IARIW 2025

IARIW 2025

Thursday, October 2 & Friday, October 3

Persistent transitions and chronic poverty in Nigeria: Accounting for the effects of classification error

Rulof Burger
(Stellenbosch University, South Africa)

Zander Prinsloo
(World bank)

Dean Jolliffe
(World bank)

rulof@sun.ac.za.

Paper prepared for the IARIW–World Bank–UEB/VNU Conference on "Improving Well-being Measurement in Data-challenged Environments in Developing Countries for Better Evidence-based Policies" October 2-3, 2025

Session 2A: Measurement Errors
Time Slot: Thursday, October 2, 1:15- 3:15 PM

Persistent transitions and chronic poverty in Nigeria: Accounting for the effects of classification error *

Rulof Burger† Zander Prinsloo ‡ Dean Jolliffe§

September, 2025

Abstract

Nigerian households experience seemingly high levels of consumption mobility and frequent poverty transitions, suggesting that poverty in Nigeria is primarily transitory. However, this may be driven by measurement error in self-reported household consumption across successive panel data waves. Serially uncorrelated measurement error in the consumption aggregate can overestimate consumption in one wave and underestimate it in the next. If true consumption is in the neighborhood of the poverty line, measurement error can cause a transient misclassification of poverty status and an upward bias on estimated poverty transitions. Even when there are no changes in consumption and poverty over time, measurement and classification error will create the appearance of change. In Nigeria, poverty transitions are high across two waves

*The authors gratefully acknowledge financial support from the UK Government through the Data and Evidence for Tackling Extreme Poverty (DEEP) Research Program. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. All errors and omissions are our own. Declarations of interest: none.

†Department of Economics, Stellenbosch University and JPAL. Email: rulof@sun.ac.za.
‡World Bank, Washington, D.C, United States
§World Bank, Washington, D.C, United States

of panel data, but comparatively low over three waves. This pattern of poverty is consistent with the classification error hypothesis.

We employ two structural estimators that use three waves of panel data to estimate the rate of persistent poverty transitions while allowing for measurement and classification error. Our results suggest that transitions out of poverty by a factor of 3, and that neglecting classification error makes approximately 40% of chronic poverty appear transitory. As robustness tests, we show that allowing for duration-dependence does not explain the transient poverty pattern and that households with unreliable responses have more frequent poverty transitions. Both findings provide further evidence that classification error upwardly biases poverty transitions and makes chronic poverty appear transitory.

Keywords: Poverty transitions, consumption mobility, measurement error

JEL Codes: D86, D91, J22

1 Introduction

Progress towards alleviating extreme poverty, the first of the Sustainable Development Goals, has stagnated in recent years (World Bank, 2022). To reignite progress towards this goal, policies need to be well targeted and address the root causes of poverty. A comprehensive understanding of the causes and consequences of poverty requires distinguishing between transitory and chronic poverty (Jalan and Ravallion, 1998). Distinguishing between chronic and transitory poverty requires estimating poverty dynamics: the movement into and out of poverty (Green and Hulme, 2005).

Identifying chronic poverty requires reliable estimation of poverty transitions. There are long-standing concerns that measurement error in the underlying welfare aggregate overestimates poverty transitions and makes chronic poverty appear transitory (Rendtel et al., 1998). This poses a challenge to the reliability of poverty assessments and the ability to create targeted policies.

Measurement errors are ubiquitous in survey data. Although these errors do occur in rich countries (Moore et al., 2000), they are particularly pervasive in low-income countries (Antman and McKenzie, 2007; Ashenfelter et al., 1986) where poverty assessments are most critical. Observed consumption aggregates are therefore noisy measures of true consumption (Ravallion, 2003; Rendtel et al., 1998). Self-reporting leads to complex and potentially large errors in welfare aggregates (Wollburg et al., 2021; Friedman et al., 2017; Battistin, 2003; Lanjouw and Lanjouw, 1997) that have both systematic and random components (Moore et al., 2000).

When continuous measurement error around the welfare aggregate causes the observed aggregate to cross the poverty line, it results in classification error in the binary poverty status. Independent classification errors, commonly assumed in the literature on error-methods (Shibata, 2022; Breen and Moisio, 2004), and low misclassification rates imply transient misclassification of poverty status. In a hypothetical three period economy where poverty is strictly chronic over all three periods, any misclassifications will be observed as transitions out of poverty and mistaken for transitory poverty. A misclassification in the opposite direction – i.e. of a non-poor person who is misclassified as poor – will not cancel out the first error as it would when estimating levels. In fact, rather than having zero transitions there will now be four: one into and one out of poverty for each of the misclassified observations. This illustrates the multiplicative bias from measurement and classification error when estimating dynamics (Magnac and Visser, 1999; Singh and Rao, 1995).

Mismeasurement causes attenuation bias in the slope parameter (Schennach, 2016). This bias becomes more complex and problematic in nonlinear settings (Hausman et al., 1995). When estimating transitions, the attenuation bias from measurement error places a downward bias on persistence, upward bias on transition estimates, and exaggerate the extent of mean reversion (Ward, 2023; Shibata, 2022; Burger, 2016; Schennach, 2016). Even infrequent classification errors can cause substantial bias in transition rates (Magnac and Visser, 1999; Singh and Rao, 1995). For example, Breen and Moisio (2004) find that poverty transitions are overestimated by up to 50% due to classification error. The implication is that observed transitions estimated over successive, high-frequency panel data waves will be very high due to classification error, but relatively lower over multiple or lower-frequency waves as transient misclassification is corrected.

The upward bias on poverty transitions and downward bias on poverty persistence due to classification error makes chronic poverty appear transitory. Chronic poverty, using the spells approach, is identified by the frequency and duration of poverty spells, and is characterised by infrequent poverty transitions (Foster, 2009). These identifying characteristics are precisely those that are biased by mismeasurement. Attaining reliable estimates of poverty dynamics, the movement into and out of poverty, is crucial for diagnosing the extent of chronic poverty.

Traditionally, limited access to panel data meant a neglect of the temporal dimension to poverty, ignoring the duration of deprivation and severity of welfare fluctuations (Baulch and Hoddinott, 2000). However, persistent, chronic poverty is detrimental to a range of important outcomes in ways that transitory poverty is not, including nutrition for both adults (McDonough and Berglund, 2003) and children, as well as children's learning outcomes and other health outcomes (Brooks-Gunn and Duncan, 1997), and the probability of remaining in poverty (Jo Bane and Ellwood, 1983). The efforts to estimate poverty transitions efforts have greatly benefited by the increased availability of multi-wave household panel data (Green and Hulme, 2005) and have aided the design of social safety nets and programs to support the vulnerable (Baulch and Hoddinott, 2000).

Chronic poverty, characterized by its duration and persistence, is evidence of deep social and structural causes. It is rooted in enduring structural issues, while transitory poverty arises from temporary setbacks (Jalan and Ravallion, 1998). These structural causes are distinct from the forces leading households to temporarily slide into poverty due to a transient shock. These two types of poverty are distinct phenomena that require uniquely tailored interventions, and attempts to eradicate extreme poverty can greatly benefits from an understanding of these differences (Aliber, 2003). Transitory poverty, while damaging to the household, resolves itself – thus, policy should be primarily concerned with persistent chronic poverty (Calvo and Dercon, 2009). Policies with either a short-term view that provide temporary support to climb above the poverty line or those that provide protections to adverse shocks will not address chronic poverty. Here, longer-term resource-intensive solutions are required and includes both financial assistance programs and the building of assets – human capital, financial, and physical – for chronically poor households (Hulme and Shepherd, 2003).

Distinguishing between chronic and transitory poverty is crucial for policy design. However, correcting for the confounding effects of measurement and classification error is challenging (Duclos et al., 2010). Many studies have turned to auxiliary data sources – reinterviews, administrative data, or national accounts – to correct measurement errors in survey data (Abowd and Zellner, 1985; Poterba and Summers, 1986; Chua and Fuller, 1987). However, these data sources have their own errors (Abowd and Stinson, 2013), and integrating data from different sources introduces additional issues, like matching errors, inconsistent interview methods, and variation in timing of interviews. These correction methods are not always feasible in low-income countries, where reinterviews can be prohibitively expensive and administrative data are scarce. These challenges necessitates the development of structural econometric approaches to dealing with measurement and classification errors (Poterba and Summers, 1986; Dinkelman, 2004; Burger, 2016).

There are a number of attempts to model mobility under measurement error in the literature. These typically rely on the Markov assumption on a continuous outcome, as in Ward (2023). However, continuous mobility regressions do not model poverty transitions directly and therefore are unable to provide direct insight into poverty dynamics and chronic poverty. There is a collection of discrete approaches based on Biemer and Bushery (2000) that use Markov latent class analysis (MLCA). Breen and Moisio (2004) applies this approach to poverty transitions, while others like Shibata (2022) and Feng and Hu (2013) apply a variation of the approach to employment transitions. These error-correction approaches are do not allow extensions to the functional forms to allow for duration-dependence or heterogeneity, for example. This potentially exaggerates the estimated effect and extent of classification errors by attributing the full violation of the Markov property to classification error.

In this study, we add to the literature in three ways. First, we present a structural classification error estimator that directly models poverty transitions in the presence of classification error. This estimator is extended to provide additional support for the identifying assumptions. The first extension is to allow for higher-order dependence by modeling poverty transitions as an AR(2) process. The second extension is into multiple equation nonlinear systems estimator, in which the probability of misclassifying poverty status allowed to depend on other inconsistent survey responses. In both cases, support is found for the hypothesis that there is significant misclassification of poverty and that this error under represents the extent of chronic poverty.

The second contribution is to use an alternative specification – modeling the dynamics of the continuous household consumption measure and using the estimates to simulate poverty dynamics for both true and mismeasured consumption aggregates. The strength of the continuous approach is that it utilizes information from the full welfare distribution. Therefore the version of the Markov assumption made in the continuous case is less heroic. Results from both approaches consistently reveal a significant overestimation of poverty transitions when conventional methods that disregard measurement or classification error are applied.

The third contribution is in the analysis of Nigerian poverty. Due to a historic lack of panel data, relatively little is known about Nigerian poverty dynamics. Dapel (2018) uses synthetic panels constructed from six cross-sectional survey data sets, and finds that poverty worsened substantially in Nigeria between 1980 and 2010, from 27% to 69%. This increase was mainly driven by an increase in transitory poverty. Idika et al. (2021) use two waves of the recently collected Nigerian General Household Survey panel data to study changes in poverty rates across different regions. They find that poverty is greater in rural areas, and lower in Southern, oil-producing regions. However, they do not estimate household-level poverty transition rates. To our knowledge, this is the first study investigating Nigerian poverty dynamics by utilizing three waves of data from the Nigerian General Household Survey. We find that approximately 20% of the population are in chronic extreme poverty – ignoring classification error would make nearly half of the chronic poverty appear transitory.

The remainder of this paper are organized as follows. Section 2 delves into the data employed in our empirical analysis, providing simple descriptive statistics on consumption and poverty trends in Nigeria. Section 3 outlines the two estimators used to gauge poverty transitions and consumption mobility in the presence of measurement and classification errors. The findings of this analysis are discussed in section 4, while section 5 presents robustness tests. It includes additional evidence supporting the notion that the elevated rate of poverty transitions is attributable to measurement errors in consumption rather than transient consumption shocks. Finally, section 6 provides a conclusion to our study.

2 Data and descriptive statistics

The Nigerian General Household Survey (GHS) is a multi-wave survey that aims to measure household welfare, with a particular focus on the agricultural sector. It is the result of a collaboration between Nigeria's National Bureau of Statistics (NBS), the Federal Ministry of Agriculture and Rural Development, the National Food Reserve Agency (NFRA), the Bill and Melinda Gates Foundation, and the World Bank. The survey forms part of the World Bank's Living Standards Measurement Study - Integrated Surveys on Agriculture (LSMS-ISA) project. Of the 22 000 households in the original cross-sectional survey (2010-2011), 5 000 form part of a panel component and are included in the second (2012-2013) and third (2015-2016) waves. We were able to construct a nationally representative three-wave balanced panel containing 4 217 of these households.

Each wave involves two visits, one post-planting and one post-harvest. The post-planting visit cycle ranges from August to October (September to November for wave 2), and the post-harvest cycle is from February to April of the next year. The GHS panel survey adopts a two-stage stratified sample selection, where enumeration areas (EAs) are primary sampling units, and 10 households are sampled within each EA using a sampling interval approach. The survey comprises twelve strata, representing the rural and urban sections of Nigeria's six geopolitical zones.

The survey has a detailed data cleaning process, including quality control during fieldwork, outlier and missing value checks at NBS offices, and matching checks to ensure consistency across survey sections. The data collection phase lasts 20-30 days and is followed by a tracking phase where attempts are made to interview relocated households in order to mitigate bias due to attrition. After each day of conducting interviews, a error report is generated and used to revisit households with clearly problematic responses.

Three different questionnaires – household, agriculture, and community/prices – are administered in each wave, enabling the construction of household consumption aggregates, calculated by the NBS as a nominal annual per capita household consumption variable. Consumption is converted to daily per capita consumption in real 2017 Naira values and further adjusted using a PPP adjustment factor to obtain daily per capita consumption in 2017 USD terms. For the analysis of poverty dynamics, we employ the $2.15 international poverty line set by the World Bank (Jolliffe et al., 2022). Table 2 illustrates that slightly over 30% of the population falls below the $2.15 extreme international poverty line in each wave.

Table 1 presents key descriptive statistics for four consumption measures across the three panel data waves. These measures include nominal annual household consumption, real annual household consumption, real daily consumption per capita adjusted for purchasing power parity (PPP), and the logarithm of real daily per capita consumption adjusted for PPP. The table includes mean, median, and standard deviation values, along with sample sizes for each wave. Across all three years, the mean values for consumption measures exceed the medians, suggesting positively skewed distributions, as anticipated. Conversely, the log consumption variable appears to exhibit a more or less symmetrical distribution. Notably, there is an increase in real consumption between waves 1 and 2, and again between waves 2 and 3. The standard deviation in log consumption also rises over time, indicating that consumption growth coincides with increased inequality.

Table 1: Consumption levels for balanced panel, by wave

Wave 1 (2011)				
	Mean	Median	Std. Dev.	N
Nominal annual household per capita consumption	101,153	79,758	80,063	4,217
Real annual household per capita consumption	176,630	139,270	139,804	4,217
Real daily household per capita consumption (PPP)	3.91	3.09	3.10	4,217
Log real daily per capita consumption (PPP)	1.13	1.13	0.68	4,217
Wave 2 (2013)				
	Mean	Median	Std. Dev.	N
Nominal annual household per capita consumption	126,438	88,610	219,838	4,217
Real annual household per capita consumption	189,490	132,798	329,466	4,217
Real daily household per capita consumption (PPP)	4.20	2.94	7.30	4,217
Log real daily per capita consumption (PPP)	1.13	1.08	0.69	4,217
Wave 3 (2015)				
	Mean	Median	Std. Dev.	N
Nominal annual household per capita consumption	169,332	120,218	182,912	4,217
Real annual household per capita consumption	195,322	138,670	210,986	4,217
Real daily household per capita consumption (PPP)	4.33	3.07	4.67	4,217
Log real daily per capita consumption (PPP)	1.17	1.12	0.73	4,217

Table 2 provides poverty transition rates for different periods. Using a $2.15 poverty line, the poverty rate remains relatively stable at around 30%. Transitions out of poverty are frequent, with nearly 40% of poor households escaping poverty between waves 1 and 2. This rate decreases to 34% between waves 2 and 3. Over the entire three-wave period, 39.4% of households exit poverty, a rate nearly identical to that between waves 1 and 2. Transitions into poverty are less common but still notable. The probability of entering poverty is 17% between waves 1 and 2 and 15% between waves 2 and 3. Over the entire three-wave period, nearly 18% of non-poor households in wave 1 transition into poverty by wave 3.

Table 2: Poverty transitions: \$2.15 poverty line

Wave 1	Wave 2		
	Non-poor	Poor	Total
Non-poor	57.51	11.42	68.94
Poor	12.12	18.94	31.06
Total	69.63	30.37	100.00
Probability of exiting poverty: 39%			
Probability of entering poverty: 16.6%			
Wave 2	**Wave 3**		
	Non-poor	Poor	Total
Non-poor	58.88	10.76	69.63
Poor	10.25	20.12	30.37
Total	69.12	30.88	100.00
Probability of exiting poverty: 33.7%			
Probability of entering poverty: 15.4%			
Wave 1	**Wave 3**		
	Non-poor	Poor	Total
Non-poor	56.89	12.05	68.94
Poor	12.24	18.83	31.06
Total	69.12	30.88	100.00
Probability of exiting poverty: 39.4%			
Probability of entering poverty: 17.5%			

The descriptive statistics in this section suggests that Nigerian society is characterised by very frequent transitions into and out of poverty. However, the three-wave transition rates are only marginally higher than what is observed over two waves, suggesting that such transitions are not persistent. One explanation for such short-lived transitions is measurement error in the consumption measure. The next section will explain how we use structural econometric methods to distinguish between persistent transitions into and out of poverty, and transient consumption shocks due to measurement error.

3 Methodology

3.1 Discrete poverty transitions

We start our analysis by estimating the extent of classification error in the binary poverty measure. Suppose a binary but unobservable true poverty status variable, s_t^*, takes on a value of 1 if a household is poor in period t and 0 if not, and is generated by the following process:

$$s_t^* \equiv \gamma_1 s_{t-1}^* + \gamma_2 (1 - s_{t-1}^*) + v_t^* \tag{1}$$

where γ_1 and γ_2 are constant unobservable parameters[1]. The unobservable random error term, v_t^*, has a mean-zero Bernoulli distribution when conditioned on true poverty with probabilities corresponding to the poverty entry rate, γ_2, and the poverty exit rate, $1 - \gamma_1$, for the case where s_t^* is equal to 0 or 1 respectively. Equation 7 makes the first-order Markov assumption on true poverty: future true poverty given current poverty is independent of past poverty. The Markov assumption is commonly made in the classification error literature (Shibata, 2022; Feng and Hu, 2013; Breen and Moisio, 2004; Biemer and Bushery, 2000). Furthermore, suppose s_t, is an observed but potentially misclassified poverty measure, which is related to true poverty as

[1]Since these parameters are restricted to fall within the unit interval, they can be represented as the output of the following probit transformations: $\gamma_i = \Phi(\gamma_i^*)$.

$$s_t \equiv \pi s_t^* + (1-\pi)(1-s_t^*) + v_t \tag{2}$$

where π is a time-invariant, unobservable misclassification parameter[2] and v_t is an unobservable random error term that follows a conditional Bernoulli distribution with $E(v_t|s_t^*) = 0$. Assuming there is no classification error implies that $\pi = 1$, which in turn implies that the observed poverty process exhibits the Markov property. The literature utilizing transition matrices often implicitly make the Markov assumption on the mismeasured process, which is more restrictive than our hypothesized process. We allow the observable mismeasured process to deviate from a Markov process due to classification error. If transitions out of poverty tend to be reversed in the following period then this violates our identifying Markov assumption. It also implies that such transitions are largely transitory rather than persistent. Since we estimate the poverty transition rates that are consistent with the Markov assumption and identify fleeting transitions as classification error, our poverty transitions rate estimates can be interpreted as the rate at which households *persistently* enter or exit chronic poverty.

Equation 2 assumes that classification error is independent of true poverty status and time as in Wolfe et al. (2003); Burger (2016); Shibata (2022) – we also relax this assumption by allowing for asymmetric misclassification, where the classification error rates depend on a person's true poverty status. Moreover, the noise process in equation 2 that generates the observable but potentially mismeasured poverty statuses assumes that observed poverty in two consecutive periods is independent conditional on true poverty as in Abowd and Zellner (1985); Chua and Fuller (1987); Poterba and Summers (1986); Singh and Rao (1995); Magnac and Visser (1999); Breen and Moisio (2004). The system can be expressed in terms of the

[2]The π parameter is also restricted to fall within the unit interval, and can thus be represented as the output of the following probit transformation: $\pi = \Phi(\pi^*)$

observed poverty statuses over three waves as below:

$$
\begin{aligned}
P(s_t = 1|s_{t-1}, s_{t-2}) = (1-\pi) &+ (2\pi - 1)(1 + \gamma_1 - \gamma_2)\gamma_2 - (\gamma_1 - \gamma_2)^2(1 - \pi) \\
&+ (\gamma_1 - \gamma_2)^2 s_{t-2} \\
&+ E(v_t|s_{t-1}, s_{t-2}) - (\gamma_1 - \gamma_2)^2 E(v_{t-2}|s_{t-1}, s_{t-2}) \\
&+ (2\pi - 1)E(v_t^*|s_{t-1}, s_{t-2}) + (2\pi - 1)(\gamma_1 - \gamma_2)E(v_{t-1}^*|s_{t-1}, s_{t-2})
\end{aligned} \tag{3}
$$

Deriving the population likelihood requires finding the observed joint probability, $P(s_{t-1}, s_{t-2})$, and the conditional expectations of the stochastic terms in equation 3 above. We estimate this model using the expectation-maximization (EM) algorithm pioneered by Dempster et al. (1977), commonly used for partially unobserved variables and mixture models. This is an iterative two-step process, first calculating the expectation of the log-likelihood over the unobserved s^* and conditional on current parameter values, then maximizing this expectation. A similar estimation is used by Shibata (2022) and Breen and Moisio (2004).

3.2 Continuous consumption mobility

The discrete approach above loses useful information about how far households are from the poverty line, but is agnostic about the underlying consumption distribution. It therefore sacrifices efficiency for robustness.

Now we present the continuous consumption measurement error approach, before using simulations to gauge the effect of these errors on observed poverty transition rates. The continuous consumption approach leverages information from the entire consumption distribution, which is advantageous from an estimator efficiency perspective. However, deriving the implications of consumption mobility for poverty transitions requires making additional distributional assumptions.

Log consumption per capita in period t, y_t^*, is assumed to follow an autoregressive process of order 1:

$$y_t^* - y_{t-1}^* = \mu_t + \beta y_{t-1}^* + u_t \tag{4}$$

where β measures the speed of unconditional beta convergence between periods $t-1$ and t, and it is assumed that $u_t \sim iid(0, \sigma_u^2)$. Importantly, this assumption implies that consumption shocks are persistent and without memory so there is no inherent inclination for such shocks to be reversed in the subsequent periods. Our interpretation of β is therefore a measure of persistent convergence. If $\beta = 0$ then there is no tendency for high and low consumption households to experience persistently different growth rates, whereas if $\beta < 0$ then we should expect consumption levels to converge[3]: initially poor households tend to experience more rapid consumption growth than rich ones.

Observed measures of consumption may be measured with error. Suppose observed consumption per adult equivalent in period t is denoted y_t and differs from the true but unobservable consumption value by $e_t \equiv y_t - y_t^*$. In addition, we assume that $e_t \sim iid(0, \sigma_e^2)$. Rewriting equation 4 in terms of the observed consumption measure produces

$$y_t - y_{t-1} = \mu_t + \beta y_{t-1} + u_t + e_t - (\beta + 1)e_{t-1} \tag{5}$$

Since y_{t-1} is positively correlated with e_{t-1} by construction, regressing $y_t - y_{t-1}$ on y_{t-1} is expected to produce an estimate of β that is downwardly biased, creating the appearance of more mobility and less persistence in consumption than is actually the case. This makes it difficult to distinguish between measurement error and exaggerated mobility in two waves of panel data. However, if fluctuations in measured consumption between waves 1 and 2 are caused by transient errors in measurement as opposed to persistent shifts in consumption,

[3]This kind of beta-convergence does not necessarily imply a reduction in inequality (sigma-convergence) over time: consumption shocks replenish the tails of the consumption distribution to produce a stable distribution.

then we would not expect such variation to have any impact on wave 3 consumption. This insight can be leveraged along with a third wave of panel data to test how much of two-wave mobility is due to actual consumption variation rather than measurement error.

There are at least seven regression coefficients that can be estimated on three waves of panel data to provide information on the extent of mobility and measurement error. Table 3 defines each of these coefficients as the outcome of a linear projection, indicating the ordinary least squares (OLS) regression that produces the coefficient estimate. It also notes the expected value of each coefficient in terms of the speed of convergence, β (in the case of no measurement error) and in terms of β and a data reliability statistic $\alpha \equiv \frac{var(y)}{var(y^*)}$ (in the case of measurement error).

Table 3: Regression coefficients and population moments

Regression coefficients		Population moments	
Coefficient	Linear projection	No measurement error	Measurement error
θ_1	$L\left(y_2 - y_1 \vert y_1\right) = \theta_1 y_1$	β	$\alpha(\beta+1)-1$
θ_2	$L\left(y_3 - y_2 \vert y_2\right) = \theta_2 y_2$	β	$\alpha(\beta+1)-1$
θ_3	$L\left(y_3 - y_1 \vert y_1\right) = \theta_3 y_1$	$(\beta+1)(\beta+1)-1$	$\alpha(\beta+1)(\beta+1)-1$
θ_4	$L\left(y_3 - y_2 \vert y_1\right) = \theta_4 y_1$	$(\beta+1)\beta-1$	$\alpha(\beta+1)\beta$
θ_5	$L\left(y_3 - y_2 \vert y_1, y_2\right) = \theta_5 y_1 + \theta_6 y_2$	0	$\frac{(\beta+1)(\beta+1)\alpha\alpha\left(\frac{1}{\alpha}-1\right)}{\frac{\alpha}{\alpha}-(\beta+1)(\beta+1)\alpha\alpha}$
θ_6	$L\left(y_3 - y_2 \vert y_1, y_2\right) = \theta_5 y_1 + \theta_6 y_2$	β	$\frac{(\beta+1)\alpha(1-(\beta+1)(\beta+1)\alpha)}{\frac{\alpha}{\alpha}-(\beta+1)(\beta+1)\alpha\alpha} - 1$
θ_7	$L\left(y_3 - y_2 \vert y_2 - y_1\right) = \theta_7 (y_2 - y_1)$	$\frac{\beta}{2}$	$\frac{1-\frac{1}{\alpha}-\beta\beta}{\frac{1}{\alpha}+\frac{1}{\alpha}-2(\beta+1)}$

The derivations in table 3 rely crucially on three assumptions: true consumption follows an AR(1) process, consumption shocks are i.i.d., and measurement error is classical. If consumption shocks have a strong transient component, or actually follow an AR(2) process – e.g. because large positive consumption shocks are commonly reversed in the following period – then this would violate our identifying assumptions. However, in either of these cases, the two-wave mobility estimates or poverty transition matrices would be an unreliable indication of the extent of *persistent* consumption mobility. Even though our estimator would then conflate transient consumption shocks and measurement error, our estimates of true mobility would still be a more reliable indication of persistent mobility and transitions.

There are two ways to use the information from these regressions to investigate the effect of measurement error on consumption mobility estimates. First, we can compare the different estimates to test their internal consistency in the absence of measurement error. For example, coefficients θ_1 and θ_2 represent the conventional two-wave estimates of consumption mobility (or the speed of beta convergence). If these coefficients accurately reflect the extent of true and persistent mobility in economic welfare, then we would expect the total mobility between waves 1 and 3 (captured by θ_3) to be substantially higher than between either two-wave period[4]. On the other hand, if all consumption variability over time is driven by measurement error, then estimated mobility between waves 1 and 3 should be no higher than mobility between waves 1 and 2 or waves 2 and 3[5].

Secondly, we could use the expected values of the regression coefficients in the presence of measurement error to estimate the most likely values of β and α parameters using a system estimator like the generalised method of moments (GMM). This would produce more precise point estimates of these parameters and facilitate hypothesis testing. One hypothesis of particular relevance is the J-test of over-identifying restrictions, which tests whether the estimated value of all seven regression coefficients can be explained with the parsimonious two-parameter model we propose.

Once the parameters of the consumption dynamics model are estimated and validated we can explore the implications of this model for poverty transitions via simulations. Specifically, we can use the estimated parameter values to simulate three waves of true (unobserved) and mismeasured (observed) household consumption data and see how the frequency of poverty

[4]More precisely, Table 3 indicates that mobility between waves 1 and 3 should be the product of mobility between waves 1 and 2 and between waves 2 and 3.

[5]The same argument can be framed in terms of transient consumption shocks: in consumption mobility between waves 1 and 3 is no higher than between waves 1 and 2, and 2 and 3, then consumption shocks must be completely transient.

transitions differ between these two simulated measures. These simulations require making additional distributional assumptions about consumption shocks and measurement error that are not required for the estimator. We can compare the poverty transition matrices for the simulated mismeasured consumption measure to the same transition matrices for the observed real-world data to check the validity of these assumptions.

4 Empirical analysis

4.1 Estimating discrete poverty transitions

The binary approach, outlined in section 3.1 above, models the discrete poverty transitions directly without requiring estimating consumption dynamics or distributional assumptions. Table 4 shows the implied probabilities, which are transformations of the estimated parameters. All three columns in table 4 use the $2.15 international poverty line. The first column restricts the model so that there is no classification error, implying that all observed poverty transitions are true transitions. When imposing this restriction, we estimate that 16% of households become poor in successive waves, while the over 36% escape poverty. The estimated poverty rate is just over 31%.

Column 2 relaxes this restriction and allows for a misclassification probability that is symmetric: the probability of a poor household being misclassified as non-poor is the same as that of a non-poor household being misclassified as poor. Here we estimate that 9.5% of households are misclassified, and that the true poverty entry and exit rates decrease by approximately a factor of three to 5.3% and 13.2% respectively. Because transient classification error will bias transition estimates towards 50%, removing the bias due to classification error decreases the poverty rate by over 4 percentage points to 26.8%.

Column 3 of table 4 allows the misclassification rate to differ based on a household's true poverty status. When doing so, we estimate that poor households have a 17% probability of being classified as non-poor, while the probability of misclassifying non-poor households as poor is only around 8.2%. This additional flexibility in classification error rates reduces the estimated true poverty entry and exit rates even further to only 3.8% and 3.5% respectively. At a poverty rate of 30%, the asymmetry in misclassification rates means that approximately the same number of poor and non-poor households are misclassified and that the bias in the estimated poverty rate is not very large. The goodness-of-fit, as seen in the log-likelihood values, shows a large improvement in fit when allowing for classification error in the model and a relatively smaller improvement when allowing for asymmetric misclassification.

Overall, the results from our binary poverty model are broadly consistent with those from the continuous consumption model and suggest that poverty transitions are overestimated by approximately a factor of three due to a misclassification rate of near 10%.

Table 4: Discrete regression \$2.15 poverty line —implied probabilities

	(1)	(2)	(3)
	No misclass.	Misclass.	Asymm. misclass.
Miscl. poor (%)	0	9.52***	17.21***
	(.)	(0.007)	(0.012)
Miscl. not poor (%)	0	9.52***	8.19***
	(.)	(0.007)	(0.008)
Entry rate (%)	16.00***	5.31***	3.80***
	(0.006)	(0.009)	(0.010)
Exit rate (%)	36.41***	13.19***	3.53***
	(0.011)	(0.023)	(0.027)
Poverty rate (%)	31.07***	26.79***	30.07***
	(0.009)	(0.011)	(0.014)
Log lik (mill.)	-48.06	-47.41	-47.22
Misclassification	No	Symmetric	Asymmetric

Note: '***' as 0.001, '**' as 0.01, '*' as 0.05

The results in table 4 illustrate the upward bias on estimated transition rates due to classification error. The most common way to display these transition rates is using transition matrices. Table 5 reports three such transition matrices based on the estimates from table 4. The first transition matrix, at the top of Table 5 uses the estimates from the model with no misclassification (column 1 in table 4). The second, in the middle of Table 5 uses the symmetric misclassification model estimates (column 2 in table 4) and reports the transition matrix for the observed but misclassified poverty variable. These two transition matrices are very similar to each other and to the transition matrix for the Nigerian data in table 2 in terms of poverty levels and transition rates. The third matrix at the bottom of table 5 also uses the symmetric misclassification model estimates (column 2 in table 4), but now calculations the transitions for the true but unobserved poverty variable.

The results of table 5 make two striking points. The first builds on the previous regression results and shows that chronic poverty is far more pervasive in Nigeria than naive transition estimates imply. Measurement error in the consumption aggregate exaggerates short term consumption variation and leads to faux transitions for households in the neighbourhood of the poverty line. Misclassifications understate the extent of persistent poverty in Nigeria across two waves: without allowing for classification error, we would mistakenly conclude that only 58% of the population are consistently poor over a two year period when this number is actually just below 70%. The approximately 10% classification error rate leads to poverty transitions in Nigeria being exaggerated by a factor of three.

The second point is a methodological one. Transition matrices, a ubiquitous approach to estimating and communicating transition rates, are highly sensitive to classification error. As discussed in section 1, it is known that classification error places an upward bias on transition estimates. Yet, they are still frequently used in the poverty literature. The results in table 5 demonstrate that one should be cautious of drawing conclusions about the extent of chronic poverty without controlling for classification error, especially because of the error-prone nature of survey data.

Table 5: Transition matrices from discrete approaches at \$2.15 poverty line

True transitions, no misclassification estimator			
	Non-poor	Poor	Total
Non-poor	57.90	11.03	68.93
Poor	11.31	19.76	31.07
Total	69.21	30.79	100.00
Probability of exiting poverty: 36.4%			
Probability of entering poverty: 16.0%			
Mismeasured transitions, misclassification estimator			
	Non-poor	Poor	Total
Non-poor	57.60	11.19	68.79
Poor	10.90	20.31	31.21
Total	68.50	31.50	100.00
Probability of exiting poverty: 34.9%			
Probability of entering poverty: 16.3%			
True transitions, misclassification estimator			
	Non-poor	Poor	Total
Non-poor	69.32	3.89	73.21
Poor	3.53	23.26	26.79
Total	72.85	27.15	100.00
Probability of exiting poverty: 13.2%			
Probability of entering poverty: 5.3%			

Table 6 extends this analysis by showing the implied proportion of the Nigerian population that follows each of eight potential poverty histories over the three waves, where ”W1” refers to wave 1 and ”P” and ”N” refer to poor and non-poor respectively. Using the spells approach to identifying chronic poverty, we regard those that are poor across all three waves as chronically poor. The fourth column (labelled "Data") shows the weighted proportion if households that follow each of the relevant poverty histories according to the Nigerian survey data. The fifth column (labelled "True, no misclassification") reports the proportions implied by the no misclassification estimator (estimated in column 1 of table 4). Although table 5 showed that the no misclassification estimator is able to accurately predict the two-wave poverty transition matrix of the Nigerian data, this is not the case for the more detailed three-wave poverty histories. Combining the Markov property with the assumption of no misclassification produces predicted poverty histories in which chronic poverty (PPP) and chronic non-poverty (NNN) occur less frequently than is actually the case. It also under-estimates the frequency of poverty histories in which wave 1 to wave 2 transitions are reversed in wave 3 (PNP and NPN), while over-estimating the occurrence of histories in which either poverty or non-poverty is somewhat persistent (PPN, PNN, NPP, and NNP). The failure of this model to accurately predict poverty history frequencies is due to the fact that it does not fit all of these moments. Rather, it imposes the first-order Markov assumption, which breaks down due to the surprising number of households that churn in and out of poverty in successive waves. The Markov property cannot be reconciled with such a high churn rate over multiple waves unless we allow for classification error.

W1	W2	W3	Data	True no misclassification	Mismeasured misclassification	True misclassification
P	P	P	14.04	12.56	15.59	20.19
P	P	N	4.90	7.19	4.72	3.07
P	N	P	4.79	1.81	2.85	0.19
P	N	N	7.33	9.50	8.05	3.35
N	P	P	6.08	7.02	4.97	3.38
N	P	N	5.34	4.02	6.22	0.51
N	N	P	5.97	9.27	8.32	3.68
N	N	N	51.54	48.64	49.28	65.64
Absolute Deviation				17.38	10.99	

Table 6: Proportions that are poor (P) and non-poor (N) by different estimators

The sixth column (labelled "Mismeasured misclassification") uses the estimates of the symmetric misclassification model to calculate the implied frequency of poverty histories for the mismeasured but observed poverty variable. This model does much better in explaining the frequency of poverty histories than the model without misclassification error. The absolute deviation, giving the extent to which the predicted poverty histories differ from the data, shows that the inclusion of a single misclassification parameter, π, brings a 37% improves in goodness-of-fit. The seventh column (lablled "True misclassification") calculates the poverty history frequencies for the same symmetric misclassification model, but now using the true poverty variable without any classification error. The first obvious conclusion from this column is that allowing for classification error decreases poverty mobility: 20% of population are in chronic poverty, i.e. poor in all three waves, which is nearly 1.5 times greater than the chronic poverty suggested by the data, and 1.6 times greater than the chronic poverty implied by the estimator that assumes no misclassification. Secondly, accounting for classification error causes all other poverty histories, in which at least one transition occurs, to be over-estimated. This is particularly true for histories with transitions and reversals (PNP and NPN) which are quite common in the actual data (with frequencies of 4.9% and 5.3%), but are predicted to be very rare once misclassification is accounted for (0.2% and 0.5%).

4.2 Estimating consumption mobility

Turning to the continuous estimator, we investigate mobility in the continuous household consumption measure. Table 7 reports estimates of the seven regression coefficients discussed in section 3. All regressions were estimated using OLS and the balanced Nigerian household panel data. The estimate in column 1 represents the convergence coefficient for waves 1 and 2 of the panel. The negative and significant coefficient estimate of $\hat{\theta_1} = -0.337$ is commonly interpreted as evidence of consumption convergence. The magnitude of the estimate provides additional information that imply a very high degree of consumption mobility between waves 1 and 2. If we the data is measured without error and consumption shocks are persistent, then $\beta = -0.337$ implies that we would expect approximately 33.7% of the consumption gap between any two households to be eliminated between successive waves. In a stationary process, this in turn implies that we can expect half of the consumption gap between any two households to be eliminated in approximately $\frac{0.69}{-\log(\beta+1)} = 1.68$ waves, or in roughly 40 months. This rate of convergence is the average across the entire period, so Column 2 reports the same regression estimated on waves 2 and 3 of the data, which points to mobility that is virtually identical (p-value = 0.839) to that observed between waves 1 and 2 ($\hat{\theta}_2 = -0.342$).

Table 7: Consumption regressions

	Δy_2	Δy_3	$y_1 - y_3$	Δy_3	Δy_3	Δy_3
	(1)	(2)	(3)	(4)	(5)	(6)
y_1	-0.337***		-0.349***	-0.0124	0.373***	
	(0.0152)		(0.0166)	(0.0183)	(0.0220)	
y_2		-0.342***			-0.580***	
		(0.0189)			(0.0230)	
Δy_2						-0.482***
						(0.0225)
N	4217	4217	4217	4217	4217	4217
R^2	0.160	0.146	0.142	0.000	0.242	0.198

Notes: OLS regressions with robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The coefficient estimate of convergence between waves 1 and 3 ($\hat{\theta}_3 = -0.349$) is reported in column 3 of Table 7. Viewed on its own, this suggests a high degree of consumption mobility between waves 1 and 3. However, this estimate also indicates that economic mobility over three waves of data is not significantly greater than what is observed over two waves, one of the tell-tale signs of measurement error. As noted in table 3, in the absence of measurement error, we would expect the total convergence between waves 1 and 3 to be $\hat{\theta}_3 = (\hat{\theta}_1 + 1)(\hat{\theta}_2 + 1) - 1 = -0.56$ (std. err. = 0.016). In other words, we would expect approximately 56% of wave 1 consumption gaps to have been eliminated in wave 3. However, the estimated regression coefficient is significantly lower than that (p-value < 0.0001), which is consistent with measurement error in consumption. Many households that seemingly experienced high degrees of mobility between waves 1 and 2 had evidently reverted back to consumption levels closer to those in wave 1 y wave 3.

The remaining four regression coefficients are similarly difficult to explain under the assumption of no measurement error. Wave 1 consumption has a surprisingly weak effect on consumption growth between waves 2 and 3 (column 4). When using all three waves of data to estimate consumption as an AR(2) process, we observe that consumption growth between waves 2 and 3 depend negatively on wave 2 consumption and positively on wave 1 consumption (column 5). There is a surprisingly large negative correlation between consumption growth between waves 1 and 2, and between waves 2 and 3 (column 6). All of these patterns are consistent with classical measurement error in consumption.

We can use the expected values of the regression coefficient estimates in the presence of measurement error from table 3 to construct a system GMM estimator that can estimate the β and α parameters from the coefficients in all six of the regressions in table 7. Table 8 reports the estimates from five different versions of this model. Column 1 presents the estimate of parameter β under the assumption of no measurement error ($\alpha = 0$). It therefore attempts to find a single value of β that is best able to explain all seven regression coefficients from table 7. We see that this estimate is significantly smaller than the two-wave convergence coefficient, but still very high. However, the J-test statistic and associated p-value for this model reveals that the over-identifying restrictions – which require that the simple estimated model offers a congruent explanation of all the estimated regression coefficients – is confidently rejected. The model in column 2 relaxes the assumption of time-invariant mobility by allowing different estimates of consumption mobility between waves 1 and 2 (β_2) and waves 2 and 3 (β_3), while still assuming that consumption is measured without error. We find no evidence that consumption mobility changed over time (p-value = 0.208), and can still confidently reject the hypothesis that this model offers a coherent account of all the estimated regression coefficients.

Table 8: GMM estimation: consumption mobility

	(1)	(2)	(3)	(4)	(5)
β	-0.256***		-0.0256		
	(0.00799)		(0.0190)		
β_2		-0.260***		-0.0305	-0.0188
		(0.00955)		(0.0199)	(0.0273)
β_3		-0.233***		-0.0111	-0.0447*
		(0.0180)		(0.0267)	(0.0244)
α			0.674***	0.674***	
			(0.0200)	(0.0199)	
α_1					0.694***
					(0.0285)
α_2					0.671***
					(0.0200)
Observations	4217	4217	4217	4217	4217
J-test statistic	186.5	184.1	1.628	1.043	5.29e-27
p-value	3.55e-40	1.04e-40	0.443	0.307	.
H_0: $\beta_2 = \beta_3$ (p-value)		0.208		0.441	0.315
H_0: $\alpha_1 = \alpha_2$ (p-value)					0.299

Notes: GMM regressions with clustered standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column 3 allows for measurement error. The point estimate of the convergence coefficient is still negative, but much smaller and not statistically different from zero. This model also indicates that consumption is measured with substantial error: only 67.4% of variation in the observed consumption measure is due to variation in actual consumption, while the rest is attributed to measurement error. The J-test reveals that this model does offer a coherent account of all the regression coefficients. Our very simple two-parameter model is therefore able to explain many of the subtle three-wave dynamics of the observed consumption measure. According to these results, the conventional two-wave convergence estimates exaggerate the extent of consumption mobility by a factor of 13.

The estimates in column 4 also allow for time-varying consumption mobility, while column 5 also allow the data reliability coefficient to change over time. We find no evidence of significant changes of mobility or measurement error across the waves.

4.2.1 Simulated poverty transitions

The implications of our GMM estimates for poverty transitions can be investigated using simulations. We start with a different version of the model that estimates the variances of consumption shocks, σ_u^2, and measurement error, σ_e^2, instead of α, along with β. These estimates are combined with the observed wave-specific means of log household consumption per capita to simulate a three-wave consumption dynamics data set containing measures for both true (unobserved) and mismeasured (observed) consumption. This simulation requires making additional distributional assumptions that were not necessary for estimating the model parameters. We assume that wave 1 log household consumption per capita is normally distributed, and that multiplicative consumption shocks and measurement errors are normally distributed. The simulation is run for 10,000,000 households and a binary poverty variable is constructed from the simulated consumption values using a poverty line of $2.15, the same poverty line that was used to describe poverty transitions in table 2. Table A2 in the Appendix reports the two and three wave transition matrices for the observed but mismeasured poverty measure from our simulated data. These rates are also summarised (in the columns labelled Continuous Model) in table 9. We can compare these simulated poverty rates and transition probabilities to the observed counterparts from the Nigerian data to assess the accuracy of our simulation model.

The simulated poverty rates are slightly below 30% for all three waves, whereas the poverty rates in the actual data is slightly above 30% for all waves. Transition rates into and out of poverty across waves 1 and 2 differ by less than 1pp between the simulated and actual data. The simulated rate of transitions into poverty between waves 2 and 3 and waves 1 and 3 are also quite similar to the actual data. However, the simulated rate of poverty exit between waves 2 and 3 is about 8pp higher, and about 5pp higher between waves 1 and 3 than in the actual data. This difference is driven by subtle changes in the distribution of the actual consumption distribution between waves that is not accurately reflected in our normality assumption. Nevertheless, the simulated data does a decent job at replicating the main attributes of the actual data.

Table A2 in the Appendix reports the transition matrices for the true but unobserved poverty measure obtained from the same simulation exercise. Transition rates are much lower than for the observed but mismeasured poverty status variable. Over two successive waves, transitions out of poverty only happen about a third as often as suggested by the observed poverty measure. The observed poverty exit rate across two successive waves, according to our simulation, is around 40%, whereas the true rate is only 13%. Over three waves the transition rate is also exaggerated, but only by a factor of 2.2. As we would expect, the effect of measurement error is particularly onerous over shorter periods. Transitions into poverty are over-estimated by an even larger factor: between 4.5 and 5.8 over two waves, and 4.4 over three waves. Observed poverty entry rates are 15%, but after accounting for measurement error this drops to 3%.

Table A3 in the Appendix also reveals the simulated probability of misclassifying households as ether poor or non-poor. Across all three waves, the probability of misclassifying a poor household as non-poor is about 24%, while the probability of misclassifying a non-poor household as poor is 13%. The results of our simulation suggests that while naive two-wave transition matrices that ignore measurement error in consumption may estimate high rates of poverty transitions, in reality such transitions may be much less common.

Table 9: Comparison of observed and simulated transition and misclassification rates

		Observed & mismeasured			Unobserved & true	
		Actual data	Simulated data			
			Continuous Model	Binary Model	Continuous Model	Binary Model
Exit rate	Wave 1 to 2	39.0%	38.1%	34.9%	12.6%	13.2%
	Wave 2 to 3	33.7%	42.0%	34.9%	14.6%	13.2%
	Wave 1 to 3	39.4%	44.2%	40.9%	19.5%	24.0%
Entry rate	Wave 1 to 2	16.6%	16.2%	16.3%	3.6%	5.3%
	Wave 2 to 3	15.4%	15.2%	16.3%	2.6%	5.3%
	Wave 1 to 3	17.5%	16.3%	19.3%	3.7%	9.6%

5 Robustness tests

5.1 Indicators of unreliable responses

The results above illustrate how the observed pattern of poverty transitions and consumption mobility violate the Markov property, and how a parsimonious model that allows for measurement or classification error can explain these dynamics. However, violation of the Markov property can occur for a variety of reasons other than measurement error, such as duration-dependence or population heterogeneity in transition rates.

Measurement error in surveys can occur for several reasons. Respondents may misunderstand questions or answer untruthfully; survey workers may hear or write down responses inaccurately; or households may be incorrectly linked across waves. In all of these cases, evidence of an unreliable response to one question is likely to increase the probability of unreliable responses to other questions. We therefore identify households and waves in which some questions were answered inconsistently and test whether these households also experienced inflated transition rates compared to households with consistent responses. We find evidence that inconsistent responses to survey questions on age, sex, and marital status are correlated with the high short term poverty transition rates, which supports the hypothesis that classification error is at least partly responsible for the observed pattern of poverty dynamics.

We construct a dummy variable, $inconsistent_t$, with a value of 1 if there is an inconsistency in either the age, sex, or marital status reported by the respondent across the two visits of a given panel wave. An inconsistency is considered to occur when: a) the respondent's age either increases by more than 2 years or decreases; b) the respondent's sex changes; or c) a respondent who previously reported being married now reports having never been married.

The estimator derived in section 3.1 is expanded to a nonlinear systems estimator to allow a test for the presence of classification error in poverty status. The untransformed π^* classification error parameter in equation 2 becomes π^*_t and is expanded as

$$\pi^*_t = \pi^*_0 + \pi^*_1 inconsistent_t \tag{6}$$

where π^*_0 and π^*_1 are unobservable parameters that determine the probaility of misclassification for households with and without inconsistent responses.

The implied probabilities from the estimation of this extended nonlinear systems estimator can be found in table 10, while the untransformed parameter estimates are shown in table A4 in the appendix. The results show that respondents with inconsistent responses are significantly more likely to reveal poverty dynamics consistent with classification error: namely, the high short term transition rates along with relatively lower longer term transitions.

The misclassification rate for those with inconsistent responses in other survey questions is 1.6 percentage points higher than those without inconsistent responses, where the former have a 9.42% and the latter a 11.02% poverty misclassification rate. This difference has strong statistical significance, as seen in table A4, and the difference in the estimates, while not exceptionally large, is non-negligible. However, we would not expect the difference between these two misclassification rates to be very large for two reasons. First, the inconsistency dummy is by no means a perfect predictor for measurement error in the consumption aggregate, but rather a noisy indicator of diminished reliability for a particular respondent-fieldworker pair. There are many reasons why, for example, age could be misreported in a particular visit and not all of them imply that the respondent's reported consumption be unreliable. Second, the most likely cause of an inconsistency as described here is that there is a clear error in one of the visits in a particular wave. If that respondent is less reliable and misreports consumption during that visit, the extent of the measurement error in consumption will be smoothed by the averaging of the consumption aggregates from the two visits when creating the wave's consumption aggregate. Therefore, large errors at the visit level will be smaller errors at the wave error, if there is no serial correlation in measurement error, due to the way in which each wave's consumption aggregate is calculated. These reasons mean we would not expect the inconsistencies to have a very substantial increase in the poverty misclassification rate, but also makes the strong statistical significance of the result all the more compelling as an argument in favour of the presence of classification error.

Table 10: Regression results including symptoms of measurement error —implied probabilities

	(1)
	\$2.15
	Inconsistencies
Miscl. no inconsistency (%)	9.42***
	(0.009)
Miscl. with inconsistency (%)	11.02***
	(0.023)
Entry Rate (%)	5.35***
	(0.011)
Exit Rate (%)	13.26***
	(0.007)
Poverty Rate (%)	26.80***
	(0.027)
Log lik (mill.)	-47.41
Misclassification	Yes
Symptom included	Inconsistencies

Note: '***' as 0.001, '**' as 0.01, '*' as 0.05

$$s_t^* \equiv \gamma_1 s_{t-1}^* + \gamma_2(1 - s_{t-1}^*) + v_t^* \tag{7}$$

5.2 AR(2) duration-dependence

To allow for higher-order dependence, we extend equation 7 by defining true poverty transitions as an autoregressive process of order 2. Thus, the Markov assumption is no longer required, and the probability of exiting poverty for the chronically poor is allowed to differ to that of the transient poor. The assumed process is now defined as

$$s_t^* \equiv \gamma_{22} + (\gamma_{12} - \gamma_{22})s_{t-1}^* + (\gamma_{21} - \gamma_{22})s_{t-2}^* + (\gamma_{11} - \gamma_{22})s_{t-1}^* s_{t-2}^* + e_t \tag{8}$$

where the γ's are unobserved and estimable parameters and e_t is an unobservable random error term that is mean independent of lagged true employment such that $E(e_t|s_{t-1}^*, s_{t-2}^*) = 0$ and follows a Bernoulli distribution when conditioned on both lagged statuses. For example, when $s_{t-1}^* = 1, s_{t-2}^* = 0$ such that an individual recently found employment, the error term would be equal to $1 - \gamma 12$ and $-\gamma 12$. This implies the following population moments

$$\begin{aligned}
&P(s_t^* = 1|s_{t-1}^* = 1, s_{t-2}^* = 1) = \gamma_{12} + \gamma_{21} + \gamma_{11} - 2\gamma_{22} \\
&P(s_t^* = 1|s_{t-1}^* = 1, s_{t-2}^* = 0) = \gamma_{12} \\
&P(s_t^* = 1|s_{t-1}^* = 0, s_{t-2}^* = 1) = \gamma_{21} \\
&P(s_t^* = 1|s_{t-1}^* = 0, s_{t-2}^* = 0) = \gamma_{22}
\end{aligned} \tag{9}$$

The results can be found in table 11. The main result is that the misclassification rate remains significant, and even small classification error rates have a large bias on transition estimates (Burger, 2016; Breen and Moisio, 2004). The unconditional poverty rate remains fairly constant. Those in chronic poverty are between 10 and 15 percentage points more likely to remain poor. However, we cannot reject the Markov assumption on the true poverty process which suggests that the more parsimonious model in equation 7 is appropriate. Overall, these results provide further evidence in favor of the hypothesis that classification error is present and makes poverty appear more transitory than it truly is.

Table 11: Structural regression results of AR(2) employment process with 4 waves of data: implied probabilities

	AR(2) with misclassification
Miscl. Rate	2.56%**
	(0.60%)
$\mathbb{E}(s_t^*\|s_{t-1}^* = 1, s_{t-2}^* = 1)$	79.08%**
	(0.60%)
$\mathbb{E}(s_t^*\|s_{t-1}^* = 1, s_{t-2}^* = 0)$	70.19%**
	(0.89%)
$\mathbb{E}(s_t^*\|s_{t-1}^* = 0, s_{t-2}^* = 1)$	64.81%**
	(1.10%)
$\mathbb{E}(s_t^*\|s_{t-1}^* = 0, s_{t-2}^* = 0)$	54.51%**
	(0.44%)
$\mathbb{E}(s_{t-2}^* = 1)$	34.55**
	(1.37%)
log-lik	- 2941.6
Misclassification	Yes

6 Conclusion

This paper considers the estimation of poverty dynamics and consumption mobility using multi-wave panel surveys where consumption aggregates are known to be measured with error. We propose two innovative structural approaches that utilize three wave panel surveys to estimate these dynamics: the first models binary poverty directly without relying on distributional assumptions, while the second models continuous household consumption and use simulations to gauge the effect on poverty dynamics. We estimate these models on the panel data component from the Nigerian General Household Survey and show that the approaches come to similar conclusions: both find evidence of measurement error in the consumption aggregate that inflates the rate of poverty transitions and causes chronic poverty to appear transitory. Although most of our results are also consistent with more complicated consumption dynamics – where consumption is an AR(2) process or most consumption shocks are transient as opposed to persistent – we also present evidence that households who provided unreliable responses to non-consumption questions reveal poverty dynamics more consistent with measurement and classification error, and we test the Markov assumption to show that it is appropriate.

The two modelling approaches provide broadly consistent results regarding the important role of measurement or classification error, and the effect of such errors on inflating short-term churn rates. However, the results from the two models are not identical. The binary estimator that models poverty dynamics directly provides a better fit the two- and three-wave poverty exit rates than the simulations from the continuous consumption estimator. One reason for this is that the continuous model requires making perhaps unwarranted distributional assumptions. There are also other, more subtle differences between the two approaches. Although both invoke some version of the Markov property, this assumption has different implications when applied to consumption and poverty dynamics. Assuming that consumption is an AR(1) process and that consumption shocks are i.i.d. implies that two households with the same wave 2 consumption level have the same prospects for future consumption growth, even if their wave 1 consumption levels were very different. The Markov property for the binary poverty measure implies that future poverty transitions do not depend on past poverty statuses. However, it is likely that households that recently transitioned out of poverty have current consumption levels that are barely above the poverty line, so that past poverty status may contain information about proximity to the poverty line. Such households would naturally have a greater inclination to fall back into poverty, even if we assume that consumption is an AR(1) process and consumption shocks are i.i.d. In this regard, the Markov property for binary poverty entails a stronger assumption about consumption persistence that is more difficult to justify. Future work should focus on allowing the binary approach to including information about proximity to the poverty line through relaxing the Markov assumption made in equation 7 and allowing true poverty to be an AR(2) process.

Our results suggest possible avenues for future work. It is possible to extend our structural estimators to account for other factors that may cause violation of the Markov assumption. One such factor is household-level heterogeneity, where a small number of households have highly volatile consumption in contrasts the relatively stable consumption of the majority of households. Re-estimating our models while explicitly allowing transition rates to depend on observed heterogeneity (e.g. self-reported consumption shocks) or unobserved heterogeneity (via a finite mixture of unobservable household types) should provide a more accurate identification of the role of measurement error.

It would also be simple to apply these techniques to other welfare measures used to define poverty, like asset-based and multidimensional poverty. These measures typically reveal less frequent poverty transitions than consumption poverty, which is often attributed to consumption being a more volatile welfare measure. However, consumption may also be more prone to measurement error than household assets or the components of multi-dimensional poverty, which provides an alternative explanation for more frequent consumption poverty transitions.

References

Abowd, J. M. and Stinson, M. H. (2013), 'Estimating measurement error in annual job earnings: A comparison of survey and administrative data', Review of Economics and Statistics **95**(5), 1451–1467.

Abowd, J. and Zellner, A. (1985), 'Estimating gross labor force flows', Journal of Business & Economic Statistics **3**, 254–283.

Aliber, M. (2003), 'Chronic poverty in south africa: Incidence, causes and policies', World development **31**(3), 473–490.

Antman, F. and McKenzie, D. (2007), 'Poverty traps and nonlinear income dynamics with measurement error and individual heterogeneity', The journal of development studies **43**(6), 1057–1083.

Ashenfelter, O., Deaton, A. and Solon, G. (1986), Collecting panel data in developing countries: Does it make sense?, Technical Report 23, The World Bank, Washington.

Battistin, E. (2003), Errors in survey reports of consumption expenditures, Technical report, IFS Working Papers.

Baulch, B. and Hoddinott, J. (2000), 'Economic mobility and poverty dynamics in developing countries', The Journal of Development Studies **36**(6), 1–24.

Biemer, P. and Bushery, J. (2000), 'On the validity of markov latent class analysis for estimating classification error in labor force data', Survey Methodology **26**(2), 139–152.

Breen, R. and Moisio, P. (2004), 'Poverty dynamics corrected for measurement error', The Journal of Economic Inequality **2**, 171–191.

Brooks-Gunn, J. and Duncan, G. J. (1997), 'The effects of poverty on children', The future of children pp. 55–71.

Burger, C. (2016), A Structural Approach to Modeling South African Labour Market Decisions, PhD thesis, Stellenbosch University.

Calvo, C. and Dercon, S. (2009), 'Chronic poverty and all that', Poverty dynamics: interdisciplinary perspectives pp. 29–58.

Chua, T. and Fuller, W. (1987), 'A model for multinomial response error applied to labor flows', Journal of American Statistical Association **82**(397), 46–51.

Dapel, Z. (2018), 'Three decades of poverty mobility in nigeria: The trapped, the freed, and the never trapped', Center for Global Development Working Paper (485).

Dempster, A., Laird, N. and Rubin, D. (1977), 'Maximum likelihood from incomplete data via the em algorithm', Journal of the Royal Statistical Society. Series B (Methodological) **39**(1), 1–38.

Dinkelman, T. (2004), 'How household context affects search outcomes of the unemployed in kwazulu-natal, south africa: A panel data analysis', South African Journal of Economics **72**(3), 484–521.

Duclos, J.-Y., Araar, A. and Giles, J. (2010), 'Chronic and transient poverty: Measurement and estimation, with evidence from china', Journal of development Economics **91**(2), 266–277.

Feng, S. and Hu, Y. (2013), 'Misclassification errors and the underestimation of U.S. unemployment rates', American Economic Review **103**(2), 1054–1070.

Foster, J. E. (2009), 'A class of chronic poverty measures', Poverty dynamics: interdisciplinary perspectives pp. 59–76.

Friedman, J., Beegle, K., De Weerdt, J. and Gibson, J. (2017), 'Decomposing response error in food consumption measurement: Implications for survey design from a randomized survey experiment in tanzania', Food Policy **72**, 94–111.

Green, M. and Hulme, D. (2005), 'From correlates and characteristics to causes: thinking about poverty from a chronic poverty perspective', World Development **33**(6), 867–879.

Hausman, J. A., Newey, W. K. and Powell, J. L. (1995), 'Nonlinear errors in variables estimation of some engel curves', Journal of Econometrics **65**(1), 205–233.

Hulme, D. and Shepherd, A. (2003), 'Conceptualizing chronic poverty', World development **31**(3), 403–423.

Idika, E., Nwosu, E., Nji, I. and Ejimonye, J. (2021), 'Poverty change and mobility in urban and rural areas of geopolitical zones in nigeria: A decomposition analysis', Journal of Home Economics Research (JHER) p. 69.

Jalan, J. and Ravallion, M. (1998), Determinants of transient and chronic poverty: evidence from rural China, number 1936, World Bank Group.

Jo Bane, M. and Ellwood, D. T. (1983), 'Slipping into and out of poverty: The dynamics of spells'.

Jolliffe, D. M., Mahler, D. G., Lakner, C., Atamanov, A. and Tetteh Baah, S. K. (2022), 'Assessing the impact of the 2017 ppps on the international poverty line and global poverty'.

Lanjouw, J. O. and Lanjouw, P. (1997), Poverty comparisons with non-compatible data: theory and illustrations, number 1709, World Bank Publications.

Magnac, T. and Visser, M. (1999), 'Transition models with measurement errors', The Review of Economics and Statistics **81**(3), 466–474.

McDonough, P. and Berglund, P. (2003), 'Histories of poverty and self-rated health trajectories', Journal of health and social behavior pp. 198–214.

Moore, J. C., Stinson, L. L., Welniak, E. J. et al. (2000), 'Income measurement error in surveys: A review', Journal of Official Statistics-Stockholm- **16**(4), 331–362.

Poterba, J. and Summers, L. (1986), 'Reporting errors and labor market dynamics', Econometrica **54**(6), 1319–1338.

Ravallion, M. (2003), 'Measuring aggregate welfare in developing countries: How well do national accounts and surveys agree?', Review of Economics and Statistics **85**(3), 645–652.

Rendtel, U., Langeheine, R. and Berntsen, R. (1998), 'The estimation of poverty dynamics using different measurements of household income', Review of Income and Wealth **44**(1), 81–98.

Schennach, S. M. (2016), 'Recent advances in the measurement error literature', Annual Review of Economics **8**(1), 341–377.

Shibata, I. (2022), 'Reassessing classification errors in the analysis of labor market dynamics', Labour Economics **78**, 102252.

Singh, A. and Rao, J. (1995), 'On the adjustment of gross flow estimates for classification error with application to data from the canadian labour force survey', Journal of the American Statistical Association **90**, 478–488.

Ward, Z. (2023), 'Intergenerational mobility in american history: Accounting for race and measurement error', American Economic Review **113**(12), 3213–3248.

Wolfe, R., Carlin, J. and Patton, G. (2003), 'Transitions in an imperfectly observed binary variable: Depressive symptomatology in adolescents', Statistics in Medicine **22**(3), 427–440.

Wollburg, P., Tiberti, M. and Zezza, A. (2021), 'Recall length and measurement error in agricultural surveys', Food Policy **100**, 102003.

World Bank (2022), Poverty and shared prosperity 2022: Correcting course, The World Bank.

A Appendix

Table A1: Simulated poverty transitions: observed poverty

Wave 1	Wave 2		
	Non-poor	Poor	Total
Non-poor	59.17	11.40	70.57
Poor	11.20	18.23	29.43
Total	70.37	29.63	100.00
Probability of exiting poverty: 38.1%			
Probability of entering poverty: 16.2%			

Wave 2	Wave 3		
	Non-poor	Poor	Total
Non-poor	59.67	10.70	70.37
Poor	12.43	17.20	29.63
Total	72.10	27.90	100.00
Probability of exiting poverty: 42.0%			
Probability of entering poverty: 15.2%			

Wave 1	Wave 3		
	Non-poor	Poor	Total
Non-poor	59.08	11.49	70.57
Poor	13.02	16.41	29.43
Total	72.10	27.90	100.00
Probability of exiting poverty: 44.2%			
Probability of entering poverty: 16.3%			

Table A2: Simulated poverty transitions: true poverty

Wave 1	Wave 2		
	Non-poor	Poor	Total
Non-poor	71.33	2.69	74.02
Poor	3.27	22.71	25.98
Total	74.60	25.40	100.00
Probability of exiting poverty: 12.6%			
Probability of entering poverty: 3.6%			

Wave 2	Wave 3		
	Non-poor	Poor	Total
Non-poor	72.67	1.93	74.60
Poor	3.70	21.70	25.40
Total	76.37	23.63	100.00
Probability of exiting poverty: 14.6%			
Probability of entering poverty: 2.6%			

Wave 1	Wave 3		
	Non-poor	Poor	Total
Non-poor	71.30	2.72	74.02
Poor	5.07	20.91	25.98
Total	76.37	23.63	100.00
Probability of exiting poverty: 19.5%			
Probability of entering poverty: 3.7%			

Table A3: Simulated poverty misclassifications

True (Wave 1)	Mismeasured (Wave 1)		
	Non-poor	Poor	Total
Non-poor	64.42	9.60	74.02
Poor	6.15	19.83	25.98
Total	70.57	29.43	100.00
Probability of misclassifying Poor as Non-poor: 23.7%			
Probability of misclassifying Non-poor as Poor: 13.0%			
True (Wave 2)	**Observed (Wave 2)**		
	Non-poor	Poor	Total
Non-poor	64.39	10.21	74.60
Poor	5.98	19.42	25.40
Total	70.37	29.63	100.00
Probability of misclassifying Poor as Non-poor: 23.5%			
Probability of misclassifying Non-poor as Poor: 13.7%			
True (Wave 3)	**Observed (Wave 3)**		
	Non-poor	Poor	Total
Non-poor	66.13	10.24	76.37
Poor	5.97	17.66	23.63
Total	72.10	27.90	100.00
Probability of misclassifying Poor as Non-poor: 25.3%			
Probability of misclassifying Non-poor as Poor: 13.4%			

Table A4: Regression results including symptoms of measurement error —parameter outputs

	(1)
	\$2.15
	Inconsistencies
π_0^* (Miscl. no inconsistency)	1.32^{***}
	(0.0004)
π_1^* (Miscl. inconsistency coef.)	-0.09^{***}
	(0.001)
γ_1^* (Povert exit)	1.11^{***}
	(0.001)
γ_2^* (Poverty entry)	-1.61^{***}
	(0.0008)
μ (Poverty rate)	-0.62^{***}
	(0.0003)
Log lik (mill.)	-47.41
Misclassification	Yes
Symptom included	Inconsistencies

Note: '***' as 0.001, '**' as 0.01, '*' as 0.05