LAIDLAW SCHOLARS

Laidlaw Scholars Undergraduate Leadership and Research Programme
Research Report

Sugarcoated Lives, Bitter Truths: Child Labor In Kenya's Sugarcane Industry

Linh Khanh Pham

Research Advisor: Professor Jiaying Gu

August 2026

STUDENT LIFE | Centre for International Experience

UNIVERSITY OF TORONTO

Sugarcoated Lives, Bitter Truths: Child Labour In Kenya's Sugarcane Industry

Linh Khanh Pham
Supervisor: Professor Jiaying Gu

August 2026

Abstract

Child Labour, Sugarcane Industry, Poverty Dynamics, Educational Deprivation, Intergenerational Justice, Kenya

This paper studies the issue of child labour through an integrated economic and philosophical lens. We combine six waves of the Kenya Demographic and Health Survey (1993–2022) to construct a birth-cohort pseudo-panel for long-run structural analysis, with the Kenya Continuous Household Survey (2020–2022) and county-level sugarcane production data for short-run behavioural analysis. We model the dynamic relationship between poverty, educational deprivation, and child labour and examine how sugarcane-sector shocks interact with this system. The long-run results imply persistence across the three outcomes, while the dominant eigenstructure does not fully support a simple self-reinforcing cycle, instead revealing a more complex pattern of co-movement. In the short-run analysis, a shift-share instrumental-variable strategy does not provide sufficient evidence to identify a causal effect of sugarcane exposure: restricting the analysis to the 15 counties with genuine sugarcane variation produces weak instruments and unbounded Anderson–Rubin confidence sets.
Philosophically, the paper argues that child labour cannot be adequately justified as a voluntary choice. Structural deprivation and adaptive preferences undermine the conditions for meaningful agency, thereby violating both rights-based and capability-based normative standards.

1 Introduction

An estimated 8.5% of Kenyan children (1.3 million) are engaged in child labour, with over 96% employed in agriculture (UNICEF, 2021; U.S. Department of Labour, 2024). Quantitative evidence from Kenya indicates that household poverty significantly increases children's likelihood

of being out of school, which can be attributed to the fact that they have to engage in child labour to provide for their families (Kenya National Bureau of Statistics and UNICEF Kenya, 2017).

These dynamics are particularly pronounced in the sugarcane industry, which faces the "Salient issue": Due to volatile prices and increasing costs, farmers often do not cover their cost of production, and thus do not earn a living income, leading to negative coping mechanisms, including relying on hazardous child labour (Fairtrade International, 2024). Under such situations, economic shocks interact with household constraints in ways that potentially reinforce persistent cycles of Poverty - Child Labour - Educational Deprivation.

Unfortunately, beyond financial constraints, social attitudes further worsen this problem: "Many parents hardly understand the importance of education and see no reason for taking their children to school." (Muriuki, 2018). This parental mindset strongly shapes children's perceptions, leading them to accept labour as a mandatory responsibility rather than question it (Abdullah et al., 2022).

2 Literature Review

2.1 Related economic models/arguments

Basu and Van (1998) develop one of the most influential theoretical models of child labour, in which households and firms each solve optimisation problems. Households decide whether to send children to work or to school, while firms determine their demand for adult versus child labour. The model yields multiple equilibria, including a "bad equilibrium" characterised by high child labour and low wages, and a "good equilibrium" in which child labour is eliminated. This framework highlights child labour as a manifestation of a poverty trap. Based on this foundation, Baland and Robinson (2000) focus on the welfare analysis of child labour and come to the conclusion that child labour is indeed socially inefficient as adult parents would fail to fully internalize its negative effects, given that bequests are zero or the capital markets are under imperfect conditions. They further show that, under certain conditions, a straightforward ban on child labour can lead to Pareto improvements. Both studies adopt a microeconomic approach, modelling child labour outcomes through household optimization behaviour and firms' labour demand decisions.

A complementary strand of literature emphasises the dynamic consequences of child labour. Hindman and Smith (1999) demonstrate, within an economic framework rooted in Adam Smith's theory of national wealth, that child labour is inefficient in both the short and long run. Even though it could increase the proportion of the productive labour in the short-run, this gain is outweighed by its long-term negative effects on human capital formation, thereby contributing to the intergenerational persistence of poverty. Another key finding is that industrialisation typically involves an initial reliance on child labour followed by its gradual decline and eventual abandonment, a pattern that supports a universalist human rights perspective rather than a culturally relativist one. On the other hand, Augeraud-Véron & Fabre (2004) propose a theoretical approach to ex-

plain the relationship of Education, Poverty and Child Labour when taking into account an OLG model with parental educational choices, which is presented as a utility maximization problem.

More recent empirical work focuses on the role of income volatility and household constraints. Beegle, Dehejia, and Gatti (2006) examine how crop shocks transmit into household income fluctuations, which in turn lead to increases in child labour. They highlight two key coping mechanisms, buffer stocks and borrowing, as channels through which households respond to income volatility. The analysis is based on data from a household panel survey in Tanzania. While this literature provides important insights into short-run coping behaviour, it primarily explains child labour through liquidity constraints and income smoothing mechanisms. By contrast, rather than treating child labour decisions as static responses to income volatility, this study models the interaction between crop shocks, schooling decisions, and human capital accumulation as a dynamic system, thereby highlighting the long-run consequences.

2.2 Existing philosophical approaches

Child labour has long been a widely-debated topics in philosophy. Daniel and Patrick (2021) examine this problem from a dual prism of Utilitarianism and Deontology. Their central finding is that these ethical theories may yield conflicting conclusions: while Deontology never considers child labour to be ethical, Utilitarianism could be ethically justifiable as long as the beneficiaries of the labour quantitatively outweigh children working or suffering.

Satz (2010), in the discussion of Child Labour from a Normative Perspective, begins by defining who should be considered to be "a child": a person who is not yet fully developed, but rather developing the cognitive, moral and affective capacities. Because of this developmental stage, children require adults, whether parents or other guardians, to act on their behalf and bear special responsibilities toward them. Building on this definition, the paper then argues why child labour should be considered to be problematic due to children's limited agency, their heightened vulnerability, and the significant harm it can cause both to the individual child and to society as a whole.

From a more conceptually broad perspective, micro-founded normative theories such as the Capability Approach, originally developed by Amartya Sen and further elaborated by Martha Nussbaum, shift the evaluative focus away from aggregate economic indicators such as GDP per capita to the capability to live a good life. Sen's framework emphasises the expansion of substantive freedoms, while Nussbaum advances a more explicit account grounded in central human capabilities and functionings. A key contribution of this approach is the concept of "conversion factors," which describe how the ability of individuals to transform resources and public goods into actual functionings depends on personal, social, and environmental conditions.

The philosophical analysis in this paper primarily operates at a macro-theoretical level, drawing in particular on Distributive Justice and Social Reproduction Theory as relevant frameworks for analysing the nature of child labour. Intergenerational Justice concerns about how moral

duties and principles of justice should be allocated across present, past, and future generations, particularly in determining what obligations current generations owe to others across time, and how the significance of intergenerational injustices can be meaningfully assessed. Social Reproduction Theory, in turn, emphasizes the structural processes through which inequalities are perpetuated, with four types of capital: economic, cultural, human, and social capital. In the context of child labour, this paper focuses primarily on economic and human capital, as these most directly capture the mechanisms presented in the economics model. This would help in the interpretation that child labour is not only an ethical issue concerning individual welfare, but also a dynamically reproduced structural condition embedded within intertemporal systems of inequality.

3 Methodology

3.1 Overview and Two-Phase Analytical Strategy

3.1.1 Data Panel

The empirical strategy is conducted using two layers, motivated by the fact that poverty/welfare, educational attainment and child labour are slow-moving structural states, whereas sugarcane production shocks operate through short-term changes in labour demand and household financial incentives:

1. **Long-run Analysis**: For this layer, data is drawn from the Kenya Demographic and Health Survey for six years: 1993, 1998, 2003, 2008, 2014, and 2022. This phase addresses whether a structural poverty-education deprivation cycle exists over a thirty-year period. Although the intervals between survey waves are not uniform (ranging from five to eight years), each transition is interpreted as one discrete period. Accordingly, the estimated transition matrix captures the average evolution of the system from one survey wave to the next, rather than an annual rate of change.

2. **Immediate Behavioural Analysis**: This short-run layer investigates the full cycle of Poverty-Education-Child Labour, which can only be obtained from the Kenya Continuous Household Survey (2020–2022), together with annual sugarcane production and area harvested, to show the full interactive effect of the system.

3.1.2 State Variables

We track three core state variables at each discrete time period t:

1. P_t: Poverty at time t

2. L_t: Child labour at time t

3. E_t: Education level attainment at time t

The sugarcane sector enters the system through a single **exposure channel**, $\widetilde{S}_{ct}$, motivated jointly by the agricultural household framework of Singh et al. (1986), capturing income effects from sugarcane activity, and by Bhalotra and Heady (2003)'s model of local agricultural labour demand.

3.2 Conceptual Economic Model

We model the joint evolution of (P_t, E_t, L_t) as a linear system of equations. Rather than deriving household optimization from first principles, this paper adopts a reduced-form model, in order to capture empirically documented feedback relationships.

Poverty equation. Future poverty depends on current poverty (persistence), current child labour (current child labour may reinforce future poverty by reducing children's human capital accumulation), current education (higher attainment reduces poverty), and a direct sugarcane-price shock:

$$P_{t+1} = a_{PP}\, P_t + a_{LP}\, L_t + a_{EP}\, E_t + \gamma_P\, S_t + u^P_{t+1}. \tag{1}$$

Education equation. Educational attainment responds to poverty (negative effect), child labour (keeps children out of school), its own persistence, and the sugarcane sector.

The coefficient γ_E on S_t is *assumed to be zero* in the baseline specification (sugarcane affects education only indirectly through poverty and labour); this restriction will be tested:

$$E_{t+1} = a_{PE}\, P_t + a_{LE}\, L_t + a_{EE}\, E_t + \gamma_E\, S_t + u^E_{t+1}, \tag{2}$$

Child labour equation. Child labour is driven by poverty, its own lagged value, and the education level (more educated households supply less child labour). The direct effect of sugarcane prices on child labour (γ_L) is left unrestricted, since harvest seasons plausibly pull children out of school:

$$L_{t+1} = a_{PL}\, P_t + a_{LL}\, L_t + a_{EL}\, E_t + \gamma_L\, S_t + u^L_{t+1}. \tag{3}$$

Throughout Section 3.2, S_t denotes the sugarcane shock in general, aggregate form, consistent with the reduced-form character of the conceptual model. Its empirical counterpart, estimated in Section 3.4 as S_{ct}, is the county-time-specific realisation of this shock – i.e. S_t indexed to a given county c and year t – constructed from the production and area channels described in Section 3.1.2 and instrumented via the shift-share strategy below.

For notational convenience, the transition equations (1)–(3) can be written in matrix form. This representation highlights the interactions among the three state variables but is used only as a

conceptual framework.

$$X_t = \begin{bmatrix} P_t \\ E_t \\ L_t \end{bmatrix} \tag{4}$$

and

$$A = \begin{pmatrix} a_{PP} & a_{EP} & a_{LP} \\ a_{PE} & a_{EE} & a_{LE} \\ a_{PL} & a_{EL} & a_{LL} \end{pmatrix}, \qquad \Gamma = \begin{pmatrix} \gamma_P \\ \gamma_E \\ \gamma_L \end{pmatrix}, \tag{5}$$

where a_{ij} denotes the effect of state variable i at time t on state variable j at time $t+1$; that is, the first subscript identifies the explanatory (lagged) variable and the second subscript identifies the dependent (current) variable.

Then, we have the compact form of the system as:

$$X_{t+1} = AX_t + \Gamma S_t + u_{t+1}, \tag{6}$$

with

$$u_{t+1} = \begin{bmatrix} u^P_{t+1} \\ u^E_{t+1} \\ u^L_{t+1} \end{bmatrix}. \tag{7}$$

3.3 Long-run Structural Analysis (KDHS)

3.3.1 Data Construction

Due to inconsistencies in household tracking across Demographic and Health Survey (DHS) waves, where different random samples of the population are collected at multiple points in time, we construct a pseudo-panel by grouping individuals into birth-year-based cohorts, which remain invariant over time. This approach allows us to treat the average value of each variable within each cohort as an observation in a synthetic panel.

The three main variables are constructed as follows:

1. **Poverty (P):** Since the wealth index was first standardized and implemented only in the late 1990s, it is unavailable for the 1993 and 1998 survey waves. To ensure data consistency across all waves, we instead construct a proxy based on the concept of the Multidimensional Poverty Index (MPI). More specifically, we construct a pooled asset index using Principal Component Analysis (PCA). The index is estimated from a common set of household asset and housing-quality indicators that are consistently observed across all six survey waves, including ownership of electricity, radio, television, refrigerator, and bicycle, together with indicators for improved sanitation, finished flooring, and finished roofing. This pooled PCA index provides a comparable measure of relative household wealth across survey years.

2. **Education (E):** This variable can be directly extracted from the survey: Educational attainment is measured using the household head's cumulative years of schooling. Household heads are first identified from the household roster, after which the corresponding years of completed education are extracted directly from the survey. Following Kenya's 8-4-4 education system, a household head is classified as educationally deprived if fewer than eight years of schooling have been completed.

3. **Child Labour (L):** This variable cannot be directly observed from the survey and therefore requires further construction. Thus, we proxy child labour using school non-attendance among children aged 5–14 years. The construction procedure is described in the Appendix. Consequently, we acknowledge limitations when it fails to exclude other reasons for school non-attendance, including disability, illness, early marriage, or household caregiving responsibilities; however, it captures one of its principal consequences and has been widely employed when direct labour measures are unavailable.

4. **Demographic Controls**: Additional demographic controls include the household dependency ratio, measured as the number of children aged 5–14 relative to working-age household members (15–64 years).

3.3.2 Bayesian Estimation Strategy

Given the data limitations of the KDHS, not all structural parameters are point identified. As such, parameters are estimated using values reported in the literature or imposed theoretical restrictions, and estimations from the household-level panel.

To incorporate existing empirical evidence into the estimation process, this study adopts a Bayesian estimation framework. Rather than treating previous literature as separate evidence, prior information from existing studies is formally incorporated into the parameter estimation process.

For each parameter θ_i:

$$\theta_i \sim N(\mu_i, \sigma_i^2)$$

where μ_i and σ_i are obtained from previous empirical studies.

The posterior distribution is obtained as:

$$p(\theta|X) \propto p(X|\theta)p(\theta)$$

where $p(\theta)$ represents the prior distribution and $p(X|\theta)$ represents the likelihood function generated from the KDHS pseudo panel.

Table 1: Empirically Informed Prior Specifications

	Economic Interpretation	**Prior Distribution**	**Empirical Motivation**	**Prior Construction Justification**
a_{PP}	Poverty persistence over time	$N(0.72, 0.15^2)$	Place et al. (2006)	Place et al. (2006) study longitudinal poverty transition across rural agricultural sites in Kenya, with a weighted average around 75%
a_{EE}	Persistence of educational attainment	$N(0.70, 0.15^2)$	Hertz et al. (2008)	In Africa, the estimated parent–child schooling coefficients range from 0.69 to 1.03, with a regional average of approximately 0.80. Hence, we adopt a conservative weakly informative prior centered below the regional average while allowing substantial uncertainty
a_{LL}	Persistence of child labour participation	$N(0.29, 0.10^2)$	Emerson and Souza (2003)	In fully controlled Probit specifications, having father/mother who is also a child labourer increases the probability that the child themselves also work by around 0.26/0.32. We will use a simple average of these results, as Brazil's sugarcane sector shares multiple structural characteristics with Kenya, making it a reliably analogous construction for the prior
a_{EP}	Effect of poverty on future education outcome	$N(0.25, 0.15^2)$	Geda et al. (2005)	Extracted directly from the presented result
a_{PE}	Effect of education outcome on future poverty	$N(0.18, 0.09^2)$	Feisal (2023)	A weakly informative prior constructed based on empirical evidence from rural Kenya
a_{LE}	Effect of child labour on future education outcome	$N(-0.025, 0.05^2)$	Beegle et al. (2006)	Data from Tanzania. However, what is directly estimated in this paper is that a crop shock reduces the probability of school enrolment by about 20 percentage points and increases hours worked by 8.93 hours. Hence, combining their estimated increase in child labour hours with the corresponding decline in enrolment implies an adverse effect of child labour on educational outcomes
a_{EL}	Effect of education on child labour participation	$N(-0.12, 0.08^2)$	Feisal (2023)	Calibrating from the presented result
a_{PL}	Effect of poverty on child labour participation	$N(-0.03, 0.10^2)$	Feisal (2023)	Adopting the presented empirical result
a_{LP}	Direct effect of child labour on future poverty	$N(0, 0.20^2)$	Bell and Gersbach (2009)	Assumption: the effect of child labour on future welfare is mainly carried through the channel of human capital accumulation, which takes time to materialize (through reduced educational attainment and future earning capacity), rather than through a direct short-run channel conditional on current poverty and education.

Bayesian Multilevel Regression

Poverty (P) and Education (E) were modelled using Gaussian regression with an identity link, whereas Child Labour (L) was modelled using a zero-inflated beta regression with a logit link, which takes into account the bounded nature of the outcome and the presence of exact zeros. Each model included lagged outcome variables, dependency ratio, female-headed household status, and urban residence as predictors, with a random intercept for birth cohort to account for cohort-level clustering. Models were estimated using Hamiltonian Monte Carlo with four chains, 4,000 iterations per chain, and 1,000 warmup iterations.

Convergence to be determined by $\hat{R}$ and effective sample size. Model fit was evaluated using posterior predictive checks by comparing simulated and observed outcome distributions.

3.3.3 Robustness Checks

Given that Bayesian inference combines prior information with observed data, a prior sensitivity analysis is conducted to evaluate whether the posterior estimates are robust to alternative prior specifications.

For each coefficient, the baseline model is re-estimated under two alternative prior structures: a Diffuse prior (i.e., a weaker prior) and an Informative prior (i.e., a stronger prior), then compare posterior estimates. The full robustness check results are reported in Appendix B.

3.3.4 Result

The Bayesian models **converge** when $\hat{R}$ less than 0.01 and sufficient effective sample size (exceeding 1,600) across parameters are recorded.

1. In the Poverty model, prior poverty, prior education, and female head are positively correlated and are statistically significant.

2. In the Education model, it can be clearly observed that prior education is strongly correlated with educational attainment, yet negatively correlated with female head.

3. In the Child Labour model, child labour is observed to move in the opposite direction from welfare level, dependency ratio, and female head, while exhibiting a strong positive correlation with prior child labour.

Posterior predictive checks suggested that the Gaussian model for poverty was broadly consistent with the observed distribution, the education outcome appeared more complex and possibly bi-modal, and the zero-inflated beta model for child labour captured the skewness and low proportion of exact zeros reasonably well.

The full Bayesian model summaries, convergence diagnostics and posterior predictive check plots for all three outcomes are presented in Appendix A.

3.4 Short-run Behavioural Analysis (KCHS)

3.4.1 Baseline Empirical Model

$$Y_{ct} = \beta_0 + \beta_1 S_{ct} + \mu_c + \lambda_t + \varepsilon_{ct}, \tag{8}$$

where Y_{ct} denotes the outcome variable of interest (poverty, education deprivation, or child labour) in county c at time t. S_{ct} represents sugarcane exposure, μ_c denotes county fixed effects, λ_t denotes year fixed effects, and ε_{ct} is the error term.

3.4.2 Identification Strategy

Shift-share Instrumental Variables (IV)

To address the problem of endogeneity, we use a shift-share instrumental variable strategy:

$$Z_{ct} = \text{Exposure}_c \times \text{SugarPrice}_t,$$

where Exposure_c is each county's fixed, pre-period (2019) sugarcane area under cultivation, held constant prior to the sample period to satisfy the shift-share design's exogeneity requirement, and SugarPrice_t is the time-varying world sugar price. This single instrument is used throughout to identify the causal effect of sugarcane exposure, S_{ct}. We consider two alternative empirical constructions of S_{ct}:

Production channel:

$$IV = Production \times SugarPrice$$

captures the interaction between world sugar price fluctuations and county-specific sugar production intensity.

Area channel:

$$IV = AreaHarvested \times SugarPrice$$

captures the interaction between world sugar prices and land devoted to sugarcane cultivation.

The verification for this choice of IV is presented in the Results section.

3.4.3 Estimation Method

The causal effects of sugarcane exposure on poverty, education deprivation, and child labour are estimated using a two-stage least squares (2SLS) framework. The endogenous sugarcane exposure variable is instrumented using the shift-share instruments described above.

In the first stage, sugarcane exposure is predicted using the interaction between county-level sugarcane exposure and world sugar prices:

$$S_{ct} = \pi_0 + \pi_1 Z_{ct} + \mu_c + \lambda_t + \nu_{ct}, \tag{9}$$

where S_{ct} denotes sugarcane exposure in county c at time t, Z_{ct} represents the shift-share instrument, μ_c are county fixed effects, and λ_t are year fixed effects.

The second stage estimates the effect of the predicted sugarcane exposure on the outcome variables:

$$Y_{ct} = \beta_0 + \beta_1 \widehat{S}_{ct} + \mathbf{X}'_{ct}\gamma + \mu_c + \lambda_t + \varepsilon_{ct}, \tag{10}$$

where Y_{ct} represents poverty, education deprivation, or child labour, $\widehat{S}_{ct}$ is the instrumented sugarcane exposure obtained from the first stage, and $\mathbf{X}_{ct}$ includes the remaining state variables in the system.

All specifications include county and year fixed effects. County fixed effects control for time-invariant differences across counties, while year fixed effects account for common national shocks. Standard errors are clustered at the county level to allow for arbitrary correlation of errors within counties over time.

3.4.4 Results

Instrument Validity This choice of shift-share IV is valid for two reasons:

1. Exogeneity of shares (Exclusion Restriction): The identifying assumption is stated as, conditional on county and year fixed effects, world sugar price fluctuations affect poverty, educational attainment and child labour only through the channel of county-level sugarcane activities. Furthermore, world sugar prices are exogenous to county-specific shocks in Kenya as they are determined by international markets.

2. Relevance of shifts: Initially, we ran the first-stage regression using all 47 counties, where non-sugarcane counties were assigned zero sugar exposure. However, this leads to an inflation in first-stage F-statistics of above 40, which could be attributed to the fact that the variation is coming from artificial assignment of zero exposure rather than actual sugarcane variation. Thus, after re-estimating the model using only 15 active sugarcane-producing counties, the F-statistic is reported as follows:

   when using

   $$AreaHarvested \times SugarPrice$$

   $$F = 5.43$$

   when using

   $$Production \times SugarPrice$$

   $$F = 9.77$$

which is close to the conventional threshold of 10.

Given the possibility of weak instruments, inference is based on Anderson–Rubin confidence intervals, which remain valid regardless of instrument strength.

Second-stage Result Using production-based sugar exposure with the raw sugar price instrument, the results are:

1. Restricting to the 15 sugarcane counties with genuine instrument variation, production-based sugar exposure with the raw sugar price instrument gives:

Outcome	2SLS Estimate	AR 95% CI	p-value at zero
Poverty (P)	-0.065	Unbounded	0.413
Education (E_{avg})	0.189	Unbounded	0.330
Child labour (L_{agri})	0.010	Unbounded	0.929

None of the three outcomes is identified: the AR confidence set is unbounded over $[-5, 5]$ in every case, so the instrument does not pin down β_S once the 32 structural-zero counties are removed. The 2SLS point estimates above should not be interpreted as evidence of an effect in either direction.

2. The same specification estimated on the full 141-observation panel (all 47 counties, including 32 with $S = Z = 0$ in every year) gives:

Outcome	2SLS Estimate	AR 95% CI	p-value at zero
Poverty (P)	-0.021	[-0.86,0.04]	0.694
Education (E_{avg})	0.350	[0.02,1.24]	0.044
Child labour (L_{agri})	-0.131	[-1.14,-0.02]	0.040

Taken at face value, this table would suggest a positive association between sugarcane exposure and educational deprivation and a negative association with agricultural child labour, both significant under AR inference. However, this table does not survive restriction to the honest sample above and should not be read as evidence of an effect in either direction.

The correlation between the two measures among sugarcane counties is:

$$Corr(S_{production}, S_{area}) = 0.9975$$

3.5 Eigenstructure Analysis

3.5.1 Stability and the Spectral Radius

A necessary and sufficient condition for the system of linear equations to converge to a stable long-run equilibrium is that the **spectral radius** of A satisfies:

$$\rho(A) = \max_i |\lambda_i| < 1, \tag{11}$$

where $\{\lambda_i\}$ are the eigenvalues of A. If $\rho(A) < 1$, the system always returns to a stable fixed point *unless* an external force counteracts convergence.

The estimated regression coefficients on lagged poverty, education and child labour presented above are mapped into the transition matrix $\hat{A}$:

$$\hat{A} = \begin{pmatrix} 0.49 & 0.11 & -0.34 \\ 0.12 & 0.51 & -0.02 \\ -0.08 & 0.002 & 0.21 \end{pmatrix}, \tag{12}$$

Two key remarks worth flagging:
(i) since the L equation (i.e., the third row) is estimated with a logit link, in order to fit the linear transition matrix, it has been linearized appropriately.
(ii) additional covariates and random effects are not included in this transition matrix.
We proceed to find the eigenvalues for $\hat{A}$, obtaining 3 distinct eigenvalues:

$$\lambda_1 \approx 0.65, \qquad \lambda_2 \approx 0.43, \qquad \lambda_3 \approx 0.13. \tag{13}$$

Hence, the spectral radius is

$$\boxed{\rho(\hat{A}) = \max_i |\lambda_i| = 0.65 < 1.} \tag{14}$$

which implies this system is **asymptotically stable**, i.e., it converges to a unique long-run equilibrium.

3.5.2 Eigenvectors and Hypothesis Mapping

Let λ^* be the dominant eigenvalue (largest in modulus) and $\mathbf{v}^* = (v_P,\, v_E,\, v_L)^\top$ the associated eigenvector. The direction of $\mathbf{v}^*$ characterises the *mode* along which the system is most sensitive:

- If $v_P,\, v_L > 0$ and $v_E < 0$, the dominant mode corresponds to the self-reinforcing cycle: rising poverty and child labour, falling education.
- The magnitude of $\rho(A)$ then indicates whether external interventions (e.g. compulsory schooling) must overcome a near-unit-root process.

For the dominant eigenvalue obtained above,

$$\rho(\hat{A}) = \max_i |\lambda_i| = 0.65 \tag{15}$$

we have the associated eigenvector to be:

$$v_{\max} = \begin{pmatrix} 0.74 \\ 0.66 \\ -0.13 \end{pmatrix}, \tag{16}$$

This does **not** fully match the hypothesis of a self-reinforcing cycle. Instead, it suggests that, despite an improvement in welfare, the child labour rate continues to increase while the education rate declines. One possible interpretation is that an increase in labour demand draws more children into the industry, outweighing the positive effects of improved welfare on school attendance. It is important to note that this eigenvector characterises the internal transition matrix $\hat{A}$, estimated from the general KDHS pseudo-panel independent of any sugarcane-specific shock. As previously discussed, the short-run causal estimate of the sugarcane shock runs in the opposite direction, and the two results should not be read as describing the same mechanism. One plausible reading is that the broader system exhibits a structural tendency toward rising child labour that welfare gains alone do not reverse, while sugarcane activity acts as a partial, sector-specific countervailing force in the short run — perhaps because sugarcane households benefit from more stable seasonal income or more formalised labour arrangements than the average agricultural household in the KDHS sample.

4 Philosophical Analysis

4.1 Social Reproduction Theory and the Persistence of Child Labour

The concept of the maintenance and continuation of existing social relations can be dated back to *Das Kapital*, in which Marx (1867/1990) analyzed capitalism, particularly the reproduction of wage-labourers for further capital production. Putting into the context of the sugarcane industry, from the data presented in previous sections, it could be observed that the relationship between poverty, education, and child labour does not conform to a simple self-reinforcing cycle. This could be attributed to the idea that an increase in agricultural income may generate competing effects: on the one hand, higher income might reduce the need for children to work and support their education; on the other hand, an expansion in agricultural activities naturally increases the demand for labour.

As such, the firm-side mechanism should be taken into analysis. In the cost structure of a firm, labour is the variable cost, which could first be adjusted to maximize profit. In the agricultural sector, for manual work, which requires a lower level of labour, children, who can typically be hired at a lower wage than adults, become the ideal choice to maximize profit. As such, the demand for child labour keeps increasing, pushing more people out of school. This phenomenon strongly correlates with the idea of classical Marxian analysis which focuses on the wage relation to the

unwaged and undervalued labour (which includes child labour) that sustains the reproduction of labouring class.
More interestingly, child labour can also be viewed as a paradigm case of Social Reproduction Theory, which extends beyond the wage relation itself when the household itself reallocates children's time from schooling to immediate labour. This can further strengthen the cycle: firms have higher demand for "cheap" low-skilled child labour, whereas households are willing to supply it for short-run gains with the price of lower human capital and lower bargaining power in the future. Those reduced opportunities can, in turn, increase the likelihood that the next generation will face similar constraints. What is reproduced is therefore not simply labour, but a structure of limited possibilities.

4.2 Normative Arguments: Why is Child Labour Problematic?

4.2.1 International Children's Rights Framework

For the sake of consistency, we use the Convention on the Rights of the Child (United Nations, 1989) as the sufficient threshold. In Article 32:

> 1. States Parties recognize the right of the child to be protected from economic exploitation and from performing any work that is likely to be hazardous or to interfere with the child's education, or to be harmful to the child's health or physical, mental, spiritual, moral or social development.
> 2. States Parties shall take legislative, administrative, social and educational measures to ensure the implementation of the present article. To this end and having regard to the relevant provisions of other international instruments, States Parties shall in particular: (a) Provide for a minimum age or minimum ages for admission to employment; (b) Provide for appropriate regulation of the hours and conditions of employment; (c) Provide for appropriate penalties or other sanctions to ensure the effective enforcement of the present article.

The first normative force of Article 32 could be best explained by Interest Theory of Rights, which is one of the two main theories of legal rights (together with the Will Theory). This was first initiated by Jeremy Bentham, who is regarded as the father of utilitarianism, then most famously propounded by Ihering (1877/1913) (or more commonly known as Jhering). He highlights the idea that the purpose of law should be the protection of individual and social interests instead of the will of the citizens. Recently, the Interest Theory has been further developed by Raz (1984) by formulating rights as follows:

> 'x has a right' if and only if x can have rights, and other things being equal, an aspect of x's well-being (his interest) is a sufficient reason for holding some other person(s) to be under a duty.

Here, Article 32 explicitly protects children from economic exploitation and any work that may harm their health, development, or education. Under the lens of Interest Theory, the question is reduced to whether the child's interest in physical safety, healthy development, and education is weighty enough to ground a duty on states and employers not to exploit their labour. And the answer is an unarguably yes. An important remark worth flagging here is that these interests are not merely subjective preferences; rather, they represent fundamental conditions that are essential for a child to develop as an autonomous individual. As such, under Raz's formulation, children's interests in bodily integrity, healthy development, and education are sufficiently important to justify imposing correlative duties on states, employers, and other actors to refrain from economic exploitation and to enact protective labour legislation.

A possible objection can question why children have rights when they lack full autonomy, and rights are now formulated based on interest. Children are indeed viewed as a strong case for Interest Theory: their present inability to fully exercise autonomy does not weaken their rights but instead, it strengthens justification for protective duties, which are necessary for them to become a fully autonomous agent.

Another perspective worth considering is whether, instead of substituting x with the children, parents have the right to "force" their children to perform the responsibility of supporting the household financial situation, especially when the households are experiencing severe economic hardship. We believe the answer is a strict no.

Such a parental right cannot be grounded under Interest Theory because the well-being (in this case, a short-term improvement in household income) cannot justify two things: (i) the deprivation of access to education and (ii) the sacrifice of the potential for a higher level of well-being in the future. Thus, it is not sufficient to ground a right.

4.2.2 Distributive Justice

We will start the discussion by taking into account the Welfare-Based Principles. Welfare-based principles evaluate a problem based on how the distribution affects the level of people's welfare (Lamont & Favor, 2017), which can be mostly illustrated through utilitarianism (which was the foundation of the previous section), which evaluates social arrangements according to their ability to maximize aggregate welfare.

This idea has been criticized by John Rawls for not taking into account how people are distinct from each other, and he instead proposes in the Difference Principle in *A Theory of Justice* (Rawls, 1971/2009) and *Political Liberalism* (Rawls, 1993):

> 1. Each person has an equal claim to a fully adequate scheme of equal basic rights and liberties, which scheme is compatible with the same scheme for all; and in this scheme the equal political liberties, and only those liberties, are to be guaranteed their fair value.
>
> 2. Social and economic inequalities are to satisfy two conditions: (a) They are to

be attached to positions and offices open to all under conditions of fair equality of opportunity; and (b) they are to be to the greatest benefit of the least advantaged members of society.

Priority rules are set when these rules may conflict in practice, with the order being that Principle (1) is most prioritized, followed by (2a) and then (2b).
Principle (1) mainly focuses on basic liberties, which can be applied to the case of hazardous child labour when the work involves physical danger or coercion. However, a more interesting point worth noticing in Principle (1) is that each person needs to be guaranteed their fair value, which is framed by Rawls as the relatively equal capacity to exercise political liberties. For a child, this liberty should not be weighted; however, current educational deprivation can be a negation of their future "fair value": Due to limited access to education, which leads to illiteracy or a lack of civic knowledge, they do not have sufficient tools to exercise their basic political liberties.
Looking into Principle (2), which has direct relevance to child labour, it requires a scheme of fair equality of opportunity regardless of social background. When a child is forced to participate in the agricultural labour market, which directly interferes with their schooling, they are forced to let go of an opportunity to access knowledge and better career opportunities in the future (this is based on the assumption that higher education does imply higher income in the future), which no longer ensures fairness in comparison to their peers coming from a more advantaged background. Taking Principle (b) into account, the question to be asked here is whether allowing child labour benefit the least advantaged. In extremely poor households, a potential counterargument could be that, banning child labour could worsen the condition of the least advantaged because children's income contributes to household survival. Therefore, allowing limited child labour may satisfy the Difference Principle if it improves the welfare of the poorest families. However, Rawls does not justify any inequality simply because it increases short-term income, and it is the point where the priority clicks in. The institution must improve the position of the least advantaged while maintaining the priority of basic liberties and fair opportunity, as outlined in Principle (1) and (2a).

In short, child labour does not appear to generate a genuine conflict between Rawls's principles because the principles operate under a priority ordering. Hence, it could not be justified under a welfare-oriented Rawlsian framework.

4.3 Discussion

4.3.1 Contradiction or Agreement?

We acknowledge that by adopting two contrasting philosophical frameworks, we are facing the risks of inconsistency. However, similar to utilitarianism and deontology, even though they are often presented as opposing ethical theories, they frequently agree on what people should morally

do in many real-world cases. We believe that Child Labour can sit in that common ground. Consequently, both the aforementioned normative framework and legislative reference come to the same conclusion: Child Labour seriously violates the child's basic rights when it undermines the fundamental interests that justify the existence of those rights in the first place.

4.3.2 Who should bear the Responsibility?

To further develop the problem, Child Labour becomes an interesting case as it contains a self-generated problem (as previously explained by Social Reproduction Theory), which forces an external intervention. Then a further problem to be addressed, we believe, should not be "Who, either parents or children, should bear the responsibility?", but it should be reframed to be "Given the unjust nature of child labour, what should be the exit for the families who have been trapped in the cycle intergenerationally, and who can make it happen?"

We will provide a short argument why it should bear the responsibility on other agents, rather than a pure household problem. First, many households experiencing child labour lack the ability to escape this cycle independently. Empirical evidence and qualitative observations suggest that poverty, limited educational access, and restricted economic opportunities significantly constrain household choices. In such circumstances, child labour may represent not a freely chosen preference, but a survival strategy under structural limitations, i.e., the household does not have actual meaningful alternatives. Second, welfare distribution cannot be understood purely as the outcome of household decisions when those decisions are made within unequal social conditions – children's opportunities are shaped by broader institutional arrangements. From a Rawlsian perspective, institutions have a responsibility to provide fair equality of opportunity by ensuring that individuals' life prospects are not determined solely by an arbitrary birth lottery. Therefore, addressing child labour requires institutional intervention aimed not only at reducing immediate harm but also at breaking the intergenerational mechanisms that reproduce disadvantage.

5 Limitations and Future Directions

Several limitations should be acknowledged, which also provide directions for further research.

Unequal survey intervals. Consecutive KDHS waves are separated by unequal calendar intervals (5–8 years). The estimated transition matrix $\hat{A}$ therefore represents the average transition between survey waves rather than a constant annual transition process. While this assumption is common in analyses based on repeated household surveys, the estimated dynamics should be interpreted at the survey-wave level rather than as annual transition parameters.

Non-equivalent measurement of child labour across layers. Because the KDHS does not directly observe child labour, L is proxied in Layer 1 by school non-attendance among children aged 5–14. This proxy captures non-attendance driven by disability, illness, early marriage, or caregiving responsibilities alongside labour itself, which likely attenuates or biases the estimated

persistence and cross-dimensional coefficients in $\hat{A}$ in an unknown direction. Layer 2, by contrast, uses a directly measured agricultural child labour rate. Because the transition matrix $\hat{A}$ and the shock vectors Γ are therefore estimated on two non-identical operationalisations of L, the compact system $X_{t+1} = AX_t + \Gamma S_t + u_{t+1}$ should be interpreted as an approximate layering of a short-run labour-demand shock onto a long-run structural process, rather than a single coherently estimated dynamical system.

Scope of the philosophical framework. Further analysis could be conducted by incorporating additional theoretical perspectives and examining how institutional, social, and ethical factors shape the relationship between poverty, education deprivation, and child labour.

The results also provide suggestion for further investigation.

The counterintuitive sugarcane-labour result. Contrary to the paper's opening hypothesis, the estimations suggest sugarcane exposure is associated with a *reduction* in agricultural child labour, while education deprivation moves in the opposite direction from what a simple poverty-relief channel would predict. The dominant eigenvector of $\hat{A}$ likewise does not fully match the hypothesised self-reinforcing cycle.

This unexpected finding motivates further philosophical and empirical investigation into the mechanisms underlying the relationship between agricultural development, household welfare, education outcomes, and child labour.

6 Conclusion

This paper set out to connect two questions that are usually asked separately: whether child labour in Kenya's sugarcane industry can be explained as a self-reinforcing dynamic system, and whether it can be normatively justified as a voluntary household choice made under hardship.

The quantitative analysis through the economic model suggests the first question deserves a qualified yes: the estimated transition matrix is asymptotically stable ($\rho(\hat{A}) < 1$), meaning poverty, education deprivation, and child labour do not spiral without bound, but the system's dominant mode of adjustment is slow, and its associated eigenvector indicates that welfare gains alone are not sufficient to unwind the cycle – child labour continues to move with, rather than against, the labour-demand channel in the short-run analysis.

The philosophical analysis argues that this matters normatively, not just descriptively. Whatever their households' income gains, children whose labour is drawn on by expanding cultivation are not exercising a meaningful choice in the sense required to justify their work as voluntary.

In conclusion, the dynamic nature of the system highlights an important policy implication: if poverty, education deprivation, and child labour are sustained, the system is unlikely to correct itself through gradual internal adjustment alone. External interventions are therefore required to disrupt these reinforcing pathways. Policies that simultaneously reduce household vulnerability,

expand educational access, and alter the incentives surrounding child labour are necessary to shift the system toward a more sustainable trajectory.

Acknowledgements

I would like to express my sincere gratitude to the **Laidlaw Scholars Foundation** for their generous financial support, which made this research possible, and to the **U of T Laidlaw Scholars Programme** for giving me the opportunity to undertake this research and develop as a researcher.

I am especially grateful to my research supervisor, **Professor Jiaying Gu**, for believing in this project from its early stages and for her continued guidance, thoughtful feedback, and encouragement throughout the research process.

I would also like to thank **Tam Le, Minh Thai Duong, and Omer Ali** for their valuable advice along the way. Their insights and encouragement were greatly appreciated throughout the development of this project.

And I am deeply grateful to my beloved family and friends for their unwavering mental support through the moments of uncertainty and challenge over this summer.

Lastly, I acknowledge the use of AI-based language tools (Claude and ChatGPT) for assistance with grammar checking and wording refinement during the preparation of this research report. All ideas, analysis, and conclusions presented in this work are entirely my own.

References

[1] Abdullah, A., Huynh, I., Emery, C. R., & Jordan, L. P. (2022). *Social norms and family child labor: A systematic literature review*. International Journal of Environmental Research and Public Health, 19(7), 4082. https://doi.org/10.3390/ijerph19074082

[2] Augeraud-Véron, E., & Fabre, A. (2004). *Education, poverty and child labour*. Econometric Society 2004 Far Eastern Meetings, Paper No. 738. http://repec.org/esFEAM04/up.9133.1080753714.pdf

[3] Bhalotra, S., & Heady, C. (2003). *Child Farm Labor: The Wealth Paradox*. The World Bank Economic Review, 17(2), 197–227. http://www.jstor.org/stable/3990136

[4] Baland, J., & Robinson, J. A. (2000). *Is Child Labor Inefficient?* Journal of Political Economy, 108(4), 663–679. https://doi.org/10.1086/316097

[5] Basu, K., & Van, P. H. (1998). *The Economics of Child Labor*. The American Economic Review, 88(3), 412–427. http://www.jstor.org/stable/116842

[6] Beegle, K., Dehejia, R. H., & Gatti, R. (2006). *Child labor and agricultural shocks.* Journal of Development Economics, 81(1), 80–96. https://doi.org/10.1016/j.jdeveco.2005.05.003

[7] Bell, C., & Gersbach, H. (2009). *Child Labor and the Education of a Society.* Macroeconomic Dynamics, 13(2), 220–249. https://doi.org/10.1017/S1365100508080036

[8] U.S. Department of Labor. (2024). *Findings on the worst forms of child labor - Kenya.* https://www.dol.gov/agencies/ilab/resources/reports/child-labor/kenya

[9] Emerson, P. M., & Souza, A. P. (2003). *Is There a Child Labor Trap? Intergenerational Persistence of Child Labor in Brazil.* Economic Development and Cultural Change, 51(2), 375–398. https://doi.org/10.1086/346003

[10] Feisal, A. I. (2023). *An Empirical Investigation of the Effect of Cash Transfers on Child Labour and Schooling in Rural Kenya.* University of Nairobi. https://erepository.uonbi.ac.ke/items/b68644cc-9dd0-4bde-9b98-06de21101301

[11] Geda, A., de Jong, N., Kimenyi, M. S., & Mwabu, G. (2005). *Determinants of Poverty in Kenya: A Household Level Analysis.* African Development Review, 17(2), 353–374.

[12] Hertz, T., Jayasundera, T., Piraino, P., Selcuk, S., Smith, N., & Verashchagina, A. (2008). *The Inheritance of Educational Inequality: International Comparisons and Fifty-Year Trends.* Advances in Economic Analysis & Policy, 7(2), Article 10. https://doi.org/10.2202/1935-1682.1775

[13] Hindman, H. D., & Smith, C. G. (1999). *Cross-Cultural Ethics and the Child Labor Problem.* Journal of Business Ethics, 19(1), 21–33. http://www.jstor.org/stable/25074071

[14] Jhering, Rudolf von. (1913). *Law as a Means to an End.* (I. Husik, Trans.). Vols. 1–2. Boston, MA: Boston Book Company. (Original work published 1877–1883)

[15] Kenya National Bureau of Statistics & UNICEF Kenya. (2017). *Child poverty in Kenya: A multidimensional approach (Report).* https://www.unicef.org/esa/media/1356/file/UNICEF-Kenya-2017-Child-Poverty.pdf

[16] Lamont, J., & Favor, C. (2017). *Distributive Justice.* In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2017/entries/justice-distributive/

[17] Marx, K. (1990). *Capital: A critique of political economy* (Vol. 1) (B. Fowkes, Trans.). Penguin Classics. (Original work published 1867)

[18] Place, F., Barrett, C. B., Marenya, P. P., McPeak, J., Minten, B., Murithi, F., Kosura, W. O., Randrianarisoa, J., Rasambainarivo, J., & Wangila, J. (2006). *Welfare Dynamics in*

Rural Kenya and Madagascar. The Journal of Development Studies, 42(2), 248–277. https://doi.org/10.1080/00220380500405394

[19] Rawls, J. (2009). *A Theory of Justice: Original Edition*. Cambridge, MA and London, England: Harvard University Press. (Original work published 1971). https://doi.org/10.4159/9780674042605.

[20] Rawls, J. (1993). *Political Liberalism*. New York, NY: Columbia University Press.

[21] Raz, J. (1984). *On the Nature of Rights*. *Mind*, 93(370), 194–214. http://www.jstor.org/stable/2254002

[22] Singh, Inderjit; Squire, Lyn; Strauss, John [editors]. (1986). *Agricultural household models: extensions, applications, and policy*. Baltimore, MD: The Johns Hopkins University Press. http://documents.worldbank.org/curated/en/621291468739297175

[23] United Nations. (1989). *Convention on the Rights of the Child*. United Nations Treaty Series, 1577, 3. Depositary notifications C.N.147.1993.TREATIES-5 (15 May 1993) [amendments to Article 43(2)] and C.N.322.1995.TREATIES-7 (7 November 1995) [amendment to Article 43(2)].

[24] UNICEF. (2021). *Let's end child labor in Kenya*. https://www.unicef.org/kenya/stories/lets-end-child-labor-kenya

A Appendix A

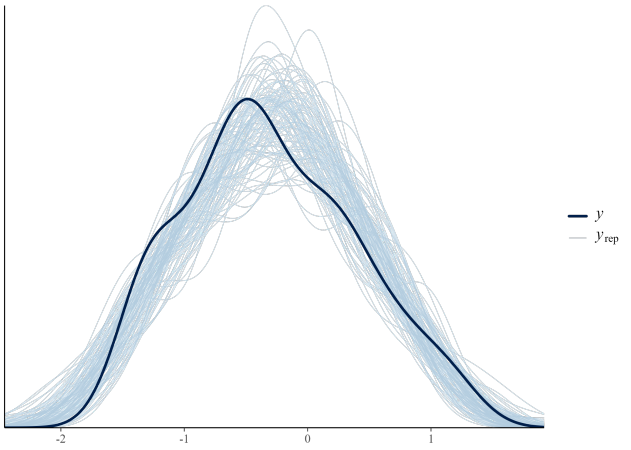

(a) Poverty (P) - density overlay

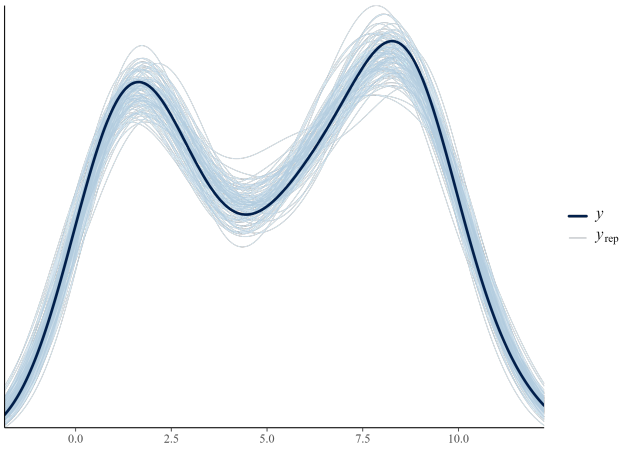

(b) Education (E) - density overlay

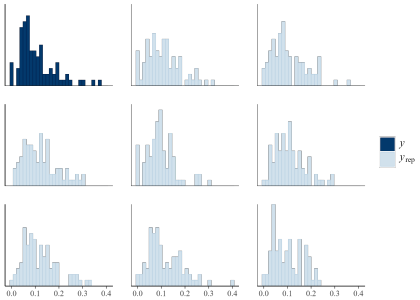

(c) Child labour (L) - histogram

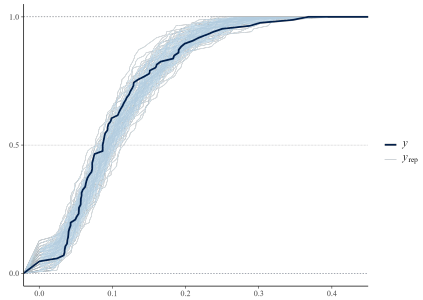

(d) Child labour (L) - ECDF overlay

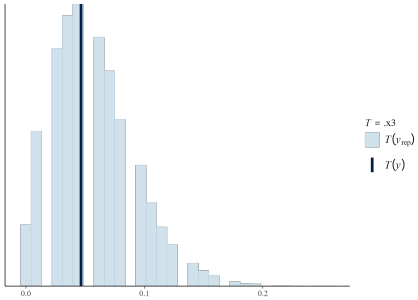

(e) Child labour (L) - zero-mass check

Figure 1: Posterior predictive checks for all three fitted models. Panels (a) and (b) show the density-overlay comparison of observed vs. model-implied draws for the Gaussian poverty and education equations. Panels (c)-(e) show three complementary checks for the zero-inflated Beta child labour equation, given its bounded outcome with a point mass at zero: a histogram comparison, an empirical CDF overlay, and a direct check of the model-implied proportion of zeros against the observed value.

Table 2: Summary Statistics - KDHS

Variable	**Obs.**	**Mean**	**Std. Dev.**	**Min**	**Max**
Household head age	108,259	44.50	16.10	9.00	99.00
Female-headed household	108,288	0.339	0.473	0.000	1.000
Household head years of schooling	107,716	7.14	4.97	0.00	25.00
Household head has no formal education	107,716	0.495	0.500	0.000	1.000
Urban household	108,289	0.342	0.474	0.000	1.000
Household size	108,289	4.18	2.53	0.00	30.00
Working-age household members	108,289	2.21	1.36	0.00	15.00
Dependency ratio	103,486	0.624	0.804	0.000	9.000
Household asset index (PCA)	107,916	0.000	1.750	-2.36	4.06
Number of children aged 5–14	108,289	1.27	1.43	0.00	15.00
Number of out-of-school children	108,289	0.176	0.568	0.000	13.00
Household reports child labour	108,289	0.118	0.322	0.000	1.000
Share of children out of school	62,622	0.125	0.283	0.000	1.000

Notes: This table reports household-level summary statistics from the harmonized Kenya Demographic and Health Surveys (KDHS) covering the 1993, 1998, 2003, 2008, 2014, and 2022 survey waves. The household asset index is constructed using principal components analysis (PCA) on harmonized household assets and dwelling characteristics and is standardized to have mean zero and unit variance. Binary variables take the value one if the stated condition holds. The share of children out of school is calculated only for households with at least one child aged 5–14.

Table 3: Summary Statistics - KCHS

Variable	**Obs.**	**Mean**	**Std. Dev.**	**Min**	**Max**
County identifier	141	24	14	1	47
Survey year	141	2021	1	2020	2022
School non-attendance rate	141	0.09	0.13	0.00	0.54
Child labour rate	141	0.08	0.08	0.00	0.49
Agricultural child labour rate	141	0.07	0.08	0.00	0.47
Educational attainment	141	5.77	0.42	4.75	7.10
Education expenditure	94	7596.96	4303.91	1568.06	24080.09
Sample size (children)	141	441	147	173	820
Valid obs (attendance)	141	441	147	173	820
Valid obs (child labour)	141	417	149	156	815
Valid obs (education)	141	387	135	161	779
Valid obs (expenditure)	141	217	181	0	588
Mean consumption	141	7098.25	2127.75	3305.79	17442.82
Poverty incidence	141	0.42	0.15	0.14	0.84
Welfare index	141	0.58	0.15	0.16	0.86
Sugar output (tonnes)	141	14897.75	32990.45	0.00	188756.00
Harvested area (ha)	141	2134.24	4744.33	0.00	27952.00
Cultivated area (ha)	141	4723.60	10409.89	0.00	57466.00
Cane deliveries	141	165033.98	365186.99	0.00	2085241.00
Sugarcane indicator	141	0.319	0.468	0.000	1.000
Sugar production index	141	0.00	0.56	-1.06	3.22
Normalized harvested area index	141	0.09	0.19	0.00	0.89

Table 4: Posterior Estimates for the Three Bayesian Panel Regressions

	Poverty (P) *Gaussian*	**Education** (E) *Gaussian*	**Child Labour** (L) *Zero-Inflated Beta*
Intercept	-1.13 [-1.68, -0.59]*	2.98 [1.81, 4.09]*	-1.86 [-2.68, -1.05]*
P_{lag}	**0.49** [0.35, 0.63]*	0.12 [-0.02, 0.26]	**-0.68** [-0.91, -0.46]*
E_{lag}	**0.11** [0.05, 0.18]*	**0.51** [0.40, 0.63]*	0.02 [-0.06, 0.10]
L_{lag}	-0.34 [-0.69, 0.01]	-0.02 [-0.12, 0.08]	**1.82** [1.14, 2.47]*
Dependency ratio	-0.23 [-0.47, 0.01]	-0.12 [-0.48, 0.23]	**-0.78** [-1.20, -0.38]*
Female head	**1.58** [0.70, 2.44]*	**-1.31** [-2.41, -0.19]*	**-1.63** [-2.86, -0.40]*
Urban	0.71 [-0.10, 1.47]	1.35 [-0.02, 2.70]	-0.22 [-1.27, 0.83]
Hyperparameters			
sd(Intercept) [cohort]	0.11 [0.00, 0.26]	1.90 [1.32, 2.71]	0.10 [0.00, 0.26]
σ (residual)	0.29 [0.24, 0.35]	0.40 [0.34, 0.49]	–
ϕ (precision)	–	–	73.02 [49.39, 101.50]
zi (zero-inflation)	–	–	0.06 [0.02, 0.11]
Observations	86	86	86
Cohorts (levels)	21	21	21
Max $\hat{R}$	< 1.01	< 1.01	< 1.01

* 95% confidence interval excludes zero. Point estimates are posterior means; brackets show the 95% confidence interval. P_{lag}, E_{lag}, L_{lag} denote the previous-wave (lagged) values of poverty (asset index, higher = wealthier), education (mean years of schooling), and child labour (out-of-school proportion) at the cohort level. All models include a cohort-level random intercept (1 — birth_cohort) and were estimated via Hamiltonian Monte Carlo (NUTS) with 4 chains, 4,000 iterations, 1,000 warmup, using literature-informed priors.

B Appendix B

To assess the sensitivity of the posterior estimates to prior strength, two alternative prior specifications are considered. The *Diffuse Prior* uses a relatively large standard deviation ($\sigma = 1.5$), rendering the prior essentially weakly informative so that inference is driven primarily by the observed data. In contrast, the *Strong Prior* adopts a very small standard deviation ($\sigma = 0.005$), concentrating nearly all prior mass around the literature-based mean and thus approximating a near-deterministic prior. These two specifications are intentionally chosen to bracket the baseline prior, spanning the spectrum from minimal prior influence to an almost complete reliance on prior information. If the resulting posterior estimates remain stable across these extremes, the findings can be interpreted as robust to the choice of prior specification.

Table 5: Posterior estimate and 95% confidence interval for each lagged/cross-lag coefficient under the baseline, diffuse ($SD = 1.5$), and informative ($SD = 0.005$) prior specifications.

Coefficient	Diffuse (SD = 1.5)	Baseline	Informative (SD = 0.005)
Poverty equation			
P_{t-1}	0.275 [0.107, 0.440]	0.476 [0.332, 0.625]	0.718 [0.709, 0.728]
E_{t-1}	0.142 [0.084, 0.204]	0.113 [0.051, 0.179]	0.247 [0.237, 0.257]
L_{t-1}	−1.580 [−2.230, −0.904]	−0.351 [−0.700, −0.003]	0.000 [−0.010, 0.010]
Education equation			
E_{t-1}	0.479 [0.344, 0.625]	0.510 [0.384, 0.642]	0.699 [0.689, 0.709]
P_{t-1}	0.166 [−0.111, 0.437]	0.099 [−0.104, 0.306]	0.179 [0.170, 0.189]
L_{t-1}	0.658 [−0.340, 1.640]	−0.066 [−0.416, 0.293]	−0.025 [−0.035, −0.015]
Child labour equation (logit scale)			
L_{t-1}	0.377 [0.259, 0.496]	0.377 [0.264, 0.492]	0.290 [0.280, 0.299]
P_{t-1}	−0.042 [−0.071, −0.014]	−0.043 [−0.071, −0.015]	−0.025 [−0.035, −0.016]
E_{t-1}	−0.005 [−0.016, 0.005]	−0.005 [−0.015, 0.005]	−0.101 [−0.112, −0.091]

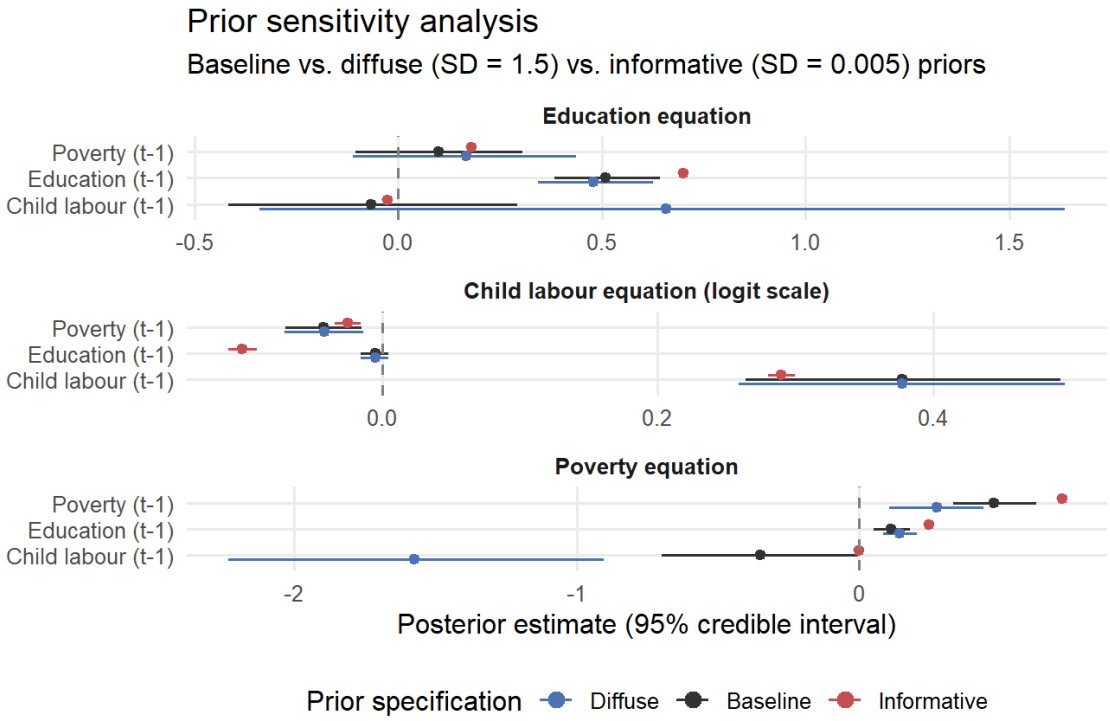

Figure 2: Prior sensitivity analysis. Posterior estimates (points) and 95% confidence intervals (bars) for the lagged and cross-lag coefficients in the poverty, education, and child labour equations, under the baseline, diffuse ($SD = 1.5$), and informative ($SD = 0.005$) prior specifications. Coefficients in the child labour equation are on the logit scale.