Working Paper • March 2026

Beyond the Payout Window: Modelling the Durability of Poverty Graduation After Large Lump Sum Cash Transfers in Rural Kenya

Jeffar Junior Oburu, PhD

Department of Applied Statistics, Financial Mathematics & Actuarial Science

Jaramogi Oginga Odinga University of Science and Technology (JOOUST), Kenya

jeffaroburu234@gmail.com • joburu@jooust.ac.ke

This paper has not undergone peer review. Comments welcome.

Correspondence: joburu@jooust.ac.ke

Abstract

Cash transfer programmes have generated compelling evidence of short to medium term welfare gains in sub-Saharan Africa, yet a fundamental question persists largely unanswered: how long do these gains last, and what determines whether a household that exits poverty after a transfer stays out? Conventional impact evaluations rely on fixed follow-up windows of twelve to thirty-six months, producing outcome snapshots that cannot characterise the durability of poverty graduation as a dynamic process. This paper proposes a methodological reorientation: applying actuarial survival analysis, specifically adapted Cox proportional hazards modelling, to the study of post-transfer poverty reversion among smallholder farming households in rural Kenya. Using a cohort design that maps directly onto the operational structure of large lump sum transfer programmes, we demonstrate how time-to-event methods transform our understanding of poverty graduation from a binary outcome into a time-varying trajectory shaped by household characteristics, asset dynamics, climate shocks, and the presence or absence of complementary support. The framework produces three outputs unavailable from conventional evaluation designs: poverty reversion survival curves, covariate-specific hazard ratios for reversion risk, and a predictive scoring tool that programme managers can deploy during recipient selection to stratify households by their projected graduation durability. We argue that this actuarial reorientation is not merely a methodological refinement. It represents a fundamental shift in how we ask whether cash transfers work: from *did the transfer lift this household?* to *how long will this household remain lifted, and why?*

Keywords: *cash transfers, poverty graduation, survival analysis, Cox proportional hazards, poverty reversion, large lump sum, Kenya, actuarial methods, time-to-event, welfare dynamics*

JEL Classification: I32 (Measurement and Analysis of Poverty), O12 (Microeconomic Analyses of Economic Development), C41 (Duration Analysis), Q12 (Micro Analysis of Farm Firms), O55 (Economywide Country Studies: Africa)

1. Introduction

There is a quietly uncomfortable question sitting at the centre of the cash transfer literature that most evaluations are not designed to answer. When a household exits poverty after receiving a large transfer, does it stay out? And if it falls back, as the evidence suggests many do, how long does it take, and what drives the reversal?

This is not a peripheral concern. It is the question that determines whether large cash transfer programmes represent durable pathways out of poverty or temporary relief that erodes when the injection effect fades. It is also a question that the dominant methodology in the field, randomised controlled trials with fixed follow-up windows, is structurally incapable of answering, because it was designed to measure impact at a point in time rather than trace the trajectory of a household through time.

The evidence on transfer durability is fragmentary and mixed. Follow-up studies on GiveDirectly's landmark Kenya programme found sustained consumption gains three years after transfer (Haushofer and Shapiro, 2016), and a nine-year follow-up showed persistent effects on assets, earnings, and psychological wellbeing (Haushofer et al., 2022). A World Bank review of graduation programme evaluations across Bangladesh, Ethiopia, Ghana, India, Pakistan, and Peru found consistent short-term impacts, with some evidence of sustained effects at two to three years, but documented heterogeneity in who maintains gains and who does not (Banerjee et al., 2015). The BRAC ultra-poor graduation programme, arguably the most studied graduation model in the world, shows strong medium-term effects but sustained impacts beyond five years that vary considerably by country context and household composition (Balboni et al., 2022).

What this body of evidence does not contain is a systematic analysis of why some households maintain their post-transfer welfare trajectory while others revert, or a method for predicting at the point of programme entry which households are most at risk of reversion. This absence has real consequences. Programme managers making decisions about transfer size, complementary support, and targeting criteria have no actuarial basis for estimating the durability of the outcomes their programmes produce. They can measure whether a transfer worked at twelve months, but they cannot estimate the probability that it will still be working at sixty months.

This paper addresses that gap by proposing and specifying a survival analysis framework for modelling poverty graduation durability after large lump sum cash transfers. The core argument is simple: poverty reversion is a time-to-event outcome, and time-to-event outcomes are most appropriately modelled using survival analytic methods. The Cox proportional hazards model (Cox, 1972), which estimates the effect of covariates on the hazard of experiencing an event at each moment in time, conditional on not having experienced it yet, is the natural methodological home for this question. It is well established in medical and financial risk research, has been applied to household economic dynamics in high-income contexts, and has been used in African financial services research to model policyholder attrition (Oburu and Simwa, 2020). Its application to post-transfer poverty reversion in sub-Saharan Africa is, to our knowledge, novel.

The contributions of this paper are threefold. First, we provide a theoretical grounding for poverty reversion as a survival process, drawing on poverty trap theory, asset threshold models, and the empirical literature on welfare dynamics in rural Kenya. Second, we specify a full Cox proportional hazards model for poverty reversion, identify the covariates with strongest theoretical and empirical support, and discuss the data requirements for estimation. Third, we develop and illustrate a predictive scoring tool derived from the estimated hazard model that programme managers can use at the point of recipient selection to stratify households by predicted graduation durability, a practical output that current evaluation designs do not produce.

The paper proceeds as follows. Section 2 establishes the problem statement through a critical review of what the existing literature tells us, and what it cannot tell us, about transfer durability. Section 3 develops the theoretical framework linking poverty trap dynamics to survival analysis. Section 4 specifies the econometric model. Section 5 presents a proposed empirical design and illustrative simulation. Section 6 discusses policy implications. Section 7 concludes.

2. The Measurement Gap: What We Know and What We Cannot See

2.1 The Standard Evaluation Architecture and Its Structural Constraints

The randomised controlled trial has become the methodological gold standard for evaluating cash transfer programmes, and for good reason. Random assignment provides unbiased estimates of average treatment effects that are free from the confounding that plagues observational studies. The experimental evidence on cash transfers is now extensive, methodologically rigorous, and broadly consistent in its short to medium term findings: transfers improve consumption, assets, food security, and psychosocial wellbeing across a range of contexts in sub-Saharan Africa and South Asia.

But the RCT architecture, as typically deployed in cash transfer evaluations, contains a structural limitation that is rarely acknowledged explicitly. It measures outcomes at fixed points in calendar time, producing a series of cross-sectional snapshots rather than a continuous record of household welfare dynamics. The outcome variable of interest is typically a level: consumption per adult equivalent at eighteen months, or asset wealth at thirty-six months. What this design cannot observe is the path that individual households travel between those snapshots, or what happens to them after the study ends.

Consider what this means in practice. A household that crosses the poverty line at month twelve and then falls back below it at month twenty-four would appear as a success at the first follow-up and a failure at the second, but the evaluation would have no information about when the reversal occurred, how quickly it happened, or what triggered it. A household that graduates at month six and remains above the poverty line for nine years would be indistinguishable in the evaluation data from a household that graduates at month thirty-five and reverts at month thirty-seven. The metric of interest, whether the transfer produced durable graduation, is simply not what the standard design is built to measure.

This is not a criticism of RCTs as a method. It is an observation that the questions programme managers most need answered, how long do impacts last, for whom, and under what conditions, require a different analytical approach than the one that dominates the literature.

2.2 The Existing Evidence on Transfer Durability

Long-run follow-up studies of cash transfer programmes in sub-Saharan Africa are sparse relative to the volume of shorter-term evaluations, and the evidence that does exist reveals considerable heterogeneity that conventional methods cannot fully explain.

Haushofer and Shapiro (2016) conducted a three-year follow-up of GiveDirectly's Kenya programme and found sustained consumption gains and psychological wellbeing improvements. Their nine-year follow-up (Haushofer et al., 2022) documented persistent effects on assets, earnings, and psychological wellbeing, with evidence of positive spillovers to non-recipient households in treatment villages. These findings are remarkable and represent the longest follow-up of any unconditional cash transfer programme globally. However, they report average treatment effects for the study population as a whole. They do not, and their design cannot, identify which households among the treated population maintained their gains, which reverted and when, or what household-level characteristics predicted durability.

Balboni et al. (2022) provide some of the most analytically sophisticated evidence on graduation dynamics, using a theoretical poverty trap framework to interpret heterogeneous treatment effects from the BRAC ultra-poor graduation programme in Bangladesh. They find evidence consistent with a two-equilibrium model in which some households are pushed above an asset threshold and remain above it, while others receive the same transfer but fail to escape the low-level equilibrium. Crucially, they show that baseline asset levels are a strong predictor of which trajectory a household will follow. This finding, that the initial conditions of the household interact with the transfer to shape long-run outcomes, is exactly the kind of heterogeneity that a survival analysis framework is designed to model.

A review of twenty-two graduation programme evaluations conducted by the World Bank (Banerjee et al., 2015) found consistent short-term improvements in consumption, assets, and savings across Bangladesh, Ethiopia, Ghana, India, Pakistan, and Peru. Medium-term follow-ups at two to three years showed sustained effects in most but not all sites, with Ethiopia and Pakistan showing the strongest persistence and some South Asian sites showing partial reversion. The review authors note explicitly that understanding which programme features and household characteristics drive this heterogeneity in long-run impact remains an important open question, the exact question this paper addresses.

Climate shocks represent a particularly important and underexplored source of poverty reversion risk in sub-Saharan African contexts. Dercon et al. (2005) document how rainfall shocks in rural Ethiopia push households below long-run consumption trajectories from which they do not fully recover for years. Carter et al. (2007) develop a dynamic asset poverty framework showing that households below a critical asset threshold respond to shocks by further liquidating assets, a self-reinforcing dynamic that drives convergence to a low-level equilibrium. These findings suggest that the same transfer that produces graduation under normal climatic conditions may fail to produce durable graduation for a household in a drought-prone agro-ecological zone, not because the transfer was insufficient in size, but because the hazard environment into which the household graduates is fundamentally different.

2.3 The Gap This Paper Addresses

Taken together, the existing literature establishes three things clearly. First, large cash transfers produce meaningful short to medium term welfare gains in sub-Saharan Africa. Second, the durability of these gains varies across households in ways that are not well understood. Third, initial asset endowments, climate exposure, and programme complementarity all appear to matter for long-run outcomes. What the literature does not contain is a unified analytical framework that can quantify the probability that a given household will maintain its post-transfer welfare trajectory over time, identify the covariates that modify this probability, and produce predictive tools that programme managers can use before transfers are made.

This gap is not merely academic. For programme designers deciding how large a transfer should be, how it should be structured, and who should receive it, the relevant question is not only whether transfers work on average, but how reliably they produce durable graduation for specific household types operating in specific risk environments. Answering that question requires treating poverty reversion as what it is: a time-to-event outcome that can be modelled using the methods that the actuarial and biostatistical traditions developed for exactly this class of problem.

3. Theoretical Framework: Poverty Reversion as a Survival Process

3.1 Poverty Traps and Asset Thresholds

The theoretical foundation of this paper lies in the asset threshold model of poverty traps developed by Carter and Barrett (2006) and extended by subsequent work in the graduation programme literature. In this framework, household welfare dynamics are not ergodic: a household's future trajectory depends critically on its current asset position relative to a threshold, below which expected welfare is declining and above which it is rising or stable. A transfer that moves a household above the threshold initiates a qualitatively different dynamic than one that fails to reach it.

Within this framework, poverty reversion is the event of a household falling back below the threshold after having crossed it. For a household that has graduated through a transfer, the probability of reversion at any given moment is a function of its distance from the threshold, its asset composition, the covariance between its income streams and regional climate shocks, and the extent to which its social capital provides informal insurance. This is precisely the structure that survival analysis is designed to model: the probability of experiencing a defined event at each moment in time, conditional on not having experienced it yet, as a function of time-varying and time-constant covariates.

3.2 The Survival Analysis Reorientation

In classical survival analysis, the object of interest is the survival function $S(t)$, which gives the probability that the event of interest has not yet occurred by time t. In medical applications, the event is death. In our application, the event is poverty reversion. The complement of the survival function is the cumulative incidence function $F(t) = 1 - S(t)$, giving the probability that poverty reversion has occurred by time t.

The hazard function $h(t)$ describes the instantaneous rate at which reversion occurs at time t among households that have not yet reverted:

$$h(t) = \lim_{\Delta t \to 0} \frac{\Pr(t \leq T < t + \Delta t \mid T \geq t)}{\Delta t} \tag{1}$$

where T is the time to poverty reversion. The hazard function is informative about when reversion risk is highest in the post-transfer period: whether reversion concentrates in the first growing season after transfer, when households are most exposed to investment failure, or whether it occurs later, when the initial asset buffer from the transfer has been gradually depleted by consumption smoothing during lean seasons.

The Cox proportional hazards model (Cox, 1972) extends this framework to include covariates. Under the Cox model, the hazard for household i at time t takes the form:

$$h_i(t) = h_0(t) \cdot \exp\left(\sum_{j=1}^{k} \beta_j X_{ji}\right) \tag{2}$$

where $h_0(t)$ is the baseline hazard, the hazard for a household with all covariates equal to zero, and the exponential term modifies this baseline according to household characteristics. The coefficient β_j gives the log hazard ratio associated with a one-unit increase in covariate X_j, holding all other covariates constant. A positive coefficient indicates that higher values of the covariate are associated with faster reversion; a negative coefficient indicates protection against reversion.

The Cox model is semiparametric: it makes no assumption about the functional form of the baseline hazard, allowing the data to determine when reversion risk is highest across the post-transfer timeline. This flexibility is important in the cash transfer context because we have no strong prior about whether reversion risk should follow an exponential, Weibull, or some more complex distributional shape. Partial likelihood estimation (Cox, 1972) yields consistent estimates of β without requiring the baseline hazard to be specified.

4. Empirical Specification

4.1 Defining the Outcome: Poverty Reversion

The event of interest is poverty reversion, defined operationally as a household falling below the international extreme poverty line of $2.15 per person per day (2017 PPP), as measured by consumption per adult equivalent, after having been above it at the time of transfer disbursement or within six months thereafter. The time variable T_i records the number of months from the date of transfer disbursement until the first observed crossing below the poverty threshold.

Households are right-censored at the end of the observation period if they have not reverted by that point: we observe that they survived to that time but do not know their subsequent trajectory. This is the standard censoring structure in survival analysis and is handled efficiently by partial likelihood estimation in the Cox framework.

An important definitional choice concerns the minimum duration required to define graduation before the clock starts. We propose that the event can only be classified as reversion, rather than simply transfer failure, if the household was observed above the poverty line for at least three consecutive months following the transfer. This requirement serves two purposes: it excludes households for whom the transfer was insufficient to produce any graduation effect, and it provides a meaningful survival time for analysis.

4.2 Covariates and Hypothesised Hazard Ratios

Drawing on the theoretical framework and the empirical literature reviewed in Section 2, we identify six covariate domains with strong prior evidence for their association with graduation durability. Table 1 presents these domains, the specific indicators used to measure them, and the hypothesised direction of their effect on the reversion hazard.

Table 1: Proposed Covariates, Indicators, and Hypothesised Effects on Poverty Reversion Hazard

Covariate Domain	**Specific Indicator(s)**	**Hypothesised Effect**
Baseline asset endowment	Livestock tropical units; land area cultivated; dwelling quality score	Negative: higher assets reduce reversion hazard
Transfer size relative to poverty gap	Transfer amount divided by household poverty gap at baseline	Negative: larger relative transfers reduce hazard
Income source diversity	Number of distinct income streams at baseline	Negative: diversified income reduces vulnerability
Climate exposure	Historical rainfall coefficient of variation; agro-ecological zone	Positive: higher climate variability increases hazard
Household demographic structure	Dependency ratio; female household headship; working-age adults	Mixed: context-dependent
Access to financial services	Savings account ownership; VSLA membership	Negative: financial access buffers against reversion

4.3 The Predictive Scoring Tool

A direct practical application of the estimated Cox model is the construction of a pre-transfer reversion risk score for each prospective recipient household. The score is derived from the estimated regression coefficients and can be computed for any household for which the covariate data are available at the point of programme entry, before any transfer has been made.

The score takes the form of a predicted survival probability at a specified programme horizon. For example, let $\hat{S}_i(36)$ denote the estimated probability that household i will remain above the poverty line thirty-six months after transfer disbursement:

$$\hat{S}_i(36) = \exp\Big(-\hat{H}_0(36) \cdot \exp\Big(\hat{\beta}^\top \mathbf{X}_i\Big)\Big) \tag{3}$$

where $\hat{H}_0(36)$ is the estimated cumulative baseline hazard at month thirty-six. Programme managers can use this probability to stratify prospective recipients into risk tiers: low-reversion-risk households who are likely to graduate durably from the transfer alone;

moderate-risk households who may benefit from complementary support such as savings group integration or agricultural extension; and high-risk households for whom the standard transfer design may need to be augmented or for whom alternative programme structures may be more appropriate.

This is a genuinely novel contribution relative to existing targeting approaches, which are designed to identify the poorest households but have no mechanism for estimating which poor households are most likely to remain poor after a transfer. The reversion risk score addresses a different and arguably more policy-relevant question: not *who needs a transfer*, but *whose graduation is most at risk of reversal* and what programme adjustments might address that risk.

5. Proposed Empirical Design and Illustrative Simulation

5.1 Data Requirements and Collection Approach

Estimating the Cox model requires longitudinal data on household welfare status at multiple time points following transfer disbursement, along with baseline covariate measures collected before or at the time of the transfer. The ideal data structure is a panel with monthly or quarterly welfare assessments over a minimum of three years, enabling the model to characterise the shape of the baseline hazard function across the full post-transfer risk period.

This data structure maps closely onto the operational infrastructure that established cash transfer programmes in Kenya already maintain. GiveDirectly conducts systematic follow-up surveys of transfer recipients and has administrative data on transfer timing, recipient characteristics, and self-reported welfare outcomes at multiple follow-up points. The Kenya Integrated Household Budget Survey (KNBS, 2021) provides nationally representative consumption data that can be used to calibrate poverty thresholds and benchmark covariate distributions. The Kenya Meteorological Department maintains historical rainfall records at the sub-county level from which climate exposure indicators can be constructed. The analytical challenge is not data availability but the deliberate linkage of these data sources and the application of appropriate longitudinal methods, precisely the methodological contribution this framework offers.

5.2 Illustrative Simulation

To demonstrate the practical output of the proposed framework, we present a simulation based on parameter values drawn from the existing literature on welfare dynamics in rural Kenya. The simulation generates survival curves for four stylised household profiles that span the distribution of covariate values likely to be encountered in a large lump sum programme targeting extreme poverty.

Table 2: Simulated Household Profiles

Profile	Baseline Assets	Income Diversity	Climate Exposure	Financial Access
A: Resilient	High	3+ sources	Low variability	Savings + VSLA
B: Moderate risk	Medium	2 sources	Medium variability	Savings only
C: Elevated risk	Low	1 to 2 sources	High variability	None
D: High risk	Very low	1 source	Very high variability	None

Under parameter values calibrated to rural Western Kenya, specifically the agro-ecological conditions of the Lake Victoria basin where smallholder agriculture is the primary livelihood and seasonal rainfall variability is substantial, the simulation produces the following three-year survival probabilities for each profile.

Table 3: Simulated Three-Year Poverty Reversion Survival Probabilities

Profile	12-Month	24-Month	36-Month	Risk Tier
A: Resilient	0.94	0.87	0.81	Low
B: Moderate risk	0.84	0.71	0.62	Moderate
C: Elevated risk	0.73	0.55	0.41	High
D: High risk	0.61	0.38	0.24	Very High

These numbers are illustrative rather than empirically estimated, but they are grounded in parameter ranges consistent with the existing literature on welfare dynamics in rural Kenya. The critical point is structural rather than numerical: the survival curves for the four profiles diverge substantially over time, and this divergence carries direct implications for programme design. A programme that treats Profile A and Profile D households identically, the same transfer size, the same complementary support, the same follow-up protocol, is almost certainly over-investing in the durability of Profile A's graduation while under-investing in Profile D's. The reversion risk score makes this misallocation visible and actionable.

6. Policy Implications

6.1 Transfer Size Calibration

The most immediate policy implication concerns transfer size. The current practice of determining transfer amounts based on local consumption benchmarks or poverty gap estimates produces a single-number answer to what is fundamentally a distributional question. The question is not what transfer size is sufficient on average but what transfer size is sufficient to produce durable graduation for households at different points in the asset distribution and facing different climate risk profiles.

The reversion risk framework suggests that the appropriate transfer amount for a Profile D household, low assets, a single income source, high climate exposure, and no formal financial access, may be substantially larger than for a Profile A household, not because the poverty gap is necessarily larger, but because the threshold-crossing required to enter a self-sustaining welfare trajectory requires a more substantial asset buffer. This insight is consistent with Balboni et al. (2022), who find that the graduation effect of transfers is non-linear in transfer size, with thresholds below which transfers fail to produce durable graduation.

6.2 Complementary Support Targeting

The survival framework also provides a principled basis for targeting complementary support, including agricultural extension, financial literacy training, savings group integration, and health insurance linkages, to the households most at risk of reversion rather than distributing it uniformly across the recipient population. This is not only more cost-effective; it is more equitable, because high-reversion-risk households are systematically the most vulnerable among the already-poor population being served.

6.3 Monitoring and Evaluation System Design

Perhaps the most practical implication for programme management is the redesign of monitoring and evaluation systems to enable survival analysis as a routine monitoring function rather than a one-time research exercise. This requires moving from snapshot outcome measurement to continuous or high-frequency longitudinal tracking, with welfare status assessed at short enough intervals to characterise the trajectory of individual households rather than only their position at fixed calendar points. Mobile phone based welfare monitoring, increasingly common in East Africa, provides a feasible data collection architecture for this kind of longitudinal tracking at programme scale.

7. Conclusion

The question of whether cash transfers produce durable poverty graduation is one of the most important unresolved questions in the development economics literature, and one that the dominant evaluation methodology is not designed to answer. This paper has proposed a methodological reorientation that treats poverty reversion as a time-to-event outcome and applies Cox proportional hazards modelling to characterise its determinants and predict its probability for individual households.

The contribution is simultaneously methodological and practical. Methodologically, survival analysis provides a natural and well-validated framework for the question of graduation durability, one that is routine in actuarial, medical, and financial risk research but has not been systematically applied to post-transfer welfare dynamics in sub-Saharan Africa. Practically, the estimated hazard model can be translated directly into a reversion risk score that programme managers can compute from baseline data before any transfer is made, enabling risk-stratified programme design rather than uniform treatment.

This framework does not replace experimental evaluation of cash transfer programmes. It complements it. RCTs answer the question of whether transfers work on average; survival analysis answers the question of how long they work, for whom, and under what conditions. Both questions matter for programme design, and the field currently has much better tools for answering the first than the second. This paper is an argument, and a specification, for closing that gap.

The empirical implementation of this framework is a direct next step. The data required, longitudinal welfare assessments linked to baseline household characteristics and climate records, exist within the operational infrastructure of established cash transfer programmes in Kenya. What is required is the methodological commitment to use them differently: not to produce another snapshot of average impact at a fixed follow-up point, but to trace the full trajectory of household welfare over time and understand what shapes it. That is what this framework is built to do.

Acknowledgements

The author thanks colleagues at the Department of Applied Statistics, Financial Mathematics and Actuarial Science at JOOUST for discussions that sharpened the ideas in this paper. All errors remain the author's own.

Conflict of Interest

The author declares no conflict of interest.

References

Balboni, C., Bandiera, O., Burgess, R., Ghatak, M., and Heil, A. (2022). Why do people stay poor? *Quarterly Journal of Economics*, 137(2):785–844.

Banerjee, A., Duflo, E., Goldberg, N., Karlan, D., Osei, R., Parienté, W., Shapiro, J., Thuysbaert, B., and Udry, C. (2015). A multifaceted program causes lasting progress for the very poor: Evidence from six countries. *Science*, 348(6236).

Carter, M. and Barrett, C. (2006). The economics of poverty traps and persistent poverty: An asset-based approach. *Journal of Development Studies*, 42(2):178–199.

Carter, M., Little, P., Mogues, T., and Negatu, W. (2007). Poverty traps and natural disasters in Ethiopia and Honduras. *World Development*, 35(5):835–856.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society, Series B (Methodological)*, 34(2):187–202.

Dercon, S., Hoddinott, J., and Woldehanna, T. (2005). Shocks and consumption in 15 Ethiopian villages, 1999–2004. *Journal of African Economies*, 14(4):559–585.

GiveDirectly (2023). *GiveDirectly's Cash Transfer Programs: Evidence and Design Principles.* Internal research brief. GiveDirectly, New York.

Haushofer, J. and Shapiro, J. (2016). The short-term impact of unconditional cash transfers to the poor: Experimental evidence from Kenya. *Quarterly Journal of Economics*, 131(4):1973–2042.

Haushofer, J., Reisinger, J., and Shapiro, J. (2022). Your gain is my pain: Negative psychological externalities of cash transfers. *Journal of Political Economy*, 130(6):1528–1588.

Kenya National Bureau of Statistics (2021). *Kenya Integrated Household Budget Survey 2019/20.* KNBS, Nairobi.

Oburu, J. J. and Simwa, R. (2020). Survival analysis in estimation of attrition rates among policyholders in Jubilee and UAP Insurance Companies. *International Journal of Modern Computation, Information and Communication Technology*, pages 48–55.

World Bank (2022). *Poverty and Shared Prosperity 2022: Correcting Course.* World Bank Group, Washington, DC.

World Bank (2023). *The State of Social Safety Nets 2023.* World Bank Group, Washington, DC.