Institute for Economic Studies, Keio University

Keio-IES Discussion Paper Series

The Impact of an Unconditional Cash Transfer Programme on Intra-Household Inequality Among Siblings: Experimental Evidence from Zambia

Shen Shen

1 August, 2026
DP2026-017
https://ies.keio.ac.jp/en/publications/27917/

Keio University

CALAMVS GLADIO FORTIOR
1858

Institute for Economic Studies, Keio University
2-15-45 Mita, Minato-ku, Tokyo 108-8345, Japan
ies-office-group@keio.jp
1 August, 2026

The Impact of an Unconditional Cash Transfer Programme on Intra-Household Inequality Among Siblings: Experimental Evidence from Zambia

Shen Shen

Keio-IES DP2026-017

1 August, 2026

JEL Classification: D13, I38, J13, O15

Keywords: Intra-household inequality; Unconditional cash transfers; Child Investment; Zambia

Abstract

Early childhood investment, or the lack thereof, profoundly affects later-life outcomes, but whether social welfare protection aimed at encouraging parental investment mitigates or exacerbates inequitable distributive patterns within families remains elusive. This paper examines the response of intra-household inequality among siblings to a cluster randomised unconditional cash transfer programme implemented in some of Zambia's poorest communities. We find that the transfers led to compensatory investment in higher birth-order children, particularly by meeting their material needs, although we do not observe a consistent pattern of internal restructuring in health and nutrition or education. These findings are robust across a spectrum of subsample analyses. We also provide tentative evidence affirming the conjecture that the transfers had raised parents' aversion to inequality after three years of programme exposure, while marginally reinforcing their preference for younger and female children. The literature regards birth-order gradients as context-dependent; therefore, our findings may not be applicable to other settings.

Shen Shen

Graduate School of Economics, Keio University

s.shen.4@keio.jp

Acknowledgement: I am deeply grateful to Professor Hiroyuki Yamada for his unwavering guidance and support over the past two years. He introduced me to the theoretical and empirical fundamentals most relevant to this work and provided invaluable feedback at every stage of this project. I also wish to thank Professor Hideo Akabayashi, whose seminar afforded me numerous opportunities to present and refine this work. All errors are my own.

The Impact of an Unconditional Cash Transfer Programme on Intra-Household Inequality Among Siblings:

Experimental Evidence from Zambia

Shen Shen

Graduate School of Economics
Keio University
s.shen.4@keio.jp

July 2026

Abstract

Early childhood investment, or the lack thereof, profoundly affects later-life outcomes, but whether social welfare protection aimed at encouraging parental investment mitigates or exacerbates inequitable distributive patterns within families remains elusive. This paper examines the response of intra-household inequality among siblings to a cluster-randomised unconditional cash transfer programme implemented in some of Zambia's poorest communities. We find that the transfers led to compensatory investment in higher birth-order children, particularly by meeting their material needs, although we do not observe a consistent pattern of internal restructuring in health and nutrition or education. These findings are robust across a spectrum of subsample analyses. We also provide tentative evidence affirming the conjecture that the transfers had raised parents' aversion to inequality after three years of programme exposure, while marginally reinforcing their preference for younger and female children. The literature regards birth-order gradients as context-dependent; therefore, our findings may not be applicable to other settings.

Keywords: Intra-household inequality; Unconditional cash transfers; Resource allocation; Child Investment; Sibling inequality; Zambia

JEL Classification: D13; I38; J13; O15

1. Introduction

Aggregate distributional accounts of income, consumption, and welfare typically rely on information collected at the household level, thereby disregarding inequality concealed within. Questions such as why some children in a family are well nourished and sent to school while others are not were thus left unanswered. Following the inception of collective models of family behaviour, the growing availability of micro data, and novel applications grounded in long-established economic principles, intra-household inequality has come to be recognised as an important dimension of welfare with significant implications for social protection systems.

In the context of parental investment in children, the short- and long-term ramifications of intra-household inequality are magnified. Indeed, given the adverse interaction between contextual sociocultural norms and who parents prioritise in developing-country settings, a patterned distribution of resources that systematically discriminates against certain child demographic groups during early childhood can severely degrade human capital accumulation among such children, entrenching them in irreversible deficits at a stage of life when formation of motor and cognitive skills is key. The sibling gap then widens over the life cycle and ultimately proliferates inequalities along these demographic lines in broader society.

Still, the intra-household redistributive impacts of cash transfer initiatives continue to elude policymakers, as rarely are their design features configured to tap into intra-household dynamics to enhance cost-effectiveness beyond designating women as primary recipients, and rarely do their impact evaluations include efforts to study changes in intra-household inequality.

Studies across diverse contexts have shown that, despite being cheaper and less paternalistic, unconditional cash transfers (UCTs) are just as effective as conditional cash transfers (CCTs) at improving a wide array of micro, meso, and macro outcomes, leading them to be increasingly championed over their conditional counterparts. However, it remains imperative to further inquire whether there are perverse effects of UCTs on intra-household inequality, since, after all, parents face the trade-off between efficiency and equality, and there is little evidence on this front.

To help address these gaps in the current body of literature, we investigate how a UCT scheme might affect recipient households' allocation of resources using data from the Child Grant Programme (CGP), a cluster-randomised UCT programme that followed young, impoverished families in Zambia for four years from 2010 to 2014. The absence of conditionality means that beneficiaries of the CGP were not stipulated to alter their behaviour in any way, such as channelling additional resources to disadvantaged members.

We find that, at the aggregate level, four years of transfers led to little detectable impact on child health and nutritional outcomes, while substantially improving material well-being and modestly increasing enrolment. To assess the impact on intra-household inequality, we begin by decomposing variance between and within households, revealing that the aggregate gains were accompanied by reductions in intra-household inequality in absolute terms but increases in relative terms.

Our core method of identification exploits variation within household-year cells to estimate the programme impact on birth-order gradients. Estimation results show that the CGP disproportionately benefitted later-born children in terms of access to material goods, compensating for their comparatively worse material well-being at baseline. In contrast, we do not observe a consistent pattern of intra-household redistribution in the domain of health and nutrition and the domain of education. Namely, among non-recipients, the health and nutrition gap between the firstborn and the fourth or higher born widened over time, and the CGP offset this divergence; the CGP also appeared to have reinforced pre-existing favourable schooling investment in the firstborn relative to the second born.

Subsample analyses suggest that these effects are shaped by the composition of and the constraints faced by the household. While birth-order effects are broadly similar across male and female sibling groups, further analyses yield some evidence for girl preference, in the form of female-biased fertility stopping and better schooling outcomes for female offspring when they have an elder brother rather than an elder sister. We also find that the selection of allocative schemes for different outcome domains is heavily shaped by the availability of resources. Among the most resource-constrained households, the CGP elicited pronounced health and nutrition gains for second-born children, muted any form of improvement in material well-being for the later born, and resulted in perverse schooling losses for fourth or higher borns. These estimates, when interpreted in context, imply at best the preservation of and at worst the exacerbation of sibling inequality. This stands in sharp contrast to our baseline estimates, corroborating the notion that allocative schemes are state-contingent and that the most resource-constrained households are more likely to pursue efficiency maximisation. Results for large households and cohort samples do not markedly deviate from the main findings. An alternative ranking method based on baseline outcomes suggests that the compensatory effects induced by the CGP may not simply be driven by parents targeting children who were initially worse off, but rather by structured changes in allocative behaviour that differentially affected offspring of various birth orders.

To understand the drivers of those results, we augment the Behrman-Pollak-Taubman model and estimate the programme impact on two structural parameters of interest – inequality aversion and child preference. The CGP is associated with greater aversion to inequality at 36 months but lower aversion at 48 months, while the effect on child preference, although small, aligns with heightened favouritism of female and higher birth-order offspring. Households in which the primary female caretaker was the sole individual making expenditure decisions on child health and nutrition exhibited more equitable distributions. These structural estimates are best interpreted as correlational rather than causal as our measure of child development is prone both to measurement error, since it is observed and reported by parents, and to endogeneity, since it is determined simultaneously with dyadic outcomes.

The setting of our study is unusual in that male offspring are not superior by a noticeable margin in any outcome domain, a pattern that is otherwise regular in developing countries. Considering that the literature has widely shown that gender preferences are shaped by contextual circumstances, our findings are unlikely to be generalisable to other settings. Regardless, this paper marks one of the rare efforts at identifying how poverty transfers affect the distribution of resources among children.

The rest of the paper proceeds as follows. Section 2 summarises the related literature on what is known about household behaviour. Section 3 describes the study sample and the construction of key dependent variables. Section 4 briefly outlines the mechanisms of action that underlie our hypotheses on the intra-household impact of the CGP. Section 5 details the empirical strategies. Section 6 covers our results. Section 7 attempts to untangle the channels that drive our results. Section 8 concludes the paper.

2. Literature Review

Beyond mutual gains from division of labour and risk sharing, opportunities to produce internally and consume household public goods provide ample incentive for individuals to form a household (Chiappori & Meghir, 2015). Before there was an abundance of micro data, concessions often had to be made to enable the empirical testing of sophisticated conjectures about family behaviour. Namely, the unitary framework aggregates individuals' preferences and treats the multi-person household as a single decision-making entity that jointly maximises a well-behaved utility function at the household level (Alderman et al., 1995). There is no consensus over how best to aggregate. One popular adaptation can be found in Becker's (1974) rotten kid theorem, in which a household consists of a benevolent dictator that makes altruistic transfers to fellow egoistic household members who, in turn, act in ways that maximise household welfare because doing so aligns with their private interests. Alternatively, positive assortative matching, whereby individuals with similar traits pair more frequently, thereby allowing preferences to converge (Siow, 2015), can be imposed to justify the synthesis of otherwise individualistic agents.

The unitary model has been subjected to numerous criticisms for its warped representation of reality. For instance, the model treats intra-household inequality as a consensual arrangement on the part of all household members, implying that the deprivation of select individuals is the result of their voluntary relinquishment of private consumption (Folbre, 1986). The inability to address various social phenomena, such as marriage, divorce, and other regularities, is also a major shortcoming (Chiappori, 1991). Browning et al. (2014) further noted that households under the unitary approach display income pooling, a dubious prerequisite stating that distributive behaviours are independent of each member's relative contribution to the pool of household resources. By extension, the unitary model implies that the characteristics of the household member receiving poverty transfers have no bearing on how those additional resources are allocated within the household.

There is no shortage of studies that have attempted to provide empirical rebuttals of the unitary framework using a wide spectrum of methods and justifications for an alternative in the collective approach. But we are able to identify only a handful of studies that investigate how the intra-household distribution of resources responds to cash transfers, fewer for UCT programmes, and even fewer that focus explicitly on inequality among siblings. Early evidence from Thomas (1990) shows that, using household survey data from Brazil and a regression-based reduced-form test of the income pooling restriction, unearned incomes controlled by mothers and fathers have disparate effects on children's health and nutrition. The findings refute the assumed equality in parental income effects under unitary households, which states that the impact of gains in household resources does not depend on the gender of the parent making decisions or of the child in whom resources are invested. In particular, female caretakers channel greater resources towards health and nutrition, and more so for girls. Using Canadian expenditure data, Browning et al. (1990) devise a structural collective model and similarly demonstrate the influence individual incomes have on resource allocation. Holding total expenditure constant, raising the wife's share of household income from 25% to 75% increases her expenditure share by 2.3 percentage points. A 60% increase in total expenditure, holding the wife's income share constant, raises her share by about 12%. Browning et al. also prove that expenditures assigned to each partner depend not only on total income as predicted by the unitary model, but also significantly on socioeconomic factors like ages and lifetime wealth. Cross-country evidence from Quisumbing and Maluccio (2003) highlights the importance of the wife's share of assets at marriage for her bargaining power, while also showing that the role of women's assets is highly context-dependent, proving effectual in some countries but not others.

Some studies exploit policy reforms that shifted the relative contribution to total household income and identify the resulting impact on child outcomes. Lundberg et al. (1997) explore the UK child benefit reform, which transferred recipientship of substantial child allowance payments from fathers to mothers, and find this transferal altered household spending patterns by increasing expenditures on women's and children's goods. These findings are directly related to Lundberg and Pollak (1993), in which a separate spheres model was developed, positing that husbands and wives operate in their own spheres and independently oversee expenditure decisions concerning distinct goods. Duflo (2003) assesses the South African old-age pension program and finds that pensions received by women improved girls' anthropometric outcomes more than pensions received by men, further highlighting the common thread that relative incomes matter. The study is unique in that the benefits were large, at about twice the median per capita income in rural areas.

There are also studies that assess changes in intra-household inequality in response to exogenous shocks, often sharing a similar context in that they focus on how rainfall shocks that adversely impact the income of agrarian families reshape the allocation of resources within the household. These studies are relevant as their findings allude to the well-connected relationship between exogenous shocks – which randomised cash transfer programmes also constitute – and the allocation of resources within households. Additionally, the exogenous nature of shocks provides ideal research designs for identification, with many studies estimating reduced-form fixed effects regressions of child outcomes on either a continuous measure of rainfall deviations from the historical average or a binary indicator denoting rainfall shock. In Uganda, Björkman-Nyqvist (2013) shows that negative rainfall deviation leading to farm income loss has a negative impact on girls' primary school enrolment rate, an effect that is more pronounced for older girls. In Ethiopia, Jose (2024) finds that negative rainfall shocks significantly reduced female-specific expenditures, but the reduction can be assuaged if women increase their temporary off-farm labour supply. Bobonis (2009) also combines randomised variation in PROGRESA enrolment with exogenous rainfall shocks and shows that changes in women's contribution to household income due to the intervention exert substantial influence on spending on food and child goods, while changes in female income due to unexpected rainfall deviations do not.

Several sophisticated demand-system methods have also been used to study intra-household allocation. Studies have estimated Engel curves by using privately assignable goods to recover resource shares. One advantage offered by the Engel curve approach is that individual expenditure data are not necessary for identification. The intuition is that if material goods can be attributed to particular members, such as men's, women's, or children's clothing, then the way clothing expenditures respond to total household resources should reveal the division of resources and bargaining power within the household. The methodological foundation for the estimation of Engel curves was provided by Lewbel and Pendakur (2008). Dunbar et al. (2013) later empirically applied their logic to Malawi, showing that the resource share of the coalition of children increases less than proportionally with size. For instance, one-child families devote around 23% of their resources to the sole child, while children of three-child families receive around only 36% of total household resources, or 12% each, which is consistent with the quality-quantity trade-off. They further show that the extra resources needed to support additional children are largely funnelled from the pool of resources that would have gone to the mothers. One broader implication is that household-level poverty measures severely obscure the true extent of child poverty. Their findings are corroborated by Lechene et al. (2021). Leveraging data from multiple countries, Lechene et al. find that men's resource shares consistently hover around 30%. However, women's and children's resource shares are both lower and more variable, with an increase in one decreasing the other.

The application of the Engel-curve approach to the context of poverty transfers is extremely limited. Casco (2023) studies the impact of Oportunidades, one of the earliest randomised implementations of a CCT programme at a national scale, and finds that transfers shifted 8 percentage points of the household budget away from fathers and to mothers and children. Their work also shows that households where the mother controls the majority of resources assign more resources to food, and that, following adverse shocks, they redirect resources from food to healthcare. Armand et al. (2020) look at another CCT programme, this time in Macedonia, and find that designating women as the recipient increases the share of expenditure dedicated to food by nearly 5%. They also obtain tentative evidence showing that the slope of the Engel curve steepened when transfers were given to the mother. Numerous studies have shown that targeting women as recipients often aligns with the objectives of poverty transfers. In fact, Bonilla et al. (2017) examine the CGP – the very UCT programme we study – and find evidence suggesting the financial empowerment of women only to a modest degree with limited practical significance. Similarly, Weaver et al. (2024) study a large-scale UCT programme in India and show that, despite results indicating significant increases in women's bargaining power and budget share put toward food, child anthropometric outcomes do not significantly improve.

Budget shares are prototypically modelled as linear in log total household expenditure. More complex demand systems, in particular the Quadratic Almost Ideal Demand System (QAIDS), were introduced by Banks et al. (1997), allowing budget shares to vary non-linearly with aggregate expenditure. Unlike the estimation of Engel curves, the QAIDS approach also models the allocation of the household budget not for specific material goods whose ownership can be distinguished, but across several consumption categories jointly. Bourguignon et al. (2009) first theoretically demonstrate the feasibility of testing for Pareto efficiency as implied by the collective model even when there is no price variation. They define distribution factors as variables that affect the intra-household bargaining process and therefore the Pareto weights, without directly affecting preferences, prices, or the household budget constraint. De Rock et al. (2022) then show that allocation decisions by recipients of Mexico's conditional PROGRESA are consistent with the collective approach, but only for the first six months. After more prolonged exposure, transfers began impacting resource allocation not just through shifting bargaining power, but also preferences directly.

Browning et al. (2013) develop a model composed of three components: individual preferences, a sharing rule that governs how much resources go to which spouse, and a consumption technology that captures the economies of scale from joint consumption. They find that wives in Canadian households control two-thirds of household resources at one end where households are wealthier, and one-half at the other end. That the sharing rule is contingent on aggregate resources is inconsistent with the unitary model. Browning et al. further remark that couples save up to one-third of total expenditures through joint consumption, which in turn suggests that singles need to individually spend up to three-fourths as much as couples to approach a standard of living comparable to that of couples. Estimation of QAIDS demand functions still commonly revolves around uncovering the resource shares of household coalitions. Mangiavacchi and Piccoli (2011) find that the receipt of unconditional transfers in Albania led to sharing rules less favourable for children, with egoistic allocations particularly prominent for non-poor beneficiaries.

Another strand of literature, one that is most pertinent to our paper, originates from Behrman et al. (1982, 1986) and is centred around reinforcing and compensatory strategies in the provision of inputs that contribute to the accumulation of a child's human capital. This marks a departure from the aforementioned scholarly efforts in that the overt aim is not to contribute to the unitary vs. collective debate. Behrman et al. develop a parental preference model, in which parents maximise household utility subject to a resource constraint and a human capital production constraint, allowing children to differ in endowments and thus in the marginal productivity of input investments. The key finding is that intra-household inequality in child investments depends on whether parents prioritise inequality minimisation or efficiency maximisation and, at the same time, whether they discriminate according to birth order, gender, or other characteristics. Further theoretical underpinnings for why parents may choose to reinforce existing inequalities were provided by Cunha and Heckman (2007), who postulated that dynamic complementarity between early endowment and later investment, such as preschooling, can explain reinforcement strategies, as investment in more endowed children yields a greater payoff.

The degree to which an allocative scheme is averse to inequality and to which child types are preferred parallel the distinction between statistical and taste-based discrimination. Inequality aversion, or statistical discrimination, is typically operationalised using birth weights as a measure of heterogeneity in genetic endowment. Early evidence from Behrman (1988) demonstrates that aversion to inequality is fluid and adjusts according to seasonal crop harvest in rural India. Behrman finds that parents maximise investment returns during the lean season when food supply is tight. Later empirical evidence has often shown that allocative behaviour espoused by parents does not fully align with reinforcement or compensation. In the same vein, from an early childhood survey on Chilean twins, Abufhele et al. (2017) find that parent investments are neutral, neither increasing nor decreasing the portion of inequality arising from birth weight differentials. Preference for specific child types, or taste-based discrimination, entails that parents attach greater welfare weights to children carrying preferred traits, so that investing in such children yields greater marginal utility for the parent.

More recent findings suggest that the frequent absence of any effect operating through inequality aversion as a channel explaining observed sibling heterogeneity may be rooted in the endogenous nature of traditional measures of endowment. Aizer and Cunha (2012) hypothesise that birth weight or other regularly used endowment measures can be affected by prenatal parental investment, and are thus necessarily serially correlated with both postnatal investment and outcome. Aizer and Cunha construct a plausibly exogenous endowment measure based on determinants such as maternal smoking, mother's pre-pregnancy weight, weight gain during pregnancy, and whether the mother was trying to conceive. Using this measure, they find that parents invested more in better-endowed children. The degree of reinforcement is increasing in household size, suggesting that higher quantity not only associates with worse average quality, but also greater variation in endowment and, by extension, greater reinforcing investments. Other studies have also attempted to mitigate the endogeneity issue in other fashions. Frijters et al. (2010) instrument for cognitive ability differentials with handedness, concluding that parents favoured the higher-ability sibling, with investment increasing by one third of a standard deviation for each standard deviation increase in cognitive gap. Similarly, Giannola (2022) implements a field experiment and finds that participants are more concerned with the efficiency of their investments than the resulting impact on sibling inequality. Their findings highlight the importance of assessing the accuracy with which parents are able to assign endowment levels to their children. Indeed, participants invest 10 % more in children whom they perceived to be 12 % more endowed.

On the other hand, there is substantial empirical evidence for parental preference for specific child types, with gender or age as the observable trait along which such preferences are explored. A broad array of quantities has been studied to quantify child preference. Bharadwaj and Lakdawala (2013) show that visits to health clinics are more frequent when Indian mothers are pregnant with a boy. Barcellos et al. (2014) document that Indian boys receive additional childcare time, are breastfed longer, and receive more vitamin supplements than their female counterparts. One crucial identification hurdle Barcellos et al. reveal is that estimates for the effect of gender on parental investment are biased when gender bias extends to fertility decisions. Jayachandran and Pande (2017) find that later-born boys in India have significantly worse anthropometric outcomes relative to their older siblings. Their results highlight the importance of saturating regressions with mother fixed effects to ensure that estimates are uncontaminated by invalid comparisons between earlier-born children of one family and later-borns of another. Parental preferences are often rationalised by poverty, but Almond et al. (2013) show that gender bias among South and East Asian immigrants to Canada is deeply rooted in cultural practices and persists even after they attain better economic outcomes.

Contextual effects likely play a crucial role, and it is unlikely these findings from India will apply to every other environment. But meta-analytic evidence from Le and Nguyen (2022) finds evidence suggesting son preference across the developing world, where girls' health and nutritional outcomes are worse by about 0.1 to 0.13 SD, depending on the outcome measure. The gender disparity also widens in rural areas where people are less educated and poorer. However, these results are not always produced from studies that carefully isolate taste-based discrimination from statistical discrimination. For instance, if boys earn higher wages, then their superior local labour market conditions, which elevate their return on investment, must be distinguished from cultural norms that may also favour boys. Our mediation analysis, whereby we compare parameter estimates between specifications with and without the aforementioned indicator for gender of the decision-maker, suggests that women's empowerment is not the primary channel through which the CGP affected resource allocation. We advise caution when interpreting the results from structural estimation, as child development, which we use to generate price variation in child investment, may be contemporaneously determined with child outcome.

3. Data

3.1 The Child Grant Programme

Zambia's CGP was an unconditional cash transfer initiative implemented by the Government of Zambia's Ministry of Community Development, Mother and Child Health and funded by three cooperating partners, including UNICEF, the Foreign, Commonwealth & Development Office of the UK government (previously Department for International Development), and Irish Aid. The American Institutes of Research (AIR) was tasked with performing the impact evaluations of the programme. Rollout of the CGP began in 2011 in three districts with high rates of child mortality and morbidity. Geographical targeting of the programme was expanded over time, ultimately reaching 100,000 households five years after it was first introduced. The CGP was also preceded by two pilot transfer schemes several years prior, accompanied by the Multiple Categorical Targeting Grant (MCTG) programme for the majority of its span of existence between the years of 2010 and 2014, and superseded in 2015 by the harmonised Social Cash Transfer (SCT) programme aimed at consolidating various transfer schemes including the CGP.

The Carolina Population Center (2025) provided the data for the CGP. The categorical targeting model of the CGP identified any household with at least one child under the age of five. But the cutoff for entry into the evaluation sample was narrower, where only those that had a child under age three were included to ensure at least two years of transfer exposure. How newborns were handled also differed between the programme entry criterion and the evaluation inclusion criterion: households were immediately enrolled after having a newborn baby but were nevertheless excluded from the evaluation sample. Eligible households were scheduled to receive transfers averaging 11 US dollars a month irrespective of size, made bimonthly and collected at designated local pay points by, in principle, the main female caretaker of the household. A household remained a beneficiary of the CGP as long as its focal child, the youngest member of the household at time of enrolment, had not turned six years of age.

Community clusters, or Community Welfare Assistance Committees (CWAC) as they were formally called, were selected by lottery from the three initial districts. Communities were then randomly assigned to either begin receiving cash transfers in December 2010 or towards the end of 2013. The delayed entry group, however, was ultimately not enrolled in the CGP, and so for practical purposes the originally planned staggered rollout provides a clean distinction of treatment and control communities throughout the entire study period. From each community, 28 households were randomly sampled to be included in the evaluation of the CGP, and the study sample totalled just over 2500 households.

Baseline data collection occurred in October 2010, with treatment assignment shortly following. The selected communities were revisited four further times, at 24, 30, 36, and 48 months after the baseline. Except for the two-and-a-half-year follow-up wave, the endline surveys were conducted around the same time of the year as the baseline when possible to minimise the influence of seasonality. Data were entered on two separate terminals as they came in from the field and screened for inconsistencies, substantially reducing the likelihood of erroneous inputs. Additionally, key variables such as age and child anthropometry were reconciled across survey waves.

3.2 Key Variables

We evaluate the intra-household impact of the CGP for three outcome domains. These, which we deem pertinent to the well-being of children and sufficiently reflective of the level of parental investment in their human capital, include health and nutrition, material needs, and schooling. For health and nutrition, we use height-for-age z-score (HAZ), weight-for-age z-score (WAZ), and weight-for-height z-score (WHZ), three continuous anthropometric measurements that express, in standard deviations (SD), the growth status of a child relative to the reference population of the same age and sex. We omit values that are larger than plus or minus 6 – a standard practice. For material needs, we employ three binary-coded indicators that take the value 1 if the child has access to a blanket, shoes, or two sets of clothes, and 0 otherwise. For schooling, we opt for a binary enrolment that is coded as 1 if the child is receiving formal or informal education at any level.

We illustrate in Figure 1 the extent of intra-household inequality at baseline for the aforementioned outcome measures. Panel A plots individual outcome against the mean outcome of the focal child's household. Using the -2 threshold for severe malnutrition, we are able to categorise children into four quadrants where the top-left sector, highlighted in black, consists of children who were suffering from arrested growth but the group of children in their household, on average, did not. The black points therefore capture severely malnourished children who would not have been identified by social welfare programmes employing household-level targeting. Conversely, dark gray points correspond to children who would have belonged to eligible households under such programmes but who were themselves adequately nourished. In fact, any point deviating from the 45-degree diagonal line encapsulates inequality between a child and their siblings and, the greater the distance, the more pronounced the degree of intra-household inequality. Overall, the figure documents similar

Figure 1: **Intra-Household Inequality in Child Outcomes at Baseline**

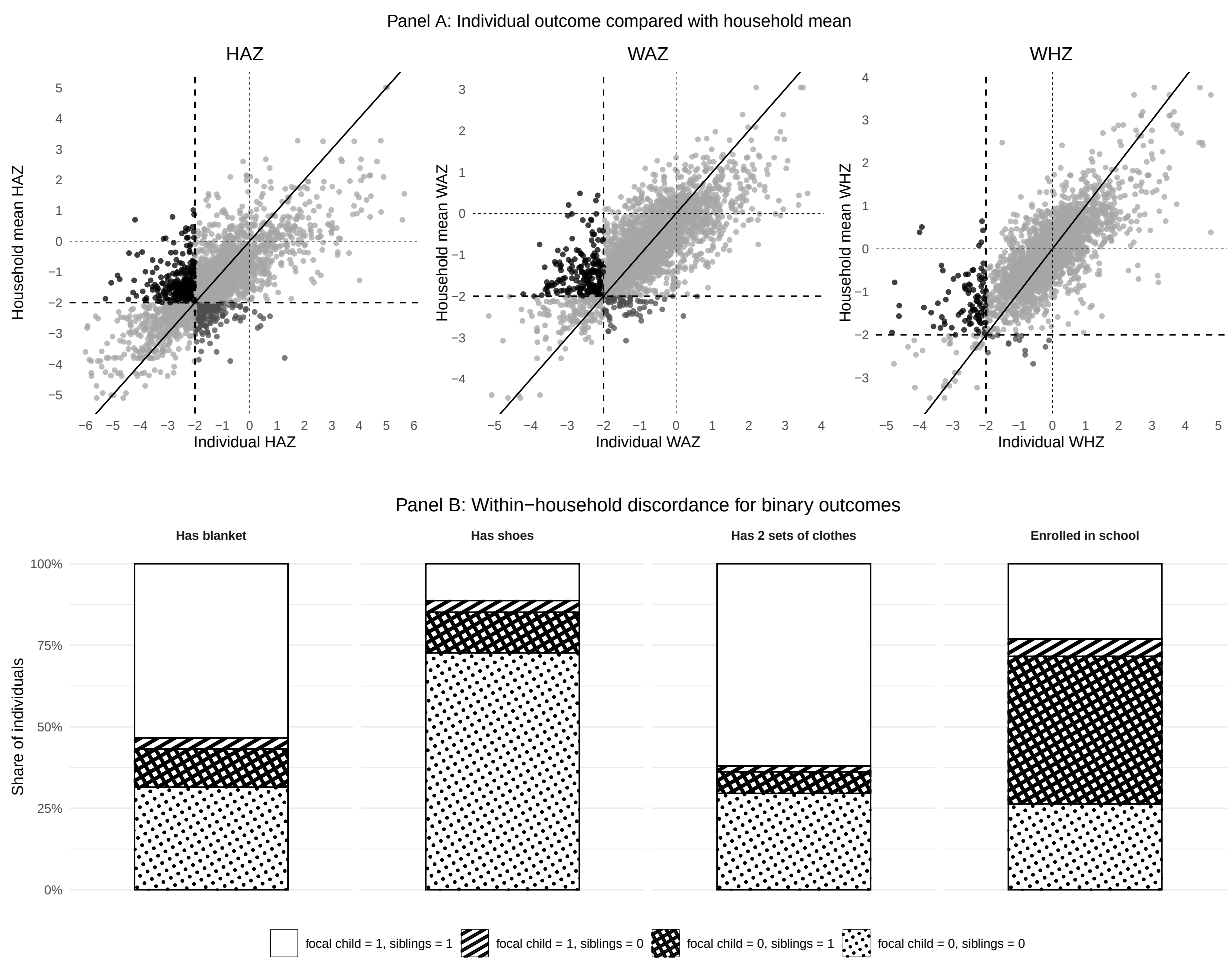

Note: Panel A plots individual outcomes on the x-axis against the mean outcome of their siblings from the same household on the y-axis. The thicker, horizontal and vertical dashed lines mark the -2 thresholds on the axes, while the thinner, horizontal and vertical dashed lines mark the 0s on the axes. Black points correspond to children who were malnourished according to the -2 threshold when their household, on average, was not. Dark gray points correspond to the opposite. Panel B classifies children into four categories for each binary outcome. The white segment represents children who have the material good or are enrolled in school and have at least one sibling who also does. The striped segment represents children who have the good or are enrolled while none of their siblings does. The criss-cross segment represents children who do not have the good or are not enrolled while at least one sibling does. The dotted segment represents children for whom neither the focal child nor any sibling has the good or is enrolled.

levels of intra-household inequality for all three anthropometric measurements, as indicated by the elongated spread of points along the diagonal line.

For binary outcomes, we use a different visualisation strategy. Children are classified into four categories according to whether the focal child and their siblings have access to the material good or are enrolled in school. Note that, for simplicity, we treat a sibling coalition in which at least one child is able to access the good or is enrolled in school as 1, and 0 otherwise. Then, the striped segment, where the focal child is on average better off than their siblings, and the cross-hatched segment, where the reverse is true, capture cases of intra-household inequality. It is immediately apparent that the baseline distribution of schooling outcomes was drastically more inequitable than the distribution of material goods. We examine these patterns in greater detail in the following sections.

In order to reduce the plurality of outcome variables for which we report effects, we use principal component analysis (PCA) to create two indices, one reflecting the child's nutritional status and another reflecting their access to material needs. We then retrieve the first principal components (PC1), which we enter into our estimating equations in conjunction with enrolment to form the core triad of child outcomes.

Next, we briefly discuss the operationalisation of the two indices, including the constituent variables, standardisation procedure, and associated PC1 loadings. Although HAZ, WAZ, and WHZ are already standardised relative to the WHO growth-reference population, we centre and scale them again within the PCA estimation sample so as to ensure that each anthropometric measure contributes on comparable sample-specific scales and that the PC1 is not driven mechanically by disparities in sample variance across the three z-scores. Table 1 provides the results for the PCA. In the health domain, HAZ, WAZ, and WHZ have loadings 0.435, 0.758, and 0.485, respectively. Given that PC1 assigns positive and significant weights to all three anthropometric outcomes and explains a satisfactory 56.9% of the total variation in the anthropometric measurements, the first component can be appropriately interpreted as an overall child nutritional status index. We extend the same argument for the material access domain. The extremely similar loadings – 0.596 for having a blanket, 0.579 for having shoes, and 0.557 for having two sets of clothes – indicate that PC1, which explains 53% of total variance, captures a common dimension of material-needs satisfaction. In both cases, the high positive correlation among all constituent variables motivates the use of PC1 as the dependent variable index.

3.3 Sample Characteristics

We now turn to an overview of the sample characteristics. To begin with, we provide in Table 2 the summary statistics and results from tests of baseline balance for the outcome measures, corresponding binary classifications in the case of anthropometric outcomes, domain PC1s, as well as other control variables of interest. Randomisation ensured comparability between treatment and control groups at baseline. At baseline, low height-for-age presented as the most common manifestation of malnutrition, with one third of the child sample stunted. At around one in every seven, a smaller

Table 1: **Principal Component Indices**

Anthropometry			
PC1 Loadings			
Variable	Mean	SD	PC1 loading
Height-for-age z-score (HAZ)	-1.418	1.434	0.435
Weight-for-age z-score (WAZ)	-0.949	1.088	0.758
Weight-for-height z-score (WHZ)	-0.134	1.285	0.485
Variance Explained			
Component	Eigenvalue	Variance explained	Cum. variance explained
PC1	1.708	0.569	0.569
PC2	1.251	0.417	0.986
PC3	0.041	0.014	1
Material Needs			
PC1 Loadings			
Variable	Mean	SD	PC1 loading
Has blanket	0.839	0.368	0.596
Has shoes	0.464	0.499	0.579
Has two sets of clothes	0.893	0.309	0.557
Variance Explained			
Component	Eigenvalue	Variance explained	Cum. variance explained
PC1	1.590	0.530	0.530
PC2	0.739	0.246	0.776
PC3	0.671	0.224	1

Notes: Mean and SD are those of a pooled sample that includes data from all survey waves, hence values may differ from those presented in the summary statistics table.

but still sizeable proportion of children suffered from underweight. Wasting was the least common form of malnourishment in our study sample, with an average WHZ of -0.15 for the control arm and -0.2 for the treatment arm and an overall prevalence rate of 6%. This is expected, as wasting reflects acute nutritional deficits rather than the chronic deprivation captured by stunting and underweight, which are more prevalent in this persistently impoverished setting. None of the anthropometric outcomes, including their corresponding binary indicators as well as the health and nutrition index (HNI) we constructed using PCA, shows statistically significant differences at baseline.

We also study fulfilment of three types of children's essential material needs. These are private assignable goods that can be consumed by only a single person within the household at any given moment, and their defining characteristics, rivalrous and excludable, render them ideal items for studying intra-household allocation of resources. Just over half of the child sample aged between 5 and 17 at baseline had access to a blanket, while around two in three had access to two sets of clothes. Shoes were relatively uncommon commodities, with 14-15% having them initially. The differences between the two intervention arms are small and statistically insignificant for all three material needs indicators and their PC1, the material needs index (MNI). In the same vein, there is no indication of baseline imbalance for school enrolment, where just under 60% of children aged 5 to 17 were enrolled in school. Turning to individual demographics, it is clear that we are working with a relatively young sample, where the average age is 15.40 in the control group and 15.91 in the treatment group. The difference of 0.52 years is statistically significant at the 5% level. On the other hand, the share of females is very similar across groups, at 53% for the control group and 54% for the treatment group.

At the household level, we observe some additional facets along which there was group imbalance at baseline. Treated households were slightly older on average and their transfer recipients (for non-treated households, the primary female caretakers) were also slightly better educated. We illustrate in fine detail the demographic composition of the sample at baseline in Figure A1. In particular, we see that a household roster typically consisted of 5 to 6 members, of which there would be 2 under-seven children and 2 children aged five to seventeen. Intra-household inequality becomes a moot notion for households without at least one child pair; such households thus make up the primary estimating sample of our study. Further, it is shown that the vast majority of

Table 2: **Baseline balance for key variables used in our analysis**

Variable	Binary	Control $\bar{\mu}$	Treatment $\bar{\mu}$	Diff	p-value	N
Panel A: Individual-level variables						
Anthropometry (age 0-7)						
health and nutrition index		0.08	0.07	-0.02	0.716	3418
HAZ		-1.34	-1.30	0.04	0.490	3418
stunted (HAZ < -2)	x	0.32	0.32	-0.01	0.661	3418
WAZ		-0.85	-0.87	-0.01	0.727	3418
underweight (WAZ < -2)	x	0.14	0.15	0.00	0.868	3418
WHZ		-0.15	-0.20	-0.05	0.212	3418
wasted (WHZ < -2)	x	0.06	0.06	-0.00	0.677	3418
Material needs (age 5-17)						
material needs index		-1.29	-1.30	-0.01	0.905	4743
has blanket	x	0.56	0.55	-0.01	0.448	4743
has shoes	x	0.14	0.15	0.01	0.378	4743
has 2 sets of clothes	x	0.64	0.64	0.00	0.928	4743
Enrolment (age 5-17)						
enrolled in school	x	0.58	0.59	0.02	0.287	4818
Demographics						
age		15.40	15.91	0.52^{**}	0.041	13690
female	x	0.53	0.54	0.01	0.349	13690
Panel B: Household-level variables						
log household size		1.66	1.68	0.02	0.315	2503
household average age		15.88	16.22	0.34^{*}	0.058	2503
proportion of household female		0.54	0.55	0.01	0.152	2503
expenditure per capita		39.47	41.49	2.01	0.120	2503
recipient age		29.65	30.09	0.44	0.251	2503
recipient highest grade		3.79	4.32	0.53^{***}	0.000	2503
recipient marital status		2.30	2.24	-0.06	0.144	2503

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Education codes are defined as follows: grades 1–12 correspond to codes 01–12; GCE A-level is coded as 13; college is coded as 14; undergraduate university is coded as 15; post-graduate certificate or diploma is coded as 16; master's degree is coded as 17; and doctoral level or above is coded as 18. Marital status codes are defined as follows: never married is coded as 1, married is coded as 2, separated is coded as 3, divorced is coded as 4, widowed is coded as 5, and co-habiting is coded as 6.

primary female caretakers were married, aged somewhere between 20 and 40 years old, and had often received no formal schooling.

3.4 Attrition

Table 3 describes the rates of and Table A1 the reasons for attrition by age group, survey wave, and intervention arm. We show that attrition rates started at around 10% at 24 months, the first post-treatment survey wave, before gradually rising over subsequent waves and peaking at around 19% at the final, 48-month follow-up wave. Attrition rates were especially high for children aged 0-7 at baseline. However, there were no statistically significant differences in the prevalence of attrition between the treatment and control groups that would suggest differential attrition. In Table A1, we provide a more detailed breakdown of the causes for attrition. For both age groups, the most common reason for attrition was moving to live with other relatives. Among children aged 0–7, the second most common cause was death, whereas among children aged 5–17, it was getting married and moving away; leaving the household with parents was the third most common cause in both age groups. Although we observe significant attrition, especially towards the later survey waves, we do not find strong evidence pointing towards the presence of disparate attrition patterns between the two intervention arms. Moreover, the number of attritors due to death was low, and as documented in Handa et al. (2019), loss to follow-up did not correlate with the vast majority of the 62 variables they evaluated, encompassing demographics, assets, and expenditures. For these reasons, we do not make any adjustment for attrition. This approach is in line with the official impact evaluation of the CGP (AIR, 2016).

Table 3: **Attrition Rates**

	All individuals			Age 0–7			Age 5–17		
Time	Control	Treat	Diff	Control	Treat	Diff	Control	Treat	Diff
24m	9.96	9.59	-0.37	11.58	10.99	-0.59	8.72	8.41	-0.31
30m	15.25	14.49	-0.76	18.58	16.83	-1.75*	12.33	12.43	0.11
36m	16.70	16.13	-0.57	21.04	19.31	-1.73	12.87	13.49	0.62
48m	18.77	18.67	-0.09	23.79	22.64	-1.16	14.76	15.56	0.80

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Newborns after baseline are excluded. Age groups are defined using age observed at baseline.

3.5 Programme Graduation

As noted earlier, the evaluation sample consists of households that at baseline had at least one child aged three or younger. This, in conjunction with the predetermined programme exit rule, under which households graduated from the CGP once the youngest child at baseline turned six, meant that each household would have had a period of continuous transfer exposure spanning no shorter than two years. Data for the earliest follow-up survey wave of the CGP evaluation sample were collected 24 months after programme commencement, but data collection for the fourth and final survey wave did not take place until two years later, indicating that, starting from the 24-month mark, an increasing proportion of households randomly assigned to receive transfers from the CGP had graduated from the programme and lost their beneficiary status.

Graduation rates over time are shown in Table 4. We infer the eligibility status of a household assigned to receive transfers at each follow-up wave from the age of its youngest child at baseline, but the computed graduation rates may not be entirely accurate. This is because, in practice, programme graduation was not administered automatically or on a monthly basis, allowing certain treated households to receive transfers for a few months longer than they should have (Handa et al., 2025). Our computed graduation rates were approximately 8% and 12% after two and three years, respectively, before rapidly ramping up to 34%, or one in three households initially assigned to the treatment arm, four years after baseline. Graduation represents a form of non-compliance, diluting estimates of the average treatment effect of treatment assignment. Consequently, we supplement our main analyses with an instrumental variable approach and discuss the rationale more closely in subsequent sections.

Table 4: **Graduation from the CGP**

	No. of treated households		
Time	Total	Graduated	Graduation rate
24m	1154	93	8.06
30m	1194	121	10.13
36m	1230	147	11.95
48m	1220	414	33.93

4. Conceptual Framework

Several frameworks have been advanced in the literature aimed at addressing the mechanisms leading to intra-household inequality. In this study, we engage with the theoretical groundwork laid by Behrman et al. (1982), in which child inequality within the household is at the intersection of two concepts: aversion to inequality and discriminatory child preference. Allocative strategies concerning inequality aversion can be broadly grouped into three categories. At one extreme, an efficiency-oriented household may choose to maximise returns on investment by channeling resources to the next child for whom the marginal return net of marginal cost is highest, and they continue to do so until no additional returns can be attained (Becker, 1975). Alternatively, an egalitarian household may opt to channel resources to the next least well-off child, and they continue to do so until there is no remaining inequality (Lamont & Favor, 2017). Households between these extremes can then be described as adopting a mixed, Cobb-Douglas approach. Allocative strategies concerning child preference are much simpler: greater welfare weights are attached to children carrying traits that are viewed more favourably by parents.

We depict in Figure 2 how the two quantities driving allocation interact. The inverted parabolic curve is the production possibility frontier that outlines the maximum combinations of child quality that can be attained across the two children. Notice that the boundary solution on the frontier is placed more outwardly for child 2. This is because we assume child 2 to be more productive in producing child quality, an assumption we make to facilitate comparison of combinations of outcomes under alternative scenarios. IC_1, IC_2, and IC_3 are indifference curves that respectively correspond to preference for traits carried by child 2, indifference, and preference for traits carried by child 1. It is immediately clear that inequality aversion alters the curvature of indifference curves, while child preference alters their symmetry.

We hypothesise that the CGP, through amelioration of the financial constraint faced by recipients, *reduces overall intra-household inequality*. We further hypothesise that the reduction, and more broadly shifts in allocative behaviour, are primarily the result of *increased aversion to inequality*, while preference for specific child types *remains unchanged*. This can be rationalised by postulating that the short-term programme exposure of two to four years is insufficient to introduce any noticeable changes in the persistent social, economic, or cultural norms that embed parents' child preferences. Conversely, compensatory or reinforcing redistributive responses to existing sibling differentials are state-contingent and vary with household resources, seasonality, or transitory shocks. Additionally, we predict that the magnitude and plausibly the direction of change in inequality among siblings *vary across different domains of outcomes*, for which households may adopt differing allocative strategies. Lastly, since the designated recipients of the CGP were invariably women, who, relative to men, tend to exhibit greater aversion to inequality, we expect the effects of the CGP to be magnified through a positive feedback loop that *empowers the more inequality-averse women*.

Figure 2: **Inequality Aversion and Child Preference**

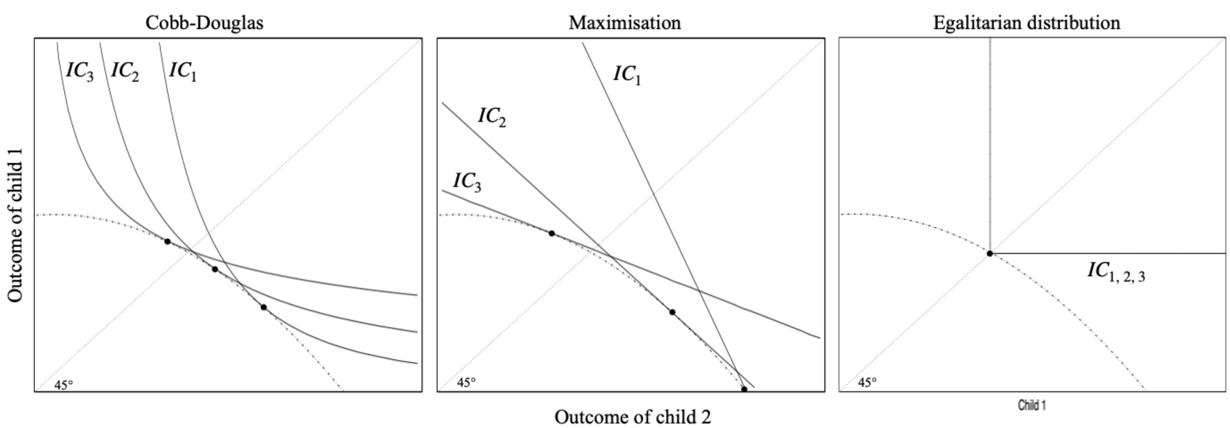

5. Methodology

5.1 Programme Impact

Before proceeding with our main analysis that pertains to intra-household inequality, we first consider how the CGP affects child outcomes at the aggregate level, following the approach of the programme's impact evaluations (Handa et al., 2018, 2025). Randomisation provides a favourable setting for identification. We employ a difference-in-differences (DID) model to examine the magnitude and direction, if any, of differential outcome trajectories exhibited over five survey waves between recipients and non-recipients. Although we do not expect our estimation results on programme impacts to deviate markedly from those in Handa et al., we mark a significant departure from their work in how we handle programme graduation, which we have discussed in detail earlier. We begin by replicating Handa et al. by estimating the wave-specific intention-to-treat (ITT) effect with the following equation:

$$Y_{ihjdt} = \beta_1 \, \text{Treat}_j + \beta_2 \, \text{Time}_t + \beta_3 \, (\text{Treat}_j \times \text{Time}_t) + \sum_{k=1}^{K} \delta_k X_{it} + \sum_{l=1}^{L} \lambda_l Z_h + \alpha_d + \varepsilon_{ihjdt} \quad (1)$$

where Y_{ihjdt} is the outcome of interest for child i in household h of community j in district d during survey wave t, $Treat_j$ is an indicator for treatment assignment, $Time_t$ is a collection of four dummy variables that each correspond to a specific follow-up survey wave, X_{it} is a vector of time-varying individual controls, Z_h is a vector of household covariates measured at baseline so as to attenuate concerns about bias, and ε_{ihjdt} is the error term. In the baseline and all subsequent specifications, we also include district fixed effects, α_d, to account for static unobserved heterogeneity between districts – the level of stratification, and cluster standard errors at the community level – the level of randomisation. For ease of interpretation, we also provide estimation results using an alternative specification for (1), where we switch $Time_t$ out for $Post_t$, thereby yielding a pooled treatment effect estimate across the four endline waves.

Our selection of covariates entered into the estimating equation is guided by Handa et al. and other empirical findings. Namely, in (1) we control for child age and gender, household per capita expenditure, demographic composition including the natural logarithm of size, average age, and proportion of female members, and the characteristics of the recipient including their age, marital status, highest education level attained. We also use the baseline values for household covariates, as they are affected by both past treatment and past outcomes, and would consequently open up backdoor paths if included in the regressions as time-varying covariates, leading to the problem of bad controls. For anthropometric outcomes, children aged seven or under are included in the regressions, while the impacts on outcomes related to schooling and material-needs access are estimated using samples of children aged between 5 and 17. All subsequent regressions use the same relevant age groups unless stated otherwise.

5.2 Intra-Household Inequality: Variance Decomposition

We devise two empirical methods that chiefly explore the response of intra-household inequality to the CGP. Indeed, specifications in the previous section conflate within-household variation with between-household variation. First and foremost, we present descriptive evidence with variance decomposition using the law of total variance[1] (Casella & Berger, 2002). Variance components are segregated by treatment assignment and survey wave to permit a tentative inspection of the temporal evolution of inequality nested within the household for those who receive the transfer and for those that do not. Because the outcomes of interest often take non-positive values, variance is our preferred choice over alternative, similarly decomposable, and regularly used measures, such as the Theil index (e.g., De Vreyer & Lambert, 2021) or the Gini coefficient (e.g., Malghan & Swaminathan, 2021), which do not allow non-positive values.

For continuous $y^{i,h} \in \mathbb{R}$, the variance can be decomposed as

$$\underbrace{\mathrm{var}\left(y^{i,h}\right)}_{\text{Total}} = \underbrace{\mathrm{var}_h\left(\bar{y}^{\,h}\right)}_{\text{Between-household}} + \underbrace{\sum_h \omega_h \ \mathrm{var}_i\left(y^{i,h} \mid i \in h\right)}_{\text{Within-household}}. \tag{2}$$

where ω_h is the weight for household h in the variance decomposition, computed as $\frac{n_h}{N}$, where n_h is the number of children contributed to the calculation by household h and N is the total number of observed children in the relevant sample. For binary $y^{i,h} \in \{0, 1\}$ with Bernoulli variance $p(1 - p)$, Equation (2) collapses to

$$\underbrace{p(1-p)}_{\text{Total}} = \underbrace{\mathrm{var}_h(p_h)}_{\text{Between-household}} + \underbrace{\sum_h \omega_h \, p_h(1-p_h)}_{\text{Within-household}}, \tag{3}$$

where $\mathrm{var}_h(p_h) \in [0, 0.25]$, and the within-household component, which is also bounded between 0 and 0.25, is maximised at $p_h = 0.5$ and minimised at $p_h = 0$ or 1. Intuitively, intra-household inequality is greatest when children in the same household are evenly split between those with and without the outcome, since this represents the maximum possible discordance. By contrast, when all children in a household have the outcome or none do, they are all just as well off or impoverished as every other sibling, representing no intra-household inequality.

Variance is translation invariant and unbounded, attractive properties that other dispersion measures do not offer. However, variance is sensitive to the scale of the measure, making it difficult to gauge the relative magnitude of within-household inequality of HNI, MNI, and enrolment. To facilitate comparison, we report and mainly focus on the proportion of within-household variance out of total variance, denoted by $\frac{W}{W+B}$.

[1] For applications of variance decomposition to binary indicators, see Breen and Ermisch (2021) and Berman et al. (2025).

5.3 Intra-Household Inequality: Fixed Effects Model

In our second empirical approach to examining responses of intra-household inequality under the CGP, we estimate a three-way interaction model saturated by household-by-time fixed effects, in which identification relies solely on comparisons across siblings within the same household and survey wave. Fixed effects models are frequently used in twin studies (e.g., Fujiwara & Kawachi, 2009; Kovas et al., 2013; Yi et al., 2015), and our approach can thus be viewed as an extension of the twin fixed effects model. The equation we estimate is

$$Y_{ihjdt} = \gamma\left(\text{Treat}_j \times \text{Time}_t \times BO_i\right) + \sum_{k=1}^{K} \theta_k X_{it} + \alpha_{ht} + \mu_d + \varepsilon_{ihjdt} \tag{4}$$

where $Time_t$ is a set of four dummy indicators, each corresponding to one of the four follow-up survey waves, BO_i denotes the birth order of child i, X_{it} includes age and whichever of birth order and gender has not been included in the interaction term, and α_{ht} is the term for household and survey wave interactive fixed effects. As before, we include district fixed effects and clustered standard errors at the community level. With interactive fixed effects, shocks common to all household members at each wave are eliminated, thereby also eliminating the need to control for baseline household characteristics.

In the regressions, birth order is classified into four levels: the firstborn, second-born, third-born, and fourth-born or higher. Considering that compound households where multiple related nuclear families live in the same residential space are not an uncommon living arrangement in the context of extreme poverty in Zambia, we decide to rank birth order among maternal siblings whom we are able to identify as having the same mother, rather than among all those who live within the same household. Birth order is defined not only with respect to siblings who satisfy the age criterion we impose for each estimation, but to siblings of all ages who share an identical maternal lineage. Additionally, we omit singleton maternal sibling groups from all estimations as opposed to assigning them firstborn rank. The impact of the omission is minor, as a singleton would have only contributed to the coefficient estimates in the specific scenario that their household is inhabited by at least another group of maternal siblings that includes at least one child who is not also firstborn.

While partial effects of receiving the treatment are not recoverable under wave-specific household fixed effects, as treatment assignment does not vary within the household, our interest lies in assessing the impact on inequality among siblings. Therefore, we estimate the treatment effect on the birth-order gradient of sibling outcomes within households using a triple interaction term, $\text{Treat}_j \times \text{Time}_t \times BO_i$. The coefficient of interest, γ, can be interpreted as the impact of the treatment on the degree to which sibling outcomes differ by birth order or gender within the household. A larger γ indicates a greater difference between the reference child and their siblings, where we set the former to either the firstborn in the birth order specification or the girl in the gender specification. It is also worth noting that all observations of a household at a specific wave are dropped if there does not exist birth order or gender variation within it. Within-household gender variation is a stricter requirement, and the estimation sample for it is therefore also smaller than that for birth order variation, which only demands that there be two valid observations within the household-wave cell.

6. Results

6.1 Programme Impact

ITT estimates for the impact of the CGP on child outcomes are presented in Table 5. We find little evidence suggesting that the programme improved children's health and nutrition. Estimated effects on the health and nutrition index are small and statistically insignificant across all follow-up waves, with a pooled effect estimate of 0.048 units. Table A2 reports the corresponding estimates for the index's constituent anthropometric outcomes, showing that the small positive estimates for the HNI are mainly driven by modest positive effects on WAZ and WHZ, of approximately 0.070 and 0.059 SDs, respectively, which are partially offset by a minor negative effect on HAZ of -0.012 SDs. Coefficient estimates for weight appear generally larger in magnitude and greater in precision than those for height. This is not surprising, as weight tends to respond to short-term shocks and exhibits greater volatility, whereas height is indicative of long-term chronic deprivation and is likely determined by a broader array of both prenatal and postnatal factors.

Table 5: **ITT Effects of the CGP**

	HNI	MNI	Enrol
Panel A: Wave-Specific Effect			
Treated × 24m	0.084	0.839***	0.010
	(0.065)	(0.165)	(0.021)
Treated × 30m	0.031	0.626***	0.057***
	(0.071)	(0.167)	(0.019)
Treated × 36m	-0.002	0.550***	0.025
	(0.078)	(0.151)	(0.020)
Treated × 48m	0.092	0.680***	0.045**
	(0.088)	(0.181)	(0.020)
Panel B: Pooled Effect			
Treated × Post	0.048	0.665***	0.035**
	(0.065)	(0.154)	(0.017)
Age range	0–7	5–17	5–17
Mean depvar at baseline	0.074	-1.293	0.586
Observations	21370	26983	27544
R^2	0.034	0.316	0.201
District FE	Yes	Yes	Yes

Note: Standard errors clustered at the CWAC level in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

We further observe from Table 5 substantial improvements in material well-being following the CGP. The pooled treatment effect on MNI is 0.665 units, statistically significant at the 1% level, with wave-specific estimates ranging from 0.550 to 0.839. Further inspection of the constituent variables in Table A2 reveals that the gains were primarily driven by large increases in the probability of having shoes and blankets – the two scarcest material goods at baseline – while the effect on clothing was positive but small and statistically insignificant.

The CGP also exerted a positive but modest impact on school enrolment, with a pooled estimate of 3.5 percentage points. The effect is statistically significant at 30 and 48 months but not at the other waves.

Because programme graduation was not uncommon, we also report estimates for the average treatment effect on the treated (ATT) of actually receiving transfers. We outline the empirical approach in detail in Section A.1 ATT Estimates of the Impact of the CGP and report the results in Table A3.

6.2 Intra-Household Inequality: Variance Decomposition

Table 6 presents the results from variance decomposition. We show, separately for baseline and each follow-up, the size of the within-household component, denoted by W; the size of the between-household component, denoted by B; and the proportion of the within component relative to total variance, denoted by $\frac{W}{W+B}$. Additionally, we report double differences for each of these three measures, obtained by first differencing across intervention arms and then across survey waves. The quantity of interest is the DID of W, which pertains to the treatment effect on outcome dispersion within the household.

We find that within-household variance in HNI fell by more among recipients than among controls across all follow-up waves, with reductions ranging from roughly 3% to 22% of the baseline control-group level. Sibling inequality in material access contracted to a noticeably greater extent, with within-household variance universally declining by between 15% and 43% relative to baseline. A common thread across these two domains is that the effectiveness of the CGP at reducing within-household inequality wanes after three years of treatment before rebounding at four years. By contrast, we report near-null effects on within-household inequality in educational investment, barring a modest increase at 30 months.

To compare the extent to which sibling inequality contributes to total inequality across outcome domains measured on different scales, we rely on the normalised within share, $\frac{W}{W+B}$. A striking pattern emerges: within-household inequality accounts for around 40–50% of total variance in HNI, only 10–20% for MNI, and over 70% for enrolment. These differences indicate that qualitatively distinct allocative schemes govern different domains of goods and services.

Turning to the DID of $\frac{W}{W+B}$, the within shares, we find that the relative contribution of within-household inequality to total HNI variance remained largely unchanged until 48 months, at which point treated households exhibited a notably higher within share than control households. For MNI, the CGP led to a consistent rise in the within share across all waves, peaking at 36 months, before contracting towards the end of the sample period — a trajectory that stands in contrast to that of HNI. We also find a small CGP-induced shift in the within share of enrolment that initially increased and subsequently decreased.

Before moving onto the results for our main identification approach, we illustrate in Figure 3 the interaction of transfer-induced changes to allocative behaviour with the degree of resource constraint endured by the household. We plot the DIDs of W against 100 approximately equally sized bins of household expenditure per capita observed at baseline. We find a greater concentration of reductions in within-household inequality in HNI among moderately constrained households and, to a lesser extent, the least constrained ones. The reductions also appear to have gradually heightened over the follow-up waves. Similarly, we observe a greater CGP-associated decrease in intra-household inequality in material well-being among the less constrained households, which increased in magnitude over time. These disparate adjustments between the extreme and non-extreme poor suggest that how constrained a household is plays a prominent role in explaining the different allocative schemes followed for different domains of goods and services. We refrain from commenting on school enrolment as the changes are an order of magnitude smaller than in the other two domains and do not present a consistent pattern regarding the moderating effects of household resource constraints.

Table 6: **Variance Decomposition by Survey Wave**

		Control			Treatment			DID		
Outcome	Wave	W	B	$\frac{W}{W+B}$	W	B	$\frac{W}{W+B}$	W	B	$\frac{W}{W+B}$
Panel A: Health and Nutrition Index										
	Baseline	0.782	1.246	38.6%	0.867	1.261	40.7%			
	FU (24m)	0.676	0.835	44.7%	0.728	0.827	46.8%	-0.033	-0.023	-0.1%
	FU (30m)	0.744	0.916	44.8%	0.715	0.825	46.4%	-0.114	-0.106	-0.6%
	FU (36m)	0.754	0.889	45.9%	0.815	0.806	50.3%	-0.024	-0.098	2.2%
	FU (48m)	0.778	0.916	45.9%	0.694	1.012	40.7%	-0.169	0.081	-7.4%
Panel B: Material Needs Index										
	Baseline	0.303	1.830	14.2%	0.339	1.901	15.1%			
	FU (24m)	0.140	1.346	9.4%	0.082	0.550	13.0%	-0.094	-0.867	2.7%
	FU (30m)	0.139	1.011	12.1%	0.115	0.537	17.6%	-0.060	-0.545	4.6%
	FU (36m)	0.130	0.868	13.1%	0.119	0.412	22.4%	-0.047	-0.527	8.4%
	FU (48m)	0.177	1.255	12.3%	0.083	0.457	15.4%	-0.130	-0.869	2.1%
Panel C: Enrolment										
	Baseline	0.155	0.045	77.3%	0.159	0.047	77.2%			
	FU (24m)	0.172	0.048	78.2%	0.177	0.045	79.6%	0.001	-0.005	1.5%
	FU (30m)	0.178	0.041	81.2%	0.191	0.039	83.2%	0.009	-0.004	2.1%
	FU (36m)	0.175	0.038	82.0%	0.181	0.037	83.0%	0.002	-0.003	1.1%
	FU (48m)	0.176	0.073	70.6%	0.178	0.072	71.3%	-0.002	-0.003	0.8%

Figure 3: **Difference-in-Differences of Within-Household Variance by Household Expenditure Per Capita**

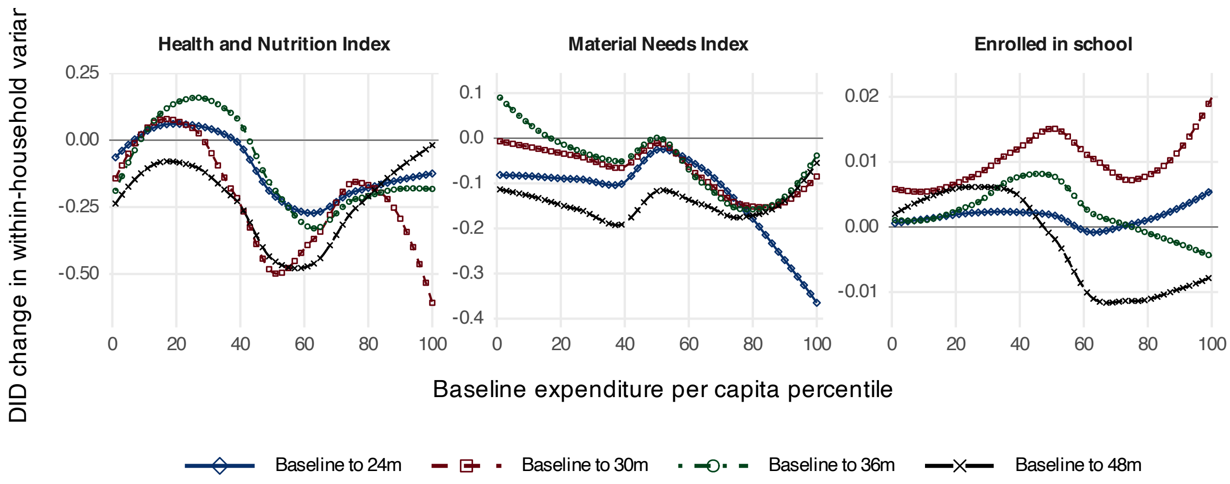

Note: Expenditure per capita percentiles are defined using values observed at baseline. Blue, red, green, and black lines correspond to the difference-in-difference of within-household variance component between baseline and 24, 30, 36, 48 months, respectively.

6.3 Intra-Household Inequality: Fixed Effects Model

6.3.1 Baseline Estimates

We next present the impact of the CGP on intra-household birth-order gradients. Ranking children by birth order within the household offers a convenient way to reliably establish sibling comparisons along an observable child trait, which is necessary for the purpose of understanding the extent of intra-household inequality. Another advantage of the birth-order approach is its relatively low data requirement, allowing us to retain the maximum number of observations. Specifically, with the inclusion of interactive household-by-time fixed effects, households only need to have at least two children for whom the precedence of birth can be established. Then, the three-way interactions can be interpreted as either the wave-specific or pooled effect of the CGP on the outcome gap between BO1 and their later-born siblings. We also combine the coefficient estimates with birth-order (BO) gaps at baseline for a more nuanced interpretation, through which we are informed of whether changes in birth-order gaps are reinforcing or compensatory in nature.

Estimation results are reported in Table 7. We find some evidence for behavioural adjustments in the allocation of health and nutritional resources induced by the CGP, but the estimates remain noisy and are largely statistically insignificant. On average, transfers raised the HNI of second-born children by 0.254 units more than they raised the HNI of firstborns within the same household, relative to the control group. At baseline, second-borns were only marginally better off than firstborns by 0.032 units, a gap that narrowed by 0.065 among control households. Against this backdrop, the positive treatment estimates indicate that the CGP reversed the convergence observed among controls and further widened the gap in favour of second-borns, reinforcing their pre-existing advantage. The treatment effect on the BO2–BO1 gap was more pronounced at the 24- and 30-month follow-ups before attenuating at later waves.

The CGP did not appear to substantially affect the BO3–BO1 gap. However, we observe a marked deterioration in the HNI of fourth-or-higher-born children relative to their firstborn siblings, with a pooled estimate of -0.508. Among control households, the BO4+–BO1 gap widened considerably over time (pooled coefficient of 0.533, significant at the 1% level), implying that the negative treatment effects reflect preventative adjustments that offset the growing health and nutrition advantage of later-borns observed among non-recipients. The overall shape of the birth-order gradient is parabolic: the impact of the CGP on gaps relative to the firstborn initially increases with birth order before turning negative.

Consistent with findings from the preceding sections, we continue to observe substantial shifts in the way resources devoted to material consumption are allocated. With coefficient estimates of 0.220 and 0.251, respectively, Table 7 indicates that the CGP considerably strengthened parents' propensity to satisfy the material needs of third- and fourth-born children relative to the oldest sibling. Since BO1 children were traditionally better off at baseline, these patterns may likewise be characterised as a form of compensation layered on top of the milder compensation already visible among control households. However, while these effects substantially reduced and in some cases inverted the direction of inequality between BO1 children and their younger siblings, the magnitude of inequality was not markedly diminished – preferential allocation that initially favoured early-borns instead shifted to favour later-borns.

We also uncover some evidence of changes in schooling investment decisions resulting from the transfers. At the 36- and 48-month follow-ups, BO2 children in treated households were significantly less likely to be enrolled than firstborns, further amplifying the baseline enrolment gap. By contrast, the BO3–BO1 gap within treated households experienced sizeable reductions early in the programme before tapering off. When considered alongside shifts among non-recipients, who exhibited a substantial narrowing of the BO4+–BO1 enrolment gap but little change elsewhere along the birth-order gradient, the evidence suggests that the CGP modestly dampened the compensatory schooling investments that non-recipient households were already making in favour of later-borns.

We also explicitly model household-level dispersion using a two-step location-scale framework, which we describe in more detail in Section A.2 Two-Step Location-Scale Framework. The reason why we adopt a location-scale framework is that the treatment effect is better understood as a change in residual variance net of household-by-wave-specific common shocks, rather than as a change in the variance of raw outcomes. In Table A5, we present the second-stage estimates for the treatment effect on residual variance, with and without the BO interaction in the first stage. While we do not find notable impacts on HNI or enrolment, the CGP appears to reduce residual intra-household inequality in material well-being more strongly once birth-order effects are accounted for. To be fully consistent with our baseline results, which suggest the CGP led to overly compensatory investment strategies, the location-scale estimates would have to imply that, overall, the CGP equalised material access, but that the way gains differ substantially over the birth-order

Table 7: **OLS Estimates of the Birth-Order Impact of the CGP**

	HNI			MNI			Enrolled in school		
	BO2	BO3	BO4+	BO2	BO3	BO4+	BO2	BO3	BO4+
Panel A: Treated									
Wave-Specific Effect									
Treated × 24m × BO	0.406	0.077	-0.171	0.026	0.211**	0.246**	0.009	0.136**	-0.012
	(0.260)	(0.324)	(0.417)	(0.065)	(0.085)	(0.123)	(0.045)	(0.064)	(0.075)
Treated × 30m × BO	0.418	0.018	-0.419	0.082	0.213**	0.181	0.012	0.094	-0.067
	(0.274)	(0.325)	(0.399)	(0.071)	(0.087)	(0.134)	(0.038)	(0.064)	(0.079)
Treated × 36m × BO	-0.001	-0.188	-0.714*	0.025	0.218***	0.228*	-0.082*	0.018	-0.121
	(0.293)	(0.354)	(0.426)	(0.066)	(0.083)	(0.123)	(0.042)	(0.068)	(0.075)
Treated × 48m × BO	0.233	-0.177	-0.657	0.063	0.235***	0.327**	-0.082*	0.034	-0.060
	(0.304)	(0.375)	(0.424)	(0.061)	(0.087)	(0.133)	(0.043)	(0.071)	(0.076)
Pooled Effect									
Treated × Post × BO	0.254	-0.075	-0.508	0.050	0.220***	0.251**	-0.039	0.065	-0.066
	(0.259)	(0.310)	(0.375)	(0.060)	(0.078)	(0.119)	(0.034)	(0.059)	(0.067)
Panel B: Control									
Wave-Specific Effect									
Control × 24m × BO	-0.154	0.101	0.424*	0.102**	0.066	0.128	0.000	-0.024	0.161***
	(0.151)	(0.192)	(0.253)	(0.044)	(0.066)	(0.088)	(0.034)	(0.049)	(0.058)
Control × 30m × BO	-0.182	0.080	0.473**	0.064	0.037	0.108	-0.025	0.011	0.219***
	(0.168)	(0.189)	(0.228)	(0.050)	(0.065)	(0.090)	(0.027)	(0.046)	(0.059)
Control × 36m × BO	0.077	0.148	0.584***	0.085*	0.002	0.015	0.030	0.033	0.236***
	(0.156)	(0.180)	(0.211)	(0.045)	(0.061)	(0.088)	(0.031)	(0.050)	(0.055)
Control × 48m × BO	-0.018	0.188	0.646***	0.107**	0.027	-0.002	0.017	0.020	0.224***
	(0.166)	(0.192)	(0.210)	(0.047)	(0.069)	(0.102)	(0.036)	(0.057)	(0.062)
Pooled Effect									
Control × Post × BO	-0.065	0.129	0.533***	0.089**	0.031	0.053	0.007	0.012	0.214***
	(0.139)	(0.154)	(0.189)	(0.041)	(0.059)	(0.085)	(0.027)	(0.046)	(0.053)
Baseline mean	-0.030	0.139	0.115	-1.365	-1.371	-1.426	0.620	0.510	0.304
Baseline gap vs. BO1	0.032	0.201	0.176	-0.095	-0.100	-0.155	-0.046	-0.156	-0.362
Age range		0–7			5–17			5–17	
Mean depvar at baseline		0.074			-1.293			0.586	
Observations		16155			21224			21644	
Within R^2		0.032			0.061			0.242	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: BO1 is the reference birth-order category. Standard errors clustered at the CWAC level in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

gradient reintroduces significant intra-household inequality. A broader implication of the more negative coefficients after purging the birth-order effects is that a majority of the CGP-induced change in outcome variation within the household likely materialises through systematic differences by birth order, not through other sociodemographic factors.

6.3.2 Male- and Female-Specific Birth-Order Effects

Strategically restricting the sample and re-estimating the same fixed effects model again allows us to potentially pinpoint child and household types most or least at risk from the programme and tease out mechanisms of action. Given that the results of four-way interactions are difficult to make sense of, we opt for this as opposed to augmenting (4) with higher-order interactions.

We begin by estimating the same specifications for male-only and female-only subsamples and provide the results in Tables A6 and A7. Each regression requires that a household have at least two same-sex siblings at a given wave. At baseline, resource allocation within male and female maternal sibling groups was broadly similar, with later-borns generally more adequately nourished in both. Although statistically insignificant, the estimates for HNI suggest that the birth-order gradient in control households acquired the familiar inverted parabolic pattern over the follow-up waves. Treated households exhibited a similarly shaped gradient, heavily favouring second-borns while disfavouring fourth-borns.

For material well-being, the CGP elicited comparable birth-order effects for both male and female sibling groups. However, because baseline birth-order gaps in MNI were smaller among boys, the treatment effects translated into a more prominent increase in sibling inequality for male households.

We also document gendered redistribution in schooling following the CGP. The baseline BO4+–BO1 enrolment gap was notably larger among boys than among girls. Male-only control households compressed this gap substantially, but the CGP moderated this convergence, resulting in treated households closing the gap less than their non-recipient counterparts. The net effect is that post-treatment levels of intra-household schooling inequality across the three BO comparisons are nearly identical for male and female subsamples.

The absence of gendered effects may be surprising given that gender inequality remains a major development challenge in Zambia, one that might be expected to intensify in the extremely poor setting of the CGP. According to the 2010 Gender Status Report published by the government of Zambia's Ministry of Gender and Child Development (2011), there is a lack of young-age gender parity in secondary and tertiary education, reproductive health risks, formal employment, and access to productive resources. A peculiarity of our sample is that boys and girls were relatively equal in all three outcome domains at baseline[2], and our results are fully in line with these patterns.

To further demonstrate the absence of significant gender preference in our setting, we illustrate in Figure 4 that families possibly exhibit behaviour consistent with female-biased fertility stopping

[2]At baseline, girls have marginally higher HNI (0.153 vs. −0.008) and MNI (−1.261 vs. −1.326) than boys, and enrolment rates are nearly identical (59.0% vs. 58.3%).

Figure 4: **Fraction Male by Age of Child**

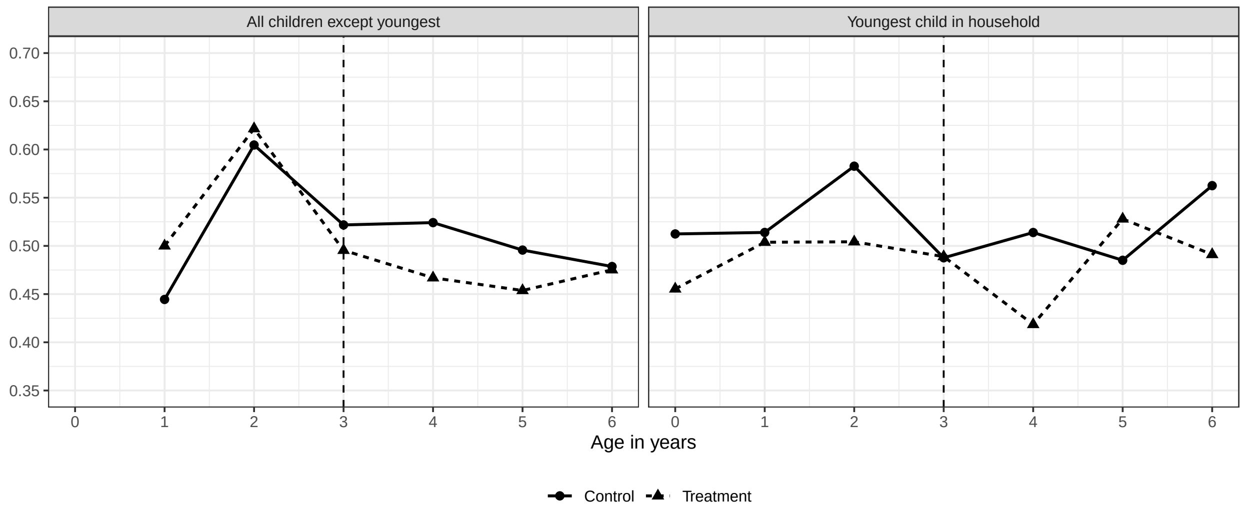

Note: This figure plots, separately for treatment and control households, the fraction of male children at the 48-month follow-up. The vertical line marks age three, children younger than this threshold were conceived during the CGP.

decisions. In the left panel, it can be seen that, compared to baseline, treated households had become more likely to have another child when their then youngest was a boy aged between three and five; in the right panel, treated families were more likely to stop having children once their youngest offspring was a girl. The differential patterns by treatment status suggest that recipients have a slightly stronger preference for girls while non-recipients prefer boys. We also find a marked spike in the proportion of children aged two who are male and are also not the youngest member of their household, implying that, towards the later years of our sample frame, families regardless of treatment status were more likely to continue having children when their most recent newborn had been male.

Following Jayachandran and Pande (2017), we test whether fertility stopping behaviour biased towards a particular gender, in this case female, leads to outcomes heterogeneous by gender. The specifications we estimate are described in Section A.3 Elder Brother. Overall, Table A8 shows that there is little evidence that the CGP led to differential outcomes due to the household's gender composition. We find that girls are significantly more likely to be enrolled when they have an elder brother, though this effect largely offsets a corresponding negative estimate among control households. Comparing outcomes by whether the child has an elder brother versus an elder sister, we do not obtain statistically significant treatment effects, although the point estimates suggest that the CGP marginally improved schooling outcomes for girls with elder brothers.

6.3.3 Most Constrained Quintile of Households

Table A9 presents estimates for the most constrained quintile of households. The results lend direct evidence corroborating the conjecture that households flexibly adopt alternative allocative schemes for different types of goods depending on their degree of financial hardship. We find that, despite a negligible gap at baseline, second-borns achieved substantial health and nutritional gains over their firstborn siblings. This exacerbation of intra-household inequality is especially notable because mean HNI in this subsample is only marginally worse than in the full sample, suggesting that the divergence arises from fundamentally different allocative approaches – ones consistent with efficiency maximisation – rather than from lower overall expenditures on food and healthcare.

We find that the CGP altered the allocation of material goods among the most constrained households differently from the rest of the sample. In Table A9, none of the MNI estimates attain statistical significance. We attribute this to the fact that these extremely constrained households lack the resources to invest in the material well-being of any child, thereby eliminating any scope for inequality. Their destitute conditions are reflected in their much lower baseline means.

We also find that the CGP-led reduction in the likelihood of fourth-borns or higher being enrolled in school relative to the firstborns is much more pronounced among the most constrained households. In particular, while the pooled BO4+–BO1 treatment effect on enrolment was -6.6 percentage points in the baseline sample, it widened to a substantial -21.1 percentage points among the most constrained households. These estimates suggest that the most constrained households opted to maintain the magnitude of intra-household inequality in schooling from baseline, a strategy we did not observe from estimates using the baseline sample. While we are not able to support this with any concrete proof, we hypothesise that from a heuristic standpoint that the CGP may have induced parents from the poorest of the poor to maximise the nutritional adequacy of later-borns, which was in part funded by resources that would have otherwise been allocated to meeting their material needs, and in part funded by those that would have otherwise been allocated to their schooling expenses.

6.3.4 Large Households

Higher fertility rates typically go hand in hand with the deprivation of human capital investment in each child due to the dilution of limited resources. Known as the child quantity-quality (QQ) trade-off (Guo et al., 2022), the phenomenon is especially relevant for the CGP as transfer size is set independently of household size. The QQ trade-off may also give rise to a spurious negative association between child outcomes and birth order – which we have shown is the case for at least material well-being and enrolment – since a higher BO necessarily implies a greater probability of belonging to a household with more siblings competing for the same resources.

In Table A10, we provide results obtained by estimating on a subset of households that at baseline had children more numerous than the sample median. Overall, we do not see clear differences in the ways households with more children behaved, with or without transfers, although non-recipients disproportionately raised their investment in schooling for later-borns during the follow-up waves. The lack of any conspicuous trade-off between the quantity and quality of children is surprising given that the transfer amount of the CGP did not vary with household size. We ascribe the anomaly to the fact that the baseline specification we re-estimate is linear in parameters and may not reveal complex relationships between quantity and quality, and to the fact that we condense all fourth-born-or-higher children into the same BO category.

6.3.5 Robustness Checks

Outcome-Order Effects Birth order is one of several ways to establish a form of comparison of siblings within the household, telling us where the child is situated in the sibling hierarchy. In order for the birth-order ranking to be consistent with a ranking that purely reflects how the CGP altered each child's outcome relative to the rest, the assumption that parents make distributional decisions based on the order of birth and nothing else (i.e., gender or genetic endowment) needs to be satisfied. To test this, we rank siblings based on their outcome at baseline using standard competition ranking[3].

Table A11 examines whether the CGP differentially affected children who were initially worse off within the household. The coefficient estimates for control households are large and positive across outcomes, implying substantial convergence to the mean over time where children who ranked below the best-performing sibling at baseline saw relatively greater improvements in later waves. For HNI and enrolment, the control group coefficients are near their corresponding gaps relative to rank 1 at baseline, which signifies compensation; while for MNI, the coefficients far exceeded the baseline gaps, signifying overcompensation and enlarged intra-household inequality. Regardless, the treatment effects are small and statistically insignificant across nearly all rank categories, suggesting that the CGP did not systematically enhance or dampen the convergence.

The birth-order results have indicated that treatment effects vary by sibling position, particularly for later-born children in some domains. The results from ranking by baseline outcome, which, by contrast, show that no effect was exerted by the CGP, suggest that these treatment differences are not simply explained by the programme promoting the parental targeting of children who were initially worse off. Instead, the birth-order effects may in fact reflect a more structural intra-household allocation pattern linked to circumstances specific to the household, survey wave, and wider societal context.

Sample Aging BO is mechanically related to age, so age effects can be masked as birth-order effects. Indeed, at baseline, our estimation sample was young, and one would expect the average age to have increased across the follow-up waves both at the aggregate level and at each level of BO. We show in Figure A2 that not only was this the case, but the increase in age had typically been greater for third-borns and fourth-or-higher-borns. Demand for food and healthcare, material goods, and schooling is age-dependent.

Hence, one could argue that the improvement to the material access of later-borns - the starkest result - attributable to the CGP was a product of those younger children reaching the age where having those material goods is more important. To counter this, we modify our baseline model in

[3] Standard competition ranking is a procedure for handling ties, which regularly occur for binary indicators or principal components constructed from binary indicators. For example, if two siblings are enrolled in school while another two are not, a ranking of 1, 1, 3, 3 would be assigned to the four-child maternal sibling group.

two ways: 1) residualise MNI and enrolment using non-parametric age splines and time and district fixed effects; 2) control for interactions of a full set of age dummies with time indicators. We conduct these exercises for models where we estimate a pooled treatment effect, the results for which are provided in Figure A3. We find that the results are largely identical to our baseline estimates.

We also modify our baseline specifications by replacing birth-order with age distance relative to the firstborn, which helps capture the local life-cycle channel and permits us to distinguish birth-order hierarchy from age spacing. It is doubly important as we have also shown that changes to allocative behaviour effected by the CGP occur through structurally differing sibling positions. Should the treatment effects vary with distance in age from the firstborn child, then the mechanism underlying our baseline results is possibly related to age-specific needs, schooling transitions, or developmental stages. If not, then our interpretation of the birth-order effects as the resource redistributive impact through structured rank of birth retains its validity.

We provide the results in Table A12. We find that each additional year in age gap with the firstborn is associated with an average increase of 0.021 units in MNI for the later-born following treatment. The practical significance of the point estimate is small, as the mean age gap from BO1 at baseline was only slightly over 3 years.

Cohort Sample We next confine the sample to a cohort of children for whom outcomes were recorded at baseline, so that the children contributing to the estimates in each period had been exposed to the CGP from the beginning of the programme up to that point in time. This restriction amplifies the magnitude of the baseline estimates. In the health and nutrition domain, however, this amplification is likely to be partially offset by the exclusion of newborns, for whom the CGP may have operated as a prenatal and early-childhood intervention with potentially heightened effects. We report the results in Table A13. As expected, the coefficient estimates are larger across the board, while the core redistributive patterns remain intact. The main deviation from the baseline estimates is the marked increase in school enrolment among later-born children relative to firstborn children. At baseline, BO4+ children were 36.2% less likely to be enrolled in school than firstborn children, an enormous schooling gap that had shifted in favour of fourth-born-or-higher children by 48 months. A plausible explanation is that, over the four-year period, firstborn children, who are older by construction, reached an age at which leaving school and entering work became more common, whereas later-born children had only just reached school age. We do not observe differences in this regard between treated and control households.

Households Perceiving the CGP as Conditional Notwithstanding that the CGP was advertised as an unconditional programme, some recipients incorrectly believed that the grant was tied to child welfare, such as buying food for children. The incidence of this misconception was quite high, but the majority faced no consequences for non-compliance, in that no repercussions would have ensued should they fail to satisfy the conditions. Therefore, we estimate on a subsample of households who believed the CGP to be conditional to assess whether holding such a misconception led to notable discrepancies in intra-household behaviour. Note that the operational performance of the CGP was retrospectively evaluated at survey waves 2, 4, and 5. Hence, causal estimates are likely biased as we are unable to adjust for baseline incomparability or, more critically, reverse causality, if past outcomes influence how the household interprets the programme. Results are provided in Table A14, which is mostly comparable to the baseline estimates. The coefficient estimates for HNI are generally more negative or less positive. This aside, however, we detect no systemic differences between the two sets of results.

Non-Graduated Households Finally, we provide in Table A15 results from estimations on a sample restricted to households that had not graduated in the wave in which they were observed. In line with our expectation, it can be seen that the magnitude of coefficient estimates is generally only slightly larger than their corresponding baseline estimates. This confirms the validity of our main specifications, in which we have not accounted for graduation from the CGP.

7. Channels

7.1 A Structural Approach

To better understand the underlying mechanisms that drive the results we obtain from the other approaches, we use an extended version of a model proposed by Behrman et al. (1982). The Behrman-Pollak-Taubman (henceforth BPT) model is an intra-household resource allocation model structurally derived from a household utility maximisation problem and is used to analyse family behaviour. To apply it to our setting, where we investigate changes in allocative behaviour in response to the CGP, we introduce several noteworthy modifications, which are expounded upon as we work through the derivation of the problem.

For exposition, we consider a household that is composed of a mother, denoted by m, a father, denoted by f, and multiple children, denoted by the subscripts 1 through K. We assume that the household utility function U has a Cobb-Douglas functional form, and that the parents' subutility functions are of the Constant Elasticity of Substitution form. It is further assumed that the children do not have decision-making power; they thus do not directly enter the household utility function, but indirectly enter via the subutility function of their parents, who derive utility from the children's human capital accumulation. The multi-input maximisation problem for this household can then be delineated as follows:

$$max \;\; U = u_m^{\mu} u_f^{1-\mu} \tag{5}$$

where

$$u_m = \left(\sum_{k=1}^{K} \alpha_k \, Q_k^{\rho} \right)^{1/\rho} \tag{6}$$

and

$$u_f = \left(\sum_{k=1}^{K} \beta_k \, Q_k^{\sigma} \right)^{1/\sigma}, \tag{7}$$

subject to the budget constraint, given as

$$M = \sum_{i=1}^{I} p_i \left(\sum_{k=1}^{K} I_{ik} \right), \tag{8}$$

and the child's human capital production function, given as

$$Q_k = A_k \prod_{i=1}^{I} I_{ik}^{c_i} \tag{9}$$

where A_k is the genetic endowment of child k that determines how effectively parental investment in input i translates into human capital Q, and

$$\sum_{i=1}^{I} c_i = 1. \tag{10}$$

It is immediately clear that any intra-household inequality among children arises from three sources of variation: relative bargaining power between the two parents, μ; the extent to which the parents exhibit discriminatory preferences for specific child types, α_k and β_k; and the extent to which the parents are averse to unequal child outcomes, ρ and σ. Therefore, it follows that the CGP alters the allocative behaviour of the household through three corresponding channels. In order to facilitate the derivation process, we have assumed away parents' private expenditures. On top of this, we also assume that parents have perfect information on each child's endowment, that the price for the same input is identical across children, and that there are also no spillover effects among them. The Marshallian demand for the quality of child k can then be derived as

$$Q_k^* = M \left[\mu \, \frac{\alpha_k^{\frac{1}{1-\rho}} P_{Q_k}{}^{\frac{-1}{1-\rho}}}{\sum_{j=1}^{K} \alpha_j^{\frac{1}{1-\rho}} P_{Q_j}{}^{\frac{-\rho}{1-\rho}}} + (1-\mu) \, \frac{\beta_k^{\frac{1}{1-\sigma}} P_{Q_k}{}^{\frac{-1}{1-\sigma}}}{\sum_{j=1}^{K} \beta_j^{\frac{1}{1-\sigma}} P_{Q_j}{}^{\frac{-\sigma}{1-\sigma}}} \right], \tag{11}$$

while the level of input i invested in that child is

$$I_{ik} = \frac{c_i}{p_{I_i}} P_{Q_k} Q_k^* = \frac{c_i M}{p_{I_i}} \left[\mu \, \frac{\alpha_k^{\frac{1}{1-\rho}} P_{Q_k}{}^{\frac{-\rho}{1-\rho}}}{\sum_{j=1}^{K} \alpha_j^{\frac{1}{1-\rho}} P_{Q_j}{}^{\frac{-\rho}{1-\rho}}} + (1-\mu) \, \frac{\beta_k^{\frac{1}{1-\sigma}} P_{Q_k}{}^{\frac{-\sigma}{1-\sigma}}}{\sum_{j=1}^{K} \beta_j^{\frac{1}{1-\sigma}} P_{Q_j}{}^{\frac{-\sigma}{1-\sigma}}} \right]. \tag{12}$$

Based on the observation that a majority of the households have one, not multiple, members making expenditure decisions, including those on both general household finances and investment in children's health, education, or material needs, we let μ be 0 or 1 depending on the gender of the decision maker. (11) and (12) then collapse to solutions equivalent to those derived from a unitary household problem. Stacking the human capital level of child k against that of child l, we have the structural equations of interest

$$\frac{Q_k^*}{Q_\ell^*} = \left(\frac{\alpha_k \, A_k}{\alpha_\ell \, A_\ell} \right)^{\frac{1}{1-\rho}} \tag{13}$$

$$\frac{I_{ik}}{I_{i\ell}} = \left(\frac{\alpha_k}{\alpha_\ell}\right)^{\frac{1}{1-\rho}} \left(\frac{A_k}{A_\ell}\right)^{\frac{\rho}{1-\rho}}. \tag{14}$$

with $\alpha_k + \alpha_l = 1$. The aim is to log-linearise and estimate via non-linear least squares (NLS) the structural parameters α_k (α_l, which equals $1 - \alpha_k$, does not need to be estimated) and ρ, and allow them to vary by treatment assignment and ideally other covariate values; for instance, $\alpha_k = \alpha_0 + \alpha_1 Treat + \varepsilon$ and $\rho = \rho_0 + \rho_1 Treat + \upsilon$. However, it is not feasible to estimate the two parameters as functions of covariates while simultaneously enforcing the inequalities of $0 < \alpha_k < 1$ and $\rho < 1$, which, as one may recall, are necessary for interpretability. It is therefore imperative that we reparametrise α_k and ρ in ways such that the necessity of explicitly imposing the bounds is relinquished. To achieve this, we rewrite the parameters as $\eta = \log\left(\frac{a}{1-a}\right)$ and $\kappa = \log(1-\rho)$, so that $\alpha = \frac{1}{1+e^{-\eta}} \in (0,1)$ and $\rho = 1 - e^{\kappa} < 1$.

The non-linear least squares problem now with proper subscripts is

$$\log\left(\frac{HNI_{kht}}{HNI_{lht}}\right) = \frac{\log\left(\frac{A_{kht}}{A_{lht}}\right) + \eta_{klht}}{\exp(\kappa_{ht})} + \varepsilon_{klht}, \tag{15}$$

with

$$\eta_{klht} = \eta_0 + \eta_1 \text{Treat}_j + \eta_2 \text{Time}_t + \eta_3 \left(\text{Treat}_j \times \text{Time}_t\right) + \sum_{m=1}^{M} \gamma X_{klht} + \sum_{r=1}^{R} \pi Z_h \equiv \log\left(\frac{\alpha_{klht}}{1-\alpha_{klht}}\right) \tag{16}$$

and

$$\kappa_{ht} = \kappa_0 + \kappa_1 \text{Treat}_j + \kappa_2 \text{Time}_t + \kappa_3 \left(\text{Treat}_j \times \text{Time}_t\right) + \sum_{r=1}^{R} \omega Z_h \equiv \log(1-\rho_{ht}), \tag{17}$$

for all $(k,l) \in \mathcal{P}_{ht}$, where $\mathcal{P}_{ht} = \{(k,l) : k \neq l,\ age_{kht} > age_{lht}\}$. We estimate (15) on a sample constructed from all possible dyads in dictatorial households, to the extent that a single member makes decisions regarding child investment at baseline, in household h at survey wave t, such that the age of the numerator child k is greater than the age of the denominator child l. This means that the results from the NLS estimations of the extended BPT model can serve as tentative evidence on the mechanisms driving our findings from the preceding fixed effects approach. We also only estimate the model for HNI, as dyadic response variables constructed from binary indicators do not provide sufficient variation. Moreover, since we hypothesised that one of the three channels through which the CGP may have resulted in resource redistribution concerns the designation of women as transfer recipients, we estimate (16) and (17) with and without $Female\ decision-maker$ as a covariate to test whether the empowerment of women, who tend to be more inequality averse and child-type indifferent, enhances the effectiveness of the CGP at mitigating intra-household inequality. Significant differences in estimates between the two specifications would suggest that women's empowerment indeed mediates the relationship between transfers and allocative strategies. For both specifications, X includes the dyad's mean age and age range, while Z includes the baseline values of household expenditure per capita, size, average age, and gender composition, and the age, highest grade, and marital status of the recipient.

Genetic endowment, A, is proxied by a child development index that we construct from 15 items in the survey questionnaire that enquired whether the child was observed to have performed activities positively pertinent to the development of their stock of motor and cognitive skills. Given inconsistencies in the questionnaire items over the survey waves, the index can only be computed for post-treatment periods 4 and 5; the estimation is consequently analogous to a joint estimation of the pair of single differences two and four years following the onset of the CGP. Because household size is informative of the estimand of interest, we elect to include all pairwise comparisons of siblings who are aged between 2 and 10 within the household as opposed to arbitrarily nominating a specific pair from each household or subsetting to households carrying exactly two children. We impose the age range because child development items in the survey questionnaire are observed at large only for those no younger than 2 and no older than 10. We also two-way cluster the standard errors at both the household level and the CWAC level to account for inter-dyad correlation, where one child may be part of multiple pairs, as well as for intra-CWAC correlation, the level of randomisation.

Next, we discuss the transformations applied to HNI and A. Foremost, child development is mechanically increasing in age; it is therefore by construction that the A ratio on average is greater than 1. As such, we residualise A by age and time and district fixed effects. Naturally, non-positive values are not permitted in the variables due to the log transformation; logically, allowing such values would also disallow any meaningful interpretation, as when either the numerator or denominator can be negative, the sign and magnitude of the ratio no longer map cleanly into a meaningful comparison. We rescale both HNI and residualised A by first adding their respective global minimum values, then a constant. Our preferred choice for the constant is 1, though we also conduct robustness checks using alternative values.

Because the relation between HNI and A is non-linear, and changes for different values of A, we leverage only the sign and statistical significance when interpreting parameter estimates. Further, we are unable to rule out the possibility that A is not a measure of primitive endowment that is orthogonal to past treatment. Instead, it reflects the cumulative and endogenous interaction of innate abilities, parental inputs, and health shocks. A also likely contains both classical and non-classical measurement error, due to contamination by parents' subjective beliefs and perceptions that are also guided by their own traits. Accordingly, we proceed with caution and interpret these results as descriptive rather than causal.

7.2 Results

We report abridged results on the impact of experiencing the CGP on inequality aversion and child preference in Table 8 and detailed results in Table A16. In the specification where we include $Female\ decision-maker$, recipients are associated with higher inequality aversion concerning child health and nutrition by 3.65 units at the 36-month follow-up wave, but are not different from non-recipients in child preference. A year later, at the 48-month mark, we find that recipients are instead significantly less averse to inequality by around 4.35 units – the sum of $Treated$ and $Treated \times 48m$ – but remain indifferent toward specific child types. Additionally, we find that the distribution of HNI in households whose primary female caretaker had been the sole decision-maker concerning child health is substantially more equitable by 7.86 units. In line with our previous findings, parents pay only minor attention to gender in formulating their discriminatory preferences regarding child type. However, how gender impacts equality depends on birth order. When the older, lower-BO child is female, α increases by 0.028 units. Considering an average value of 0.467 for α across all households, we interpret the effect as improving intra-household equality. When the younger, higher-BO child is female, equality deteriorates by 0.021 units. Recall that α denotes the preference for the elder sibling from the dyad; both estimates indicate that girls have slightly better outcomes than boys.

In the specification where we omit $Female\ decision-maker$, the portion of CGP-led resource redistribution within the household operating through women's empowerment is absorbed into the treatment effects. We find that coefficient estimates for ρ attenuated towards 0, appearing both positive at 36 months and, slightly more so, less negative at 48 months. This implies that women as decision-makers marginally moderated the 36-month effect of the CGP on improving

Table 8: **Marginal Effects of the CGP on Child Preference and Inequality Aversion (Extended Results)**

	Child preference α		Inequality aversion ρ	
	(1)	(2)	(1)	(2)
Intercept (36m, control)	-0.042^{**}	-0.056^{**}	-14.114^{***}	-15.295^{***}
	(0.020)	(0.024)	(2.875)	(3.117)
Treated	0.024	0.041^{*}	3.652^{**}	3.290^{*}
	(0.016)	(0.022)	(1.652)	(1.851)
48m	0.006	0.005	3.349^{**}	2.995^{*}
	(0.016)	(0.019)	(1.537)	(1.753)
Treated × 48m	-0.022	-0.019	-8.002^{***}	-6.322^{***}
	(0.021)	(0.024)	(2.085)	(2.214)
Female decision-maker	-0.024		7.860^{***}	
	(0.017)		(2.254)	
Older child female	0.028^{**}	0.015		
	(0.012)	(0.015)		
Younger child female	-0.021^{*}	-0.036^{*}		
	(0.012)	(0.019)		
Observations	2777	2777	2777	2777
Residual standard error	0.216	0.217	0.216	0.217
Convergence iterations	43	49	43	49
Mean	0.467	0.450	-6.158	-6.456
Median	0.466	0.450	-4.518	-5.462
Minimum	0.294	0.219	-42.908	-29.567
Maximum	0.635	0.668	-0.296	-0.592

Notes: Standard errors two-way clustered at the household and CWAC levels in parentheses. Column (1) reports estimated marginal effects with $Female\ decision-maker$ included – an indicator that is coded as 1 if the primary female care-taker is the household's sole decision-maker regarding child health and nutrition, or 0 if her husband is involved in decision-making. Column (2) reports estimated marginal effects without $Female\ decision-maker$. Dyadic-specific variables, which include age and gender, are not entered into the inequality aversion equation so as to ensure the parameter varies at the household level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

equality and, conversely, mitigated the 48-month effect of the CGP on deteriorating equality. We also find some evidence suggesting that, with the coefficient estimate for α at 36 months increasing from 0.024 to 0.041 once $Female\ decision - maker$ is eliminated, households in which women were the decision-makers exhibited a slightly stronger preference for later-born children. However, the size of the difference in coefficient estimates with and without $Female\ decision - maker$ pales in comparison to the magnitude of the coefficient estimate on $Female\ decision - maker$ itself, suggesting that rapid, broader societal changes that elevated women's bargaining power in treated and control households alike played a much bigger role than the CGP raising the relative income contribution of female recipients.

Chen and Roth (2024) have demonstrated that an elasticity estimate of any magnitude is attainable through simply rescaling the outcome variable by adding some constant ϵ before log transformation. Hence, we illustrate in Figure 5 coefficient estimates under alternative ϵ values. We find that the magnitude of estimates is largest for smaller ϵ. This is not surprising, as a smaller ϵ in turn means a smaller denominator value in the construction of log-transformed dyadic HNI. Both the sign and statistical significance of the treatment effects remain consistent for all values of ϵ we test and for both specifications we run.

Figure 5: **Marginal Effects Using Alternative ϵ Values**

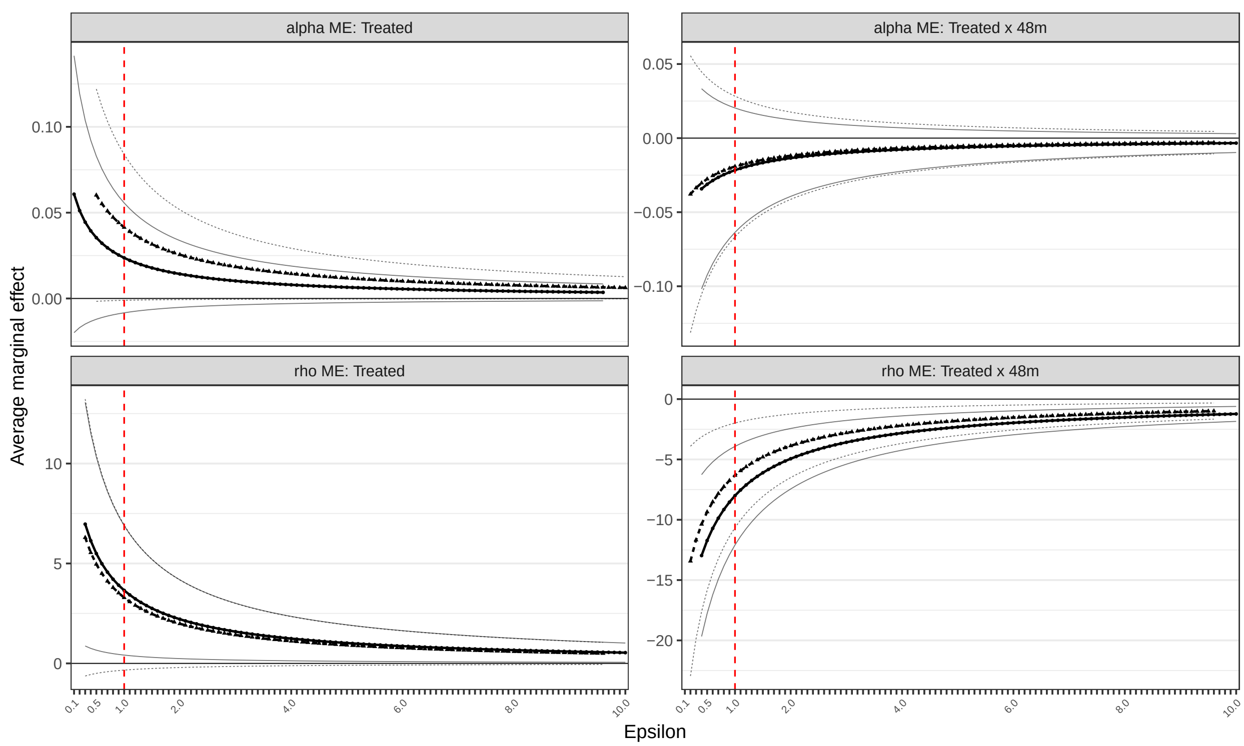

Note: Solid line corresponds to marginal effects recovered from estimations with female decision-maker indicator, while dashed line corresponds to marginal effects recovered from estimations without it. Vertical red dashed line marks $\epsilon = 1$.

8. Concluding Remarks

Together with their conditional counterparts, UCT programmes have become one of the fundamental building blocks of most social welfare protection systems, and their effectiveness at improving a wide array of outcomes has been documented in a wide range of studies. Although most cash schemes are designed to target women as recipients as a means to empower them within the household, there is little empirical evidence explicitly informing us of the intra-household allocative responses induced by transfers. Even less is known about the impact on inequality among siblings within the same household, especially of cash schemes that do not impose restrictions on how transfers are to be spent.

We formally examine the effect on intra-household inequality among siblings of a randomised UCT programme implemented in extremely poor Zambia. Overall, four years of exposure to the programme did not result in drastic changes in health and nutrition or schooling, but access to material goods saw sizeable improvements. Against this backdrop, we examine whether the CGP induced significant resource redistribution within the household that could partly explain the absence of clear programme impact in some outcome domains, and whether the observed redistribution follows patterns that could have perverse implications for domains where programme impact is documented.

We begin with a decomposition of variance into between- and within-household components, and find that absolute inequality within the household had decreased following the treatment, but, overweighed by a comparatively much larger reduction in absolute inequality between households, the proportion by which the within-household component contributes to total inequality had increased as a result of the transfers. We also observe evidence supporting that different allocative schemes are followed for different types of goods. In particular, schooling inequality is dominated by discrepancies within the household, whereas material-needs inequality is dominated by those occurring between households.

Next, we employ a difference-in-differences strategy augmented with household-by-time fixed effects, identifying the change in outcome heterogeneity along the birth-order gradient arising from the CGP. Our estimation results suggest that the CGP increased, by a substantial margin, parents' inclination to meet the material needs of later-borns, and the resource redistribution overcompensates for the baseline gap between them and the firstborn. In contrast, the CGP offset the substantial widening of the health and nutrition gap between firstborns and later-borns that was observed among control households, a pattern consistent with preventative adjustments in allocative behaviour. We do not find any consistent pattern in how transfers altered the birth-order gradient of enrolment rates. During the first two years of programme exposure, the enrolment advantage of third-borns over firstborns was substantially larger among treated households than among controls. After three to four years of exposure, we instead document schooling losses for the second-born. These findings are mostly robust to a number of subsample analyses.

Our analysis on gender suggests that, despite the study being set in severely impoverished regions where women face persistent disadvantages, households may have adopted fertility stopping decisions that appear to be biased towards girls. We also find that households tend to exhibit stronger compensatory actions concerning inequality among female offspring, although on the whole we do not find gender-specific responses. We also look at whether the child quantity-quality trade-off also applies to households in our study sample. Overall, there is no evidence that households with more children allocated resources along the birth-order gradient differentially between the treatment and control groups. Per-child investment for each of the three outcome domains is also mostly independent of the number of children.

Last but not least, we estimate an extended Behrman-Pollak-Taubman-type structural model to ascertain the drivers of the aforementioned findings. We show that the CGP increased parents' inequality aversion with respect to child health and nutrition at 36 months, but reduced it at 48 months. We also document some evidence suggesting parental preference for girls and later-borns, which is consistent with our results from other methods. Women's empowerment, the third mechanism of action, refers to the possibility that, by designating women – who tend to be more inequality averse and child-type indifferent – as recipients, the CGP may indirectly affect resource distribution within the household. We find that, while dyadic outcomes were more equitable if the primary female caretaker of a household was the sole decision-maker on child health and nutrition, only a small portion of the total effect occurs through the empowerment of women's decision-making power.

A distinct feature of our study sample is that girls are better off than boys, an anomaly that diminishes the generalisability of our findings. Another caveat of our empirical strategy is that we only assess sibling inequality arising over the birth-order gradient. This framing is necessary for forming comparisons within the household, but it means that we are unable to claim that similar patterns of resource redistribution along other child traits would be induced by transfers. We are also prevented from making causal interpretations based on our structural estimates due to the presence of endogeneity. Thus, further studies will be necessary to offer a more holistic, broadly applicable view of the intra-household impact of unconditional poverty transfers, and to more thoroughly uncover the mechanisms driving changes in parents' allocative behaviour.

References

Abufhele, A., Behrman, J. R., & Bravo, D. (2017). Parental preferences and allocations of investments in children's learning and health within families. *Social Science & Medicine, 194*, 76–86.

Aizer, A., & Cunha, F. (2012). The production of human capital: Endowments, investments and fertility. *NBER Working Paper No. 18429*. National Bureau of Economic Research.

Alderman, H., Chiappori, P.-A., Haddad, L., Hoddinott, J., & Kanbur, R. (1995). Unitary versus collective models of the household: Is it time to shift the burden of proof? *The World Bank Research Observer, 10*(1), 1–19.

Almond, D., Edlund, L., & Milligan, K. (2013). Son preference and the persistence of culture: Evidence from South and East Asian immigrants to Canada. *Population and Development Review, 39*(1), 75–95.

American Institutes for Research. (2016). *Zambia's Child Grant Program: 48-month impact report*. American Institutes for Research.

Armand, A., Attanasio, O., Carneiro, P., & Lechene, V. (2020). The effect of gender-targeted conditional cash transfers on household expenditures: Evidence from a randomized experiment. *The Economic Journal, 130*(631), 1875–1897.

Banks, J., Blundell, R., & Lewbel, A. (1997). Quadratic Engel curves and consumer demand. *The Review of Economics and Statistics, 79*(4), 527–539.

Barcellos, S. H., Carvalho, L. S., & Lleras-Muney, A. (2014). Child gender and parental investments in India: Are boys and girls treated differently? *American Economic Journal: Applied Economics, 6*(1), 157–189.

Becker, G. S. (1974). A theory of social interactions. *Journal of Political Economy, 82*(6), 1063–1093.

Becker, G. S. (1975). *Human capital: A theoretical and empirical analysis, with special reference to education* (2nd ed.). Columbia University Press.

Behrman, J. R. (1988). Intrahousehold allocation of nutrients in rural India: Are boys favored? Do parents exhibit inequality aversion? *Oxford Economic Papers, 40*(1), 32–54.

Behrman, J. R., Pollak, R. A., & Taubman, P. (1982). Parental preferences and provision for progeny. *Journal of Political Economy, 90*(1), 52–73.

Behrman, J. R., Pollak, R. A., & Taubman, P. (1986). Do parents favor boys? *International Economic Review, 27*(1), 31–52.

Berman, N., Rotunno, L., & Ziparo, R. (2025). Sweet child of mine: Parental income, child health and inequality. *The Economic Journal*.

Breen, R., & Ermisch, J. (2021). Estimating sibling models for categorical outcomes. *Sociological Methods & Research, 50*(2), 491–521.

Bharadwaj, P., & Lakdawala, L. K. (2013). Discrimination begins in the womb: Evidence of sex-selective prenatal investments. *Journal of Human Resources, 48*(1), 71–113.

Björkman-Nyqvist, M. (2013). Income shocks and gender gaps in education: Evidence from Uganda. *Journal of Development Economics, 105*, 237–253.

Bobonis, G. J. (2009). Is the allocation of resources within the household efficient? New evidence from a randomized experiment. *Journal of Political Economy, 117*(3), 453–503.

Bonilla, J., Zarzur, R. C., Handa, S., Nowlin, C., Peterman, A., Ring, H., Seidenfeld, D., & the Zambia Cash Transfer Evaluation Team. (2017). Cash for women's empowerment? A mixed-methods evaluation of the government of Zambia's Child Grant Program. *World Development, 95*, 55–72.

Bourguignon, F., Browning, M., & Chiappori, P.-A. (2009). Efficient intra-household allocations and distribution factors: Implications and identification. *The Review of Economic Studies, 76*(2), 503–528.

Browning, M., Bourguignon, F., Chiappori, P.-A., & Lechene, V. (1994). Income and outcomes: A structural model of intrahousehold allocation. *Journal of Political Economy, 102*(6), 1067–1096.

Browning, M., Chiappori, P.-A., & Lewbel, A. (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *The Review of Economic Studies, 80*(4), 1267–1303.

Browning, M., Chiappori, P.-A., & Weiss, Y. (2014). *Economics of the family*. Cambridge University Press.

Browning, M., Bourguignon, F., Chiappori, P.-A., & Lechene, V. (1990). A collective approach to household behavior. Working paper.

Carolina Population Center. (2025). *Zambia Child Grant Programme (CGP)* (Version 1) [Data set]. UNC Dataverse. https://doi.org/10.15139/S3/Z7WF9D

Casco, J. L. (2023). The effects of cash and in-kind transfers on intra-household inequality: Insights from a randomized experiment. *Economics Bulletin, 43*(3), 1519–1527.

Casella, G., & Berger, R. L. (2002). *Statistical inference* (2nd ed.). Duxbury.

Chen, J., & Roth, J. (2024). Logs with zeros? Some problems and solutions. *The Quarterly Journal of Economics, 139*(2), 891–936.

Chiappori, P.-A. (1991). Nash-bargained household decisions: A comment. *International Economic Review, 32*(4), 791–796.

Chiappori, P.-A., & Meghir, C. (2015). Intrahousehold inequality. In A. B. Atkinson & F. Bourguignon (Eds.), *Handbook of income distribution* (Vol. 2, pp. 1369–1418). Elsevier.

Cunha, F., & Heckman, J. J. (2007). The technology of skill formation. *American Economic Review, 97*(2), 31–47.

De Chaisemartin, C., & D'Haultfoeuille, X. (2018). Fuzzy differences-in-differences. *The Review of Economic Studies, 85*(2), 999–1028.

De Rock, B., Vermeulen, F., & others. (2022). Efficient intra-household allocations and distribution factors: Implications and identification. Working paper.

De Vreyer, P., & Lambert, S. (2021). Inequality, poverty, and the intra-household allocation of consumption in Senegal. *The World Bank Economic Review, 35*(2), 414–435.

Duflo, E. (2003). Grandmothers and granddaughters: Old-age pensions and intrahousehold allocation in South Africa. *The World Bank Economic Review, 17*(1), 1–25.

Dunbar, G. R., Lewbel, A., & Pendakur, K. (2013). Children's resources in collective households: Identification, estimation, and an application to child poverty in Malawi. *American Economic Review, 103*(1), 438–471.

Fujiwara, T., & Kawachi, I. (2009). Is education causally related to better health? A twin fixed-effect study in the USA. *International Journal of Epidemiology, 38*(5), 1310–1322.

Folbre, N. (1986). Hearts and spades: Paradigms of household economics. *World Development, 14*(2), 245–255.

Frijters, P., Johnston, D. W., Shah, M., & Shields, M. A. (2010). To work or not to work? Child development and maternal labor supply. Working paper.

Giannola, M. (2022). Parental investments and intra-household inequality in child human capital: Evidence from a lab-in-the-field experiment. Working paper.

Guo, R., Yi, J., & Zhang, J. (2022). The child quantity–quality trade-off. *IZA Discussion Paper No. 15233*. IZA Institute of Labor Economics.

Handa, S., Natali, L., Seidenfeld, D., Tembo, G., Davis, B., & the Zambia Cash Transfer Evaluation Team. (2018). Can unconditional cash transfers raise long-term living standards? Evidence from Zambia. *Journal of Development Economics, 133*, 42–65.

Handa, S., Peterman, A., Seidenfeld, D., & Tembo, G. (2019). Income transfers and maternal health: Evidence from a national randomized social cash transfer program in Zambia. *Health Economics, 28*(2), 175–188.

Handa, S., Natali, L., Seidenfeld, D., Tembo, G., Davis, B., & the Zambia Cash Transfer Evaluation Team. (2025). In search of the holy grail: Post-intervention effects of an unconditional cash transfer program in Zambia. *Journal of Development Economics*.

Jayachandran, S., & Pande, R. (2017). Why are Indian children so short? The role of birth order and son preference. *American Economic Review, 107*(9), 2600–2629.

Jose, A. (2024). *Income shocks and intrahousehold resource allocation: Evidence from rural Ethiopia*. SSRN Working Paper.

Kovas, Y., Voronin, I., Kaydalov, A., Malykh, S. B., Dale, P. S., & Plomin, R. (2013). Literacy and numeracy are more heritable than intelligence in primary school. *Psychological Science, 24*(10), 2048–2056.

Lamont, J., & Favor, C. (2017). Distributive justice. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Stanford University.

Le, K., & Nguyen, M. (2022). Son preference and health disparities in developing countries. *SSM – Population Health, 17*, 101036.

Lechene, V., Pendakur, K., & Wolf, A. (2021). *OLS estimation of the intra-household distribution of expenditure*. CEPR Discussion Paper No. DP15955.

Lewbel, A., & Pendakur, K. (2008). Estimation of collective household models with Engel curves. *Journal of Econometrics, 147*(2), 350–358.

Lundberg, S., & Pollak, R. A. (1993). Separate spheres bargaining and the marriage market. *Journal of Political Economy, 101*(6), 988–1010.

Lundberg, S. J., Pollak, R. A., & Wales, T. J. (1997). Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit. *Journal of Human Resources, 32*(3), 463–480.

Malghan, D., & Swaminathan, H. (2021). Global trends in intra-household gender inequality. *Journal of Economic Behavior & Organization, 189*, 515–546.

Mangiavacchi, L., & Piccoli, L. (2011). Improving child nutrition status in rural Albania: Does income matter? Working paper.

Ministry of Gender and Child Development. (2011). *2010 Gender Status Report: A baseline review.* Government of the Republic of Zambia.

Murphy, K. M., & Topel, R. H. (1985). Estimation and Inference in Two-Step Econometric Models. *Journal of Business & Economic Statistics, 3*(4), 370–379.

Pagan, A. (1984). Econometric issues in the analysis of regressions with generated regressors. *International Economic Review, 25*(1), 221–247.

Quisumbing, A. R., & Maluccio, J. A. (2003). Resources at marriage and intrahousehold allocation: Evidence from Bangladesh, Ethiopia, Indonesia, and South Africa. *Oxford Bulletin of Economics and Statistics, 65*(3), 283–327.

Siow, A. (2015). Testing Becker's theory of positive assortative matching. *Journal of Labor Economics, 33*(2), 409–441.

Thomas, D. (1990). Intra-household resource allocation: An inferential approach. *Journal of Human Resources, 25*(4), 635–664.

Weaver, J., Sukhtankar, S., Niehaus, P., & Muralidharan, K. (2024). *Maternal cash transfers for gender equity and child development: Experimental evidence from India.* NBER Working Paper No. 32093. National Bureau of Economic Research.

Yi, J., Heckman, J. J., Zhang, J., & Conti, G. (2015). Early health shocks, intrahousehold resource allocation, and child outcomes. *The Economic Journal, 125*(588), F347–F371.

Appendix

A.1 ATT Estimates of the Impact of the CGP

Because programme graduation was not uncommon, we follow de Chaisemartin and D'Haultfoeuille (2018) and estimate an instrumented DID model to take into account graduation from the CGP. In particular, we construct a wave-specific treatment indicator, $Eligible_{ht}$, that equals 1 if household h was receiving transfers from the CGP at time t, and 0 otherwise, thus permitting treated households to transition back to non-recipient status upon programme graduation. Because graduation timing is plausibly endogenous, the average treatment effect on the treated (ATT), or the causal effect of actually receiving the transfer, is estimated using a two-stage least squares DID framework, where actual receipt $Eligible_{ht}$ is instrumented by the interaction between randomised treatment assignment and time. De Chaisemartin and D'Haultfoeuille note that, algebraically, the ATT point estimates are equivalent to dividing the ITT point estimates by the actual treatment receipt rate of units initially assigned to treatment. The system of regression equations is as follows:

$$Eligible_{ht} = \pi\,(Treat_j \times Post_t) + \sum_{k=1}^{K} \theta_k X_{it} + \sum_{l=1}^{L} \rho_l Z_h + \alpha_d + \nu_{ihjdt}. \tag{18}$$

$$Y_{ihjdt} = \beta_1\, Treat_j + \beta_2\, Post_t + \beta_3\, \widehat{Eligible}_{ht} + \sum_{k=1}^{K} \delta_k X_{it} + \sum_{l=1}^{L} \lambda_l Z_h + \alpha_d + \varepsilon_{ihjdt}, \tag{19}$$

where $Post_t$ is an indicator equal to 1 for all follow-up survey waves and 0 for the baseline wave, and thus the assignment-post interaction in (18) isolates the exogenous portion of variation in actual receipt. In both stages, we control for the same set of covariates as in (1), as well as district fixed effects. It is also worth keeping in mind that we only have an imperfect approximation of the graduation date of each household, proxied by the age of the youngest member in the household. This is because, as per Handa et al. (2025), exit from the CGP was formally administered at irregular intervals rather than continuously, often allowing households to remain eligible for transfers for several months more than intended. For comparability, we also provide estimation results using an alternative specification for (1), where we switch $Time_t$ out for $Post_t$, thereby yielding a pooled treatment effect estimate across the four endline waves.

We report the results in Table A3. ATT estimates are between 18% and 23% larger than the corresponding ITT pooled effect estimates in terms of magnitude, as expected when the ITT effect is diluted by discontinued treatment exposure. However, the differences are modest and the statistical significance of the estimates is broadly unchanged, even though a non-trivial share of households in the treatment group were no longer eligible for transfers toward the end of the evaluation period. To see how the ITT and ATT estimates relate, first compute the weighted average of actual treatment receipt rates across the four follow-up survey waves, which comes out to be around 84.6% for HNI, 82.3% for MNI, and 82.4% for enrolment. Then, dividing the pooled ITT point estimates of 0.048, 0.665, and 0.035 by these proportions, we attain 0.0567, 0.808, and 0.0425, which are approximately equal to the ATT estimates.

In the same vein, ITT and ATT estimates for the constituent variables, which have been provided respectively in Table A2 and Table A4, are also highly comparable. The similarity between the ITT and ATT estimates demonstrates that declining programme compliance did not materially attenuate the estimated effects of assignment to treatment, lending credibility to the use of treatment assignment as our preferred treatment indicator throughout the rest of the study.

A.2 Two-Step Location-Scale Framework

We complement the fixed effects model, where the unit of analysis is the individual, by first extracting the residuals and then aggregating them to the household level by computing the variance in residuals for observations from the same household:

$$\hat{\sigma}^2_{hjdt} = \frac{1}{N_{hjdt} - 1} \sum_{i=1}^{N_{hjdt}} \left(\hat{\varepsilon}_{ihjdt} - \bar{\hat{\varepsilon}}_{hjdt} \right)^2 \tag{20}$$

where N_{hjdt} is the number of dyads contributed by household h from community j in district d at time t. Net of the conditional household-time mean in (4), $\hat{\sigma}^2_{hjdt}$ reflects the variation of sibling outcomes about the modelled household mean structure. We then estimate the log-linear equation below to identify the impact of the programme on residual dispersion within the household.

$$\log\left(\hat{\sigma}^2_{hjdt}\right) = \beta_1 \,\text{Treat}_j + \beta_2 \,\text{Time}_t + \beta_3 \left(\text{Treat}_j \times \text{Time}_t\right) + \sum_{l=1}^{L} \lambda_l Z_h + \mu_d + u_{hjt} \tag{21}$$

Together, the two-step procedure constitutes a location-scale framework, in which the first stage models the mean and the second stage models within-household dispersion around that mean. Furthermore, with the three-way interaction term in the mean model, the residuals ε_{ihjdt} and, by extension, the coefficient estimates of the treatment effect on ε_{ihjdt} reflect intra-household inequality after purging the portion attributable to birth order. Thus, we repeat the procedure twice, once with birth order in the first-stage three-way interaction, whereby the second stage estimates the treatment effect on residual intra-household inequality net of birth-order effects, and once more without the interaction altogether so that the second stage estimates the treatment effect on total intra-household inequality.

Because $\hat{\sigma}^2_{hjdt}$ is a measure generated from a regression, its sampling error is not captured by the standard errors of the second-stage regression, leading to overly liberal inference. The problems associated with regressions using generated regressors have been discussed in Pagan (1984), and the remedy in the classic two-step variance correction in Murphy and Topel (1985). To ensure that second-stage standard errors account for first-stage estimation, we cluster bootstrap both stages of the procedure with replacement at the community level. We expect the second-stage results to be similar to those obtained by straightforwardly modelling the variance of raw outcomes among siblings within the same household as a function of treatment assignment, owing to the fact that variance is invariant to the translation effects of within-household and within-time demeaning.

A.3 Elder Brother

We estimate the following equation

$$\begin{aligned} Y_{ihjdt} = {} & \gamma_1 \left(\text{Treat}_j \times \text{Post}_t \times \text{Girl}_i\right) + \gamma_2 \left(\text{Treat}_j \times \text{Post}_t \times \text{NoElderBrother}_i\right) \\ & + \gamma_3 \left(\text{Treat}_j \times \text{Post}_t \times \text{Girl}_i \times \text{NoElderBrother}_i\right) + \sum_{k=1}^{K} \theta_k X_{it} + \alpha_{ht} + \mu_d + \varepsilon_{ihjdt}. \end{aligned} \tag{22}$$

where $NoElderBrother_i$ takes the value 1 if a child does not have any elder brothers. The interaction between $treat_j$, $post_t$, and $girl_i$ captures, across the follow-up waves, whether girls with an elder brother responded differently to the CGP than boys with an elder brother. The interaction between $treat_j$, $post_t$, and $NoElderBrother_i$ captures whether boys with an elder brother responded to the CGP differently than boys with an elder sister; and the four-way interaction is analogous for girls.

A.4 Figures

Figure A1: **Baseline Demographic Distribution**

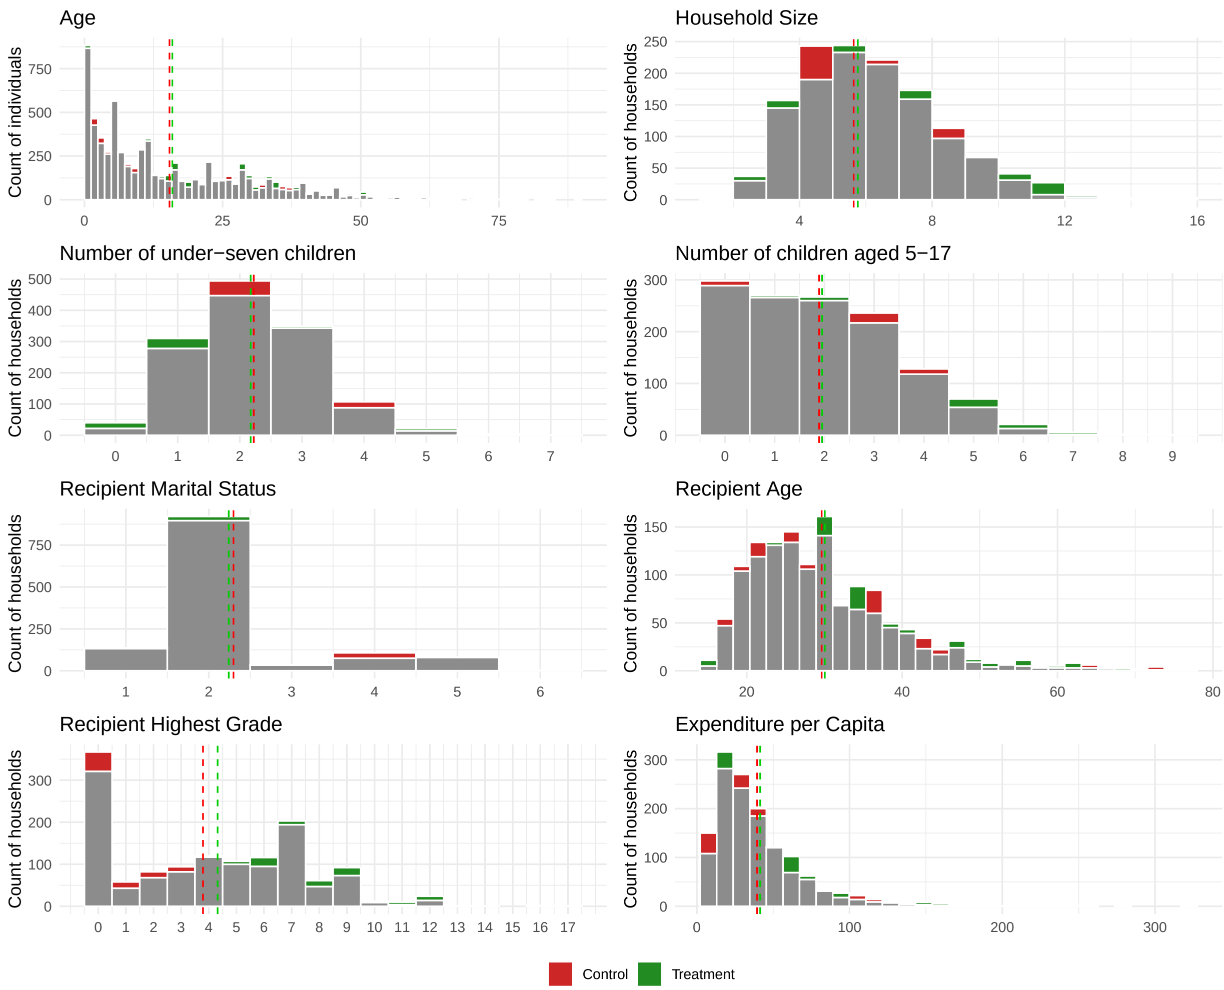

Note: Red and green represent control-group- and treatment-group-specific excess over the overlap portions for each category in gray, respectively. Vertical dashed lines represent group means. Education codes are defined as follows: grades 1–12 correspond to codes 01–12; GCE A-level is coded as 13; college is coded as 14; undergraduate university is coded as 15; post-graduate certificate or diploma is coded as 16; master's degree is coded as 17; and doctoral level or above is coded as 18. Marital status codes are defined as follows: never married is coded as 1, married is coded as 2, separated is coded as 3, divorced is coded as 4, widowed is coded as 5, and co-habiting is coded as 6.

Figure A2: **Non-Parametric Relationship Between Outcome and Age**

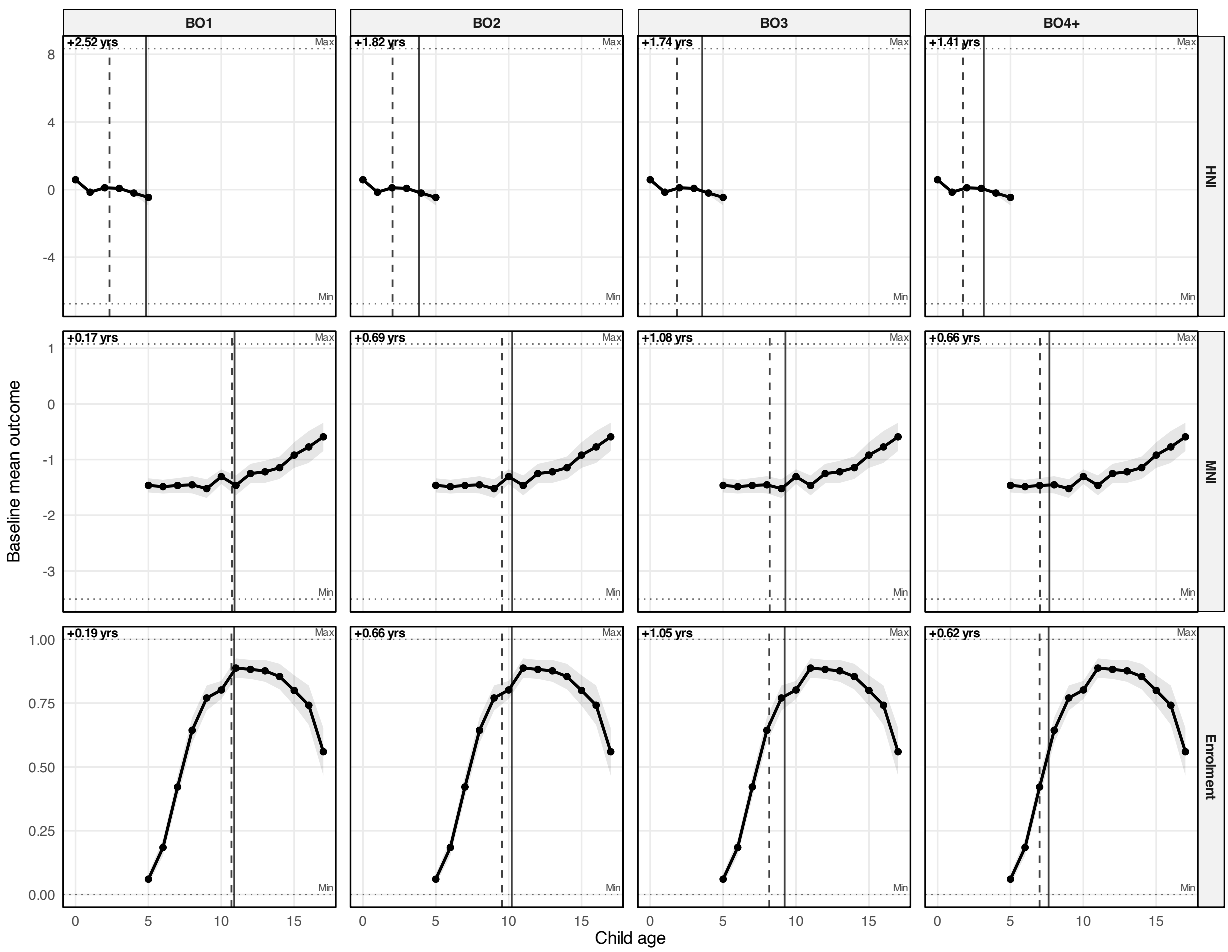

Note: The figure illustrates the relationship between child outcome and age in solid lines. Each panel then separately demonstrates the change in the average age of each level of BO, with the vertical dashed line corresponding to baseline and vertical solid line corresponding to the follow-up waves.

Figure A3: **Comparison of Regression Results Adjusting for the Association Between Outcome and Age**

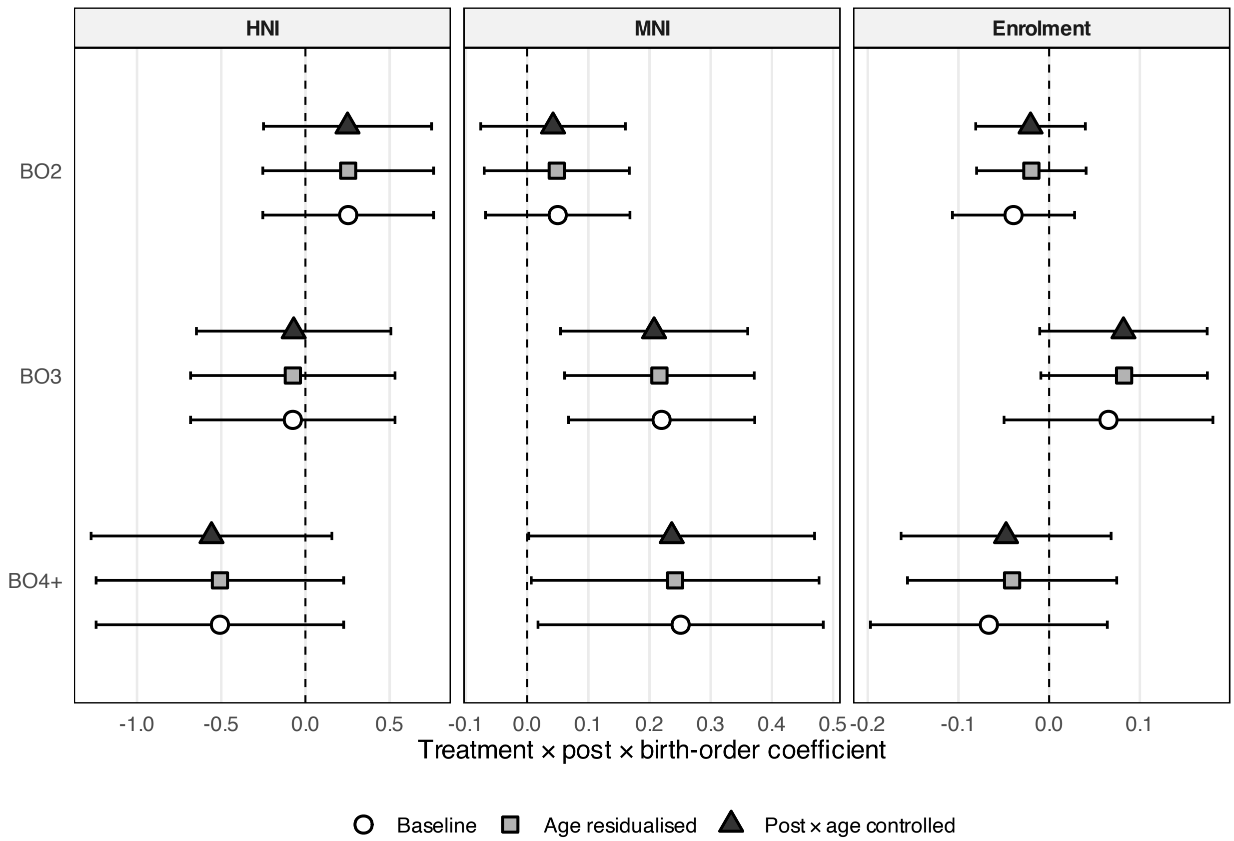

Note:

A.5 Tables

Table A1: **Reason for Attrition**

		24m		30m		36m		48m	
Age group	Reason	Control	Treat	Control	Treat	Control	Treat	Control	Treat
Age 0–7	Left to find a job	0	0	0	0	0	0	0	1
	Separated	0	0	0	0	1	0	3	1
	Married away	1	0	2	0	3	0	4	2
	Deceased	54	47	86	76	104	85	125	96
	Divorced	0	0	1	0	1	0	1	1
	Living with other relatives	79	85	129	137	198	188	273	278
	Established own home	0	0	3	0	4	0	6	0
	Left with mother/father	16	28	30	48	50	59	76	76
	Attend school	2	4	5	11	7	11	8	14
	Hired worker left	0	0	0	0	0	0	0	0
	To look after parent(s)	0	0	0	0	0	0	0	0
	Disagreed with head	0	0	0	0	0	0	0	0
	Imprisoned	0	0	0	1	0	1	0	1
	Reason not given	10	3	10	3	10	3	12	3
	Never a member	0	0	16	20	39	37	50	53
Age 5–17	Left to find a job	5	1	5	3	8	3	9	9
	Separated	0	0	0	0	1	0	1	1
	Married away	17	31	42	58	55	66	78	103
	Deceased	12	11	20	21	23	25	29	49
	Divorced	0	1	0	3	0	3	2	4
	Living with other relatives	136	145	235	228	320	306	449	446
	Established own home	0	3	12	8	14	14	29	30
	Left with mother/father	26	28	41	53	61	61	87	67
	Attend school	13	8	19	23	27	25	34	32
	Hired worker left	0	0	0	0	0	0	0	0
	To look after parent(s)	0	0	1	2	2	2	4	2
	Disagreed with head	0	0	0	0	0	0	0	0
	Imprisoned	0	0	0	2	0	2	0	2
	Reason not given	12	6	12	6	13	6	15	7
	Never a member	2	0	11	16	24	39	39	54

Notes: Newborns after baseline are excluded. Age groups are defined using age observed at baseline.

Table A2: **ITT Effects of the CGP: Extended Outcomes**

	Child Anthropometry			Material Needs			Enrolment
	HAZ	WAZ	WHZ	Blanket	Shoes	Clothes x2	Enrol
Panel A: Wave-Specific Effect							
Treated × 24m	0.013	0.101*	0.082	0.193***	0.322***	0.083	0.010
	(0.068)	(0.057)	(0.060)	(0.052)	(0.054)	(0.052)	(0.021)
Treated × 30m	0.012	0.046	0.005	0.143***	0.287***	0.033	0.057***
	(0.076)	(0.059)	(0.065)	(0.052)	(0.037)	(0.052)	(0.019)
Treated × 36m	-0.048	0.044	0.023	0.126**	0.259***	0.025	0.025
	(0.083)	(0.065)	(0.074)	(0.050)	(0.041)	(0.052)	(0.020)
Treated × 48m	-0.019	0.100	0.140	0.138***	0.280***	0.072	0.045**
	(0.086)	(0.071)	(0.097)	(0.049)	(0.044)	(0.056)	(0.020)
Panel B: Pooled Effect							
Treated × Post	-0.012	0.070	0.059	0.148***	0.284***	0.052	0.035**
	(0.070)	(0.053)	(0.061)	(0.048)	(0.035)	(0.052)	(0.017)
Age range	0–7	0–7	0–7	5–17	5–17	5–17	5–17
Mean depvar at baseline	-1.316	-0.828	-0.177	0.557	0.141	0.636	0.586
Observations	22201	22922	21370	27006	27003	26997	27544
Within R^2	0.037	0.031	0.009	0.211	0.200	0.181	0.201
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors clustered at the CWAC level in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A3: **ATT Effects of the CGP**

	HNI	MNI	Enrol
$\hat{\text{Eligible}}$	0.057	0.809***	0.043**
	(0.077)	(0.189)	(0.021)
Age range	0–7	5–17	5–17
Mean depvar at baseline	0.074	-1.293	0.586
Observations	21370	26983	27544
R^2	0.035	0.316	0.195
District FE	Yes	Yes	Yes
First-stage F-statistic	9725.513	11247.628	11495.036

Note: Standard errors clustered at the CWAC level in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A4: **ATT Effects of the CGP: Extended Outcomes**

	Child Anthropometry			Material Needs			Enrolment
	HAZ	WAZ	WHZ	Blanket	Shoes	Clothes x2	Enrol
$\widehat{\text{Eligible}}$	-0.015	0.082	0.070	0.179***	0.346***	0.064	0.043**
	(0.083)	(0.063)	(0.072)	(0.059)	(0.044)	(0.063)	(0.021)
Age range	0–7	0–7	0–7	5–17	5–17	5–17	5–17
Mean depvar at baseline	-1.316	-0.828	-0.177	0.557	0.141	0.636	0.586
Observations	22201	22922	21370	27006	27003	26997	27544
Within R^2	0.037	0.033	0.009	0.211	0.200	0.181	0.195
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-stage F-statistic	9663.901	10332.512	9725.513	11261.327	11268.010	11261.216	11495.036

Standard errors clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A5: **Two-Step Location-Scale Estimates of the CGP Impact on Within-Household Residual Dispersion**

	HNI	MNI	Enrol
Panel A: Without 1st stage BO term			
Wave-Specific Effect			
Treated × 24m	-0.000	-0.327	-0.080
	(0.133)	(0.199)	(0.113)
Treated × 30m	-0.023	-0.316	0.003
	(0.130)	(0.214)	(0.115)
Treated × 36m	-0.039	-0.116	-0.066
	(0.135)	(0.204)	(0.110)
Treated × 48m	0.020	-0.435**	-0.106
	(0.130)	(0.206)	(0.102)
Pooled Effect			
Treated × Post	-0.010	-0.301*	-0.063
	(0.103)	(0.182)	(0.097)
Panel B: With 1st stage BO term			
Wave-Specific Effect			
Treated × 24m	0.057	-1.698**	-0.001
	(0.145)	(0.722)	(0.135)
Treated × 30m	0.036	-1.700**	0.003
	(0.138)	(0.745)	(0.131)
Treated × 36m	0.008	-1.069	-0.019
	(0.142)	(0.736)	(0.116)
Treated × 48m	0.032	-1.988***	-0.061
	(0.152)	(0.707)	(0.114)
Pooled Effect			
Treated × Post	0.019	-1.856***	-0.029
	(0.120)	(0.557)	(0.108)
Age range	0–7	5–17	5–17
Cluster bootstrap	Yes	Yes	Yes
Bootstrap replications	500	500	500

Note: The dependent variable in the second stage is the log variance of first-stage residuals within each household-time cell, $\hat{\sigma}^2_{hjdt}$. Second-stage regressions control for baseline household expenditure per capita, recipient education, recipient age, recipient marital status, log household size, household age composition, and household gender composition. Standard errors are cluster bootstrap standard errors at the CWAC level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A6: **Birth-Order Effects (Male Only Sample)**

	HNI			MNI			Enrolled in school		
	BO2	BO3	BO4+	BO2	BO3	BO4+	BO2	BO3	BO4+
Panel A: Treated									
Treated × Post × BO	0.199	-0.546	-0.707	0.065	0.270**	0.282	-0.104	-0.013	-0.169*
	(0.456)	(0.607)	(0.775)	(0.092)	(0.125)	(0.182)	(0.066)	(0.091)	(0.088)
Panel B: Control									
Control × Post × BO	-0.201	0.050	0.480	0.021	-0.035	0.036	0.055	0.086	0.374***
	(0.255)	(0.329)	(0.399)	(0.075)	(0.093)	(0.121)	(0.049)	(0.065)	(0.060)
Baseline mean	-0.148	0.005	0.062	-1.350	-1.390	-1.430	0.626	0.469	0.255
Baseline gap vs. BO1	0.011	0.163	0.220	-0.003	-0.047	-0.085	-0.058	-0.214	-0.429
Age range		0–7			5–17			5–17	
Observations		4797			8106			8277	
Within R^2		0.019			0.067			0.274	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: BO1 is the reference birth-order category. Standard errors clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A7: **Birth-Order Effects (Female Only Sample)**

	HNI			MNI			Enrolled in school		
	BO2	BO3	BO4+	BO2	BO3	BO4+	BO2	BO3	BO4+
Panel A: Treated									
Treated × Post × BO	0.248	-0.269	-0.388	0.148	0.345**	0.169	-0.105	0.046	0.049
	(0.461)	(0.596)	(0.677)	(0.100)	(0.137)	(0.173)	(0.078)	(0.108)	(0.132)
Panel B: Control									
Control × Post × BO	-0.318	-0.204	-0.077	0.087	-0.002	0.120	0.009	-0.002	0.077
	(0.322)	(0.391)	(0.408)	(0.068)	(0.109)	(0.142)	(0.058)	(0.088)	(0.111)
Baseline mean	0.086	0.264	0.173	-1.390	-1.350	-1.410	0.615	0.551	0.352
Baseline gap vs. BO1	0.062	0.240	0.149	-0.187	-0.149	-0.208	-0.034	-0.097	-0.296
Age range		0–7			5–17			5–17	
Observations		4960			7993			8222	
Within R^2		0.060			0.083			0.203	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: BO1 is the reference birth-order category. Standard errors clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A8: **Heterogeneity by the Gender of Older Siblings**

	HNI	MNI	Enrol
Panel A: Treated			
Treated × Post × Girl	-0.039	-0.066	0.123**
	(0.186)	(0.085)	(0.060)
Treated × Post × No elder brother	-0.381	0.065	0.087
	(0.338)	(0.093)	(0.081)
Treated × Post × Girl × No elder brother	0.062	0.056	-0.120
	(0.403)	(0.119)	(0.088)
Panel B: Control			
Control × Post × Girl	0.083	-0.011	-0.119***
	(0.145)	(0.050)	(0.045)
Control × Post × No elder brother	0.267	0.041	-0.094
	(0.215)	(0.050)	(0.059)
Control × Post × Girl × No elder brother	-0.449	-0.098	0.114*
	(0.292)	(0.095)	(0.058)
Age range	0–7	5–17	5–17
Mean depvar at baseline	0.074	-1.293	0.586
Observations	8478	20125	20515
Within R^2	0.032	0.065	0.241
District FE	Yes	Yes	Yes
HH × Time FE	Yes	Yes	Yes

Note: Standard errors are clustered at the CWAC level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A9: **Birth-Order Effects (Most Constrained Quintile of Households)**

	HNI			MNI			Enrolled in school		
	BO2	BO3	BO4+	BO2	BO3	BO4+	BO2	BO3	BO4+
Panel A: Treated									
Treated × Post × BO	1.206**	0.415	-0.427	-0.079	0.025	0.130	-0.123	-0.026	-0.211*
	(0.456)	(0.473)	(0.722)	(0.111)	(0.134)	(0.193)	(0.075)	(0.106)	(0.108)
Panel B: Control									
Control × Post × BO	-0.434	0.082	0.259	0.062	0.018	-0.071	0.057	0.082	0.249***
	(0.276)	(0.295)	(0.381)	(0.073)	(0.103)	(0.127)	(0.053)	(0.067)	(0.081)
Baseline mean	-0.225	-0.015	0.180	-1.813	-1.808	-1.832	0.623	0.492	0.296
Baseline gap vs. BO1	-0.055	0.155	0.350	-0.051	-0.046	-0.070	-0.063	-0.193	-0.390
Age range		0–7			5–17			5–17	
Observations		3690			5569			5663	
R^2		0.573			0.930			0.511	
Within R^2		0.044			0.069			0.265	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: BO1 is the reference birth-order category. Standard errors clustered at the CWAC level in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A10: **Birth-Order Effects (Large Households)**

	HNI			MNI			Enrolled in school		
	BO2	BO3	BO4+	BO2	BO3	BO4+	BO2	BO3	BO4+
Panel A: Treated									
Treated × Post × BO	0.409	0.106	-0.353	0.051	0.228***	0.254**	-0.037	0.064	-0.060
	(0.568)	(0.616)	(0.646)	(0.066)	(0.081)	(0.119)	(0.042)	(0.064)	(0.072)
Panel B: Control									
Control × Post × BO	-0.647	-0.216	0.161	0.078*	0.016	0.039	0.064**	0.090*	0.272***
	(0.445)	(0.438)	(0.455)	(0.044)	(0.061)	(0.084)	(0.031)	(0.049)	(0.057)
Baseline mean	0.309	0.185	0.112	-1.359	-1.376	-1.422	0.677	0.519	0.301
Baseline gap vs. BO1	0.291	0.167	0.094	-0.145	-0.162	-0.208	-0.079	-0.237	-0.456
Age range		0–7			5–17			5–17	
Observations		8490			16488			16742	
Within R^2		0.035			0.065			0.239	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: BO1 is the reference birth-order category. Standard errors are clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A11: **Baseline Outcome-Rank Effects**

	HNI			MNI			Enrolled in school		
	Rank 2	Rank 3	Rank 4+	Rank 2	Rank 3	Rank 4+	Rank 2	Rank 3	Rank 4+
Panel A: Treated									
Treated × Post × Rank	0.108	0.351	-0.877	-0.019	0.090	0.151	0.038	-0.035	0.055
	(0.094)	(0.250)	(0.645)	(0.114)	(0.198)	(0.276)	(0.042)	(0.036)	(0.048)
Panel B: Control									
Control × Post × Rank	0.878***	1.132***	2.187***	1.337***	1.600***	1.935***	0.713***	0.780***	0.738***
	(0.068)	(0.202)	(0.533)	(0.068)	(0.120)	(0.155)	(0.030)	(0.027)	(0.033)
Baseline mean	-0.559	-0.654	-1.255	-1.559	-1.779	-2.001	0.000	0.000	0.000
Baseline gap vs. rank 1	-1.293	-1.388	-1.989	-0.323	-0.543	-0.766	-0.823	-0.823	-0.823
Age range		0–7			5–17			5–17	
Observations		6313			13899			17461	
Within R^2		0.310			0.388			0.334	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: Rank 1 is the reference baseline outcome-rank category. Standard errors clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A12: **Age-Gap Effects**

	HNI	MNI	Enrol
Panel A: Treated			
Treated × Post × Age gap from BO1	-0.058	0.021*	0.000
	(0.044)	(0.012)	(0.007)
Panel B: Control			
Control × Post × Age gap from BO1	0.048*	0.010	0.012**
	(0.028)	(0.009)	(0.005)
Mean depvar at baseline	0.060	-1.336	0.577
Mean age gap from BO1 at baseline	6.858	3.276	3.277
Age range	0–7	5–17	5–17
Observations	14702	19550	19951
Within R^2	0.023	0.055	0.227
District FE	Yes	Yes	Yes
HH × Time FE	Yes	Yes	Yes

Note: Standard errors clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A13: **Birth-Order Effects (Cohort Sample)**

	HNI			MNI			Enrolled in school		
	BO2	BO3	BO4+	BO2	BO3	BO4+	BO2	BO3	BO4+
Panel A: Treated									
Treated × Post × BO	0.257	-0.274	-0.708*	0.067	0.227***	0.264**	-0.020	0.074	-0.047
	(0.255)	(0.310)	(0.382)	(0.060)	(0.077)	(0.122)	(0.034)	(0.053)	(0.075)
Panel B: Control									
Control × Post × BO	-0.072	0.225	0.638***	0.066	-0.003	-0.009	0.205***	0.355***	0.728***
	(0.147)	(0.166)	(0.214)	(0.043)	(0.059)	(0.090)	(0.026)	(0.041)	(0.060)
Baseline mean	-0.030	0.142	0.118	-1.365	-1.368	-1.418	0.620	0.511	0.303
Baseline gap vs. BO1	0.035	0.207	0.183	-0.097	-0.100	-0.150	-0.045	-0.154	-0.362
Age range		0–7			5–17			5–17	
Observations		9665			14206			14337	
Within R^2		0.030			0.064			0.115	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: BO1 is the reference birth-order category. Standard errors clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A14: **Birth-Order Effects (Perceived the CGP as Conditional Households)**

	HNI			MNI			Enrolled in school		
	BO2	BO3	BO4+	BO2	BO3	BO4+	BO2	BO3	BO4+
Panel A: Treated									
Treated × Post × BO	0.202	-0.149	-0.537	0.057	0.204***	0.253**	-0.020	0.094	-0.040
	(0.268)	(0.316)	(0.378)	(0.059)	(0.077)	(0.117)	(0.036)	(0.060)	(0.068)
Panel B: Control									
Control × Post × BO	-0.066	0.126	0.528***	0.088**	0.030	0.054	0.006	0.011	0.214***
	(0.139)	(0.154)	(0.189)	(0.041)	(0.059)	(0.085)	(0.027)	(0.046)	(0.053)
Baseline mean	-0.001	0.019	-0.004	-0.015	-0.076	-0.162	0.037	-0.101	-0.327
Baseline gap vs. BO1	0.011	0.031	0.008	-0.118	-0.179	-0.264	-0.095	-0.233	-0.458
Age range		0–7			5–17			5–17	
Mean depvar at baseline		0.074			-1.293			0.586	
Observations		12894			17023			17350	
Within R^2		0.035			0.064			0.240	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: BO1 is the reference birth-order category. Standard errors clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A15: **Birth-Order Effects (Always Treated Households)**

	HNI			MNI			Enrolled in school		
	BO2	BO3	BO4+	BO2	BO3	BO4+	BO2	BO3	BO4+
Panel A: Treated									
Treated × Post × BO	0.280	-0.047	-0.493	0.071	0.239***	0.302**	-0.047	0.074	-0.062
	(0.255)	(0.304)	(0.369)	(0.063)	(0.081)	(0.119)	(0.036)	(0.061)	(0.067)
Panel B: Control									
Control × Post × BO	-0.064	0.129	0.534***	0.089**	0.031	0.053	0.006	0.012	0.214***
	(0.139)	(0.154)	(0.189)	(0.041)	(0.059)	(0.085)	(0.027)	(0.046)	(0.053)
Baseline mean	-0.001	0.019	-0.004	-0.018	-0.074	-0.174	0.041	-0.104	-0.334
Baseline gap vs. BO1	0.012	0.031	0.008	-0.124	-0.181	-0.280	-0.091	-0.236	-0.466
Age range		0–7			5–17			5–17	
Mean depvar at baseline		0.075			-1.312			0.584	
Observations		15583			20320			20726	
Within R^2		0.032			0.064			0.247	
District FE		Yes			Yes			Yes	
HH × Time FE		Yes			Yes			Yes	

Note: BO1 is the reference birth-order category. Standard errors clustered at the CWAC level in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A16: **Marginal Effects of the CGP on Child Preference and Inequality Aversion (Extended Results)**

	Child preference α		Inequality aversion ρ	
	(1)	(2)	(1)	(2)
Intercept (36m, control)	-0.042^{**}	-0.056^{**}	-14.114^{***}	-15.295^{***}
	(0.020)	(0.024)	(2.875)	(3.117)
Treated	0.024	0.041^{*}	3.652^{**}	3.290^{*}
	(0.016)	(0.022)	(1.652)	(1.851)
48m	0.006	0.005	3.349^{**}	2.995^{*}
	(0.016)	(0.019)	(1.537)	(1.753)
Treated × 48m	−0.022	−0.019	-8.002^{***}	-6.322^{***}
	(0.021)	(0.024)	(2.085)	(2.214)
Age range	-0.036^{***}	-0.049^{***}		
	(0.009)	(0.016)		
Age mean	−0.013	−0.012		
	(0.008)	(0.010)		
Older child female	0.028^{**}	0.015		
	(0.012)	(0.015)		
Younger child female	-0.021^{*}	-0.036^{*}		
	(0.012)	(0.019)		
Female decision-maker	−0.024		7.860^{***}	
	(0.017)		(2.254)	
Expenditure per capita	−0.006	−0.008	0.499	0.621
	(0.003)	(0.005)	(0.579)	(0.537)
Recipient age	−0.004	−0.000	1.080	−0.772
	(0.006)	(0.009)	(0.913)	(0.927)
Recipient highest grade	-0.014^{**}	-0.025^{**}	-1.388^{*}	-3.276^{***}
	(0.007)	(0.011)	(0.792)	(0.953)
Recipient married	0.007	−0.004	-4.343^{**}	−0.596
	(0.019)	(0.017)	(2.109)	(1.890)
Household size	−0.009	-0.018^{*}	−0.603	0.427
	(0.007)	(0.010)	(0.852)	(0.804)
Household average age	0.010	0.007	−1.057	0.191
	(0.006)	(0.006)	(0.796)	(0.792)
Fraction household female	0.013	0.017^{*}	1.254^{*}	0.981
	(0.008)	(0.010)	(0.732)	(0.785)
Observations	2777	2777	2777	2777
Residual standard error	0.216	0.217	0.216	0.217
Convergence iterations	43	49	43	49
Mean	0.467	0.450	-6.158	-6.456
Median	0.466	0.450	-4.518	-5.462
Minimum	0.294	0.219	-42.908	-29.567
Maximum	0.635	0.668	-0.296	-0.592

Notes: Standard errors two-way clustered at the household and CWAC levels in parentheses. Column (1) reports estimated marginal effects with $Female\ decision-maker$ included – an indicator that is coded as 1 if the primary female care-taker is the household's sole decision-maker regarding child health and nutrition, or 0 if her husband is involved in decision-making. Column (2) reports estimated marginal effects without $Female\ decision-maker$. Dyadic-specific variables, which include age and gender, are not entered into the inequality aversion equation so as to ensure the parameter varies at the household level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.