	1663	2354	4179	14619
50th	4444	3098	5172	8607	21316	5096	3757	5809	11833	28801
75th	12558	7651	9254	17161	38844	14982	9690	11466	26754	55495
90th	27145	15799	16374	26713	62468	35672	23279	23063	57330	101920
95th	40773	21155	24633	39686	88582	59496	35222	40768	164346	190490
99th	86521	34241	41538	62210	154007	189826	76440	485035	207025	3967236
n	8,496,020	6,057,916	658,717	228,238	1,551,149	13,635,281	10,436,201	1,146,969	334,613	1,717,498

¹ Unit of observation is per capita household total income.

Table 4: Household Log Incomes in South Africa: 1993 and 2008

	1993						2008					
	All	African	Coloured	Asian	White	All	African	Coloured	Asian	White		
Mean	8.44	8.03	8.54	9.13	9.94	8.58	8.29	8.62	9.18	10.22		
Std. Dev.	1.41	1.26	1.02	1.00	1.09	1.54	1.35	1.39	1.81	1.48		
Percentiles												
1st	4.82	4.75	5.80	5.24	6.18	4.47	5.04	4.13	4.39	5.72		
5th	6.18	5.96	6.90	7.57	8.08	6.28	6.23	6.64	4.39	8.29		
10th	6.71	6.52	7.27	8.11	8.65	6.76	6.64	7.17	7.35	8.83		
25th	7.49	7.20	7.97	8.57	9.44	7.57	7.42	7.77	8.34	9.59		
50th	8.46	8.06	8.60	9.14	10.09	8.55	8.24	8.67	9.38	10.27		
75th	9.47	8.95	9.17	9.83	10.63	9.62	9.18	9.35	10.19	10.92		
90th	10.23	9.68	9.73	10.25	11.08	10.49	10.06	10.05	10.96	11.53		
95th	10.63	9.97	10.15	10.59	11.43	10.99	10.47	10.62	12.01	12.16		
99th	11.39	10.44	10.70	11.04	11.95	12.15	11.24	13.09	12.24	15.19		
n	7,323,317	5,228,275	592,242	193,320	1,209,480	12,468,985	9,406,025	1,071,028	318,660	1,606,743		

¹ Unit of observation is per capita household total income.

Table 5: Logit Regression for Reweighting

Variable	Coefficient	Standard Error	z
Household Size	-.211995	.039686	-5.34
African	1.473985	.060685	24.29
Coloured	1.155165	.083275	13.87
Asian	.9407829	.1223874	7.69
Number w/ Formal Jobs	-.6269353	.0277821	-22.57
Number of Adults	-.0474666	.0446675	-1.06
Highest Education in HH	.1236585	.0062593	19.76
Gender of HH Head	-.3668922	.042225	-8.69
SOAP eligible	-.1605647	.0366606	-4.38
Number of children under 14	.1361046	.0438766	3.10
Urban	.858294	.0509862	16.83
Metro	1.586205	.0511592	31.01
Constant	-.5383919	.1522273	-3.54

¹ Dependent Variable is a 1 if year is 2008, 0 if 1993.

² Whites are the excluded population group.

³ Highest Education is given in years.

⁴ Rural is the excluded region-type category.

Table 6: Unemployment Rates by per capita Income Deciles

Decile	1993	2008
1	49.1%	69.4%
2	33.6%	46.0%
3	26.8%	46.7%
4	22.0%	36.9%
5	23.4%	30.3%
6	18.7%	26.1%
7	14.5%	20.1%
8	9.4%	16.4%
9	4.3%	9.0%
10	1.5%	4.5%
Overall	13.7%	24.4%

¹ Source: Leibbrandt, Woolard, McEwen and Koep (2010a)

Figure 1: Household Income CDF's




Figure 2: Household Income CDF's- African Only

Figure 3: Log Per Capita Household Incomes




Figure 4: Log Per Capita Household Incomes-- African Only




Figure 5: The “Change in Endowments” Explanation




Figure 6: The “Change in Returns” Explanation




Figure 7: The “Child Support Grant” Explanation




Figure 8: Social Expenditures as a percentage of GDP