Conditioning Out the Poor? Consumption Inequality and the Design of Cash Transfer Programs*

Santosh Anagol† Thomas Fujiwara‡ Martin Navarrete§

October 2025

Abstract

Conditionality can exclude poor households from receiving transfers. Reanalyzing five randomized evaluations of conditional cash transfers (CCTs), we find that 9–37% of eligible recipients fail to meet conditions, and they typically have lower baseline consumption. We assess the welfare implications of budget-neutral shifts from CCTs to unconditional cash transfers. Conditionality exacerbating inequality among eligible recipients can be quantitatively important for welfare impacts. We quantify how large conditionality-induced human capital gains must be for the welfare benefits of CCTs to outweigh the inequality costs generated by conditionality.

*We thank Dan Bjorkegren, Rema Hanna, and Katrina Kosec for help with data access and Pedro Cabra Acela and Jeremias A. Manzano Quiroga for excellent research assistance. We also thank Shing-Yi Wang, Ben Lockwood, and Ben Sprung-Keyser for their insightful comments.

†The Wharton School, University of Pennsylvania. Email: anagol@wharton.upenn.edu.

‡Department of Economics and SPIA, Princeton University. Email: fujiwara@princeton.edu.

§The Wharton School, University of Pennsylvania. Email: martinnc@wharton.upenn.edu.

1 Introduction

Conditional cash transfers (CCTs) are "*perhaps the single largest innovation in social protection programs in the developing world in the past 30 years*" (Banerjee et al., 2024). Proponents argue that by conditioning cash benefits on actions such as enrolling children in school, the policy provides income support while promoting human capital investments. Moreover, conditionality can lead to welfare gains when households underestimate human capital returns, discount future generations' utility too much, or fail to internalize externalities.

However, if households not meeting the conditions are poorer than those that do, the program effectively targets richer households and increases inequality among eligibles. Under standard social welfare functions, this reduces welfare gains. To our knowledge, no prior work quantifies the welfare consequences of conditionality that exacerbates consumption inequality; this is the focus of our paper.

Theoretically, it is unclear whether poorer or richer households are more likely to meet conditions. Poorer households derive higher marginal utility from transfers, but may face higher opportunity costs, rely more on child labor, or underestimate returns to schooling more severely.

We analyze data from five randomized evaluations of CCTs in Malawi (SIHR), Mexico (Progresa), Morocco (Tayssir), Nicaragua (RPS), and Tanzania (TASAF).[1] We present six results for each context. First, we report the fraction of conditions-met and conditions-unmet children by age. Conditions-unmet children are eligible children who do not meet their program's most basic educational condition: school enrollment. Among eligible children assigned to a CCT treatment, the shares not meeting conditions are 36.8% in Malawi, 19.7% in Mexico, 20.8% in Morocco, 15.6% in Tanzania, and 9.0% in Nicaragua. The likelihood of not meeting conditions varies by age, rising during the teenage years in most contexts.

Second, we compare the baseline (before transfers start) household consumption distributions for eligible children who meet versus do not meet the conditions. Although

[1] Only these programs met our requirements for program design and data availability (Section 2).

straightforward, this comparison has not been systematically reviewed and is not usually reported in CCT evaluations. In Malawi, conditions-unmet children have, on average, 25% lower consumption. In Mexico, Tanzania, and Morocco, conditions-unmet children are also poorer, but differences in averages are more modest (between 4%-6%). In Nicaragua, the average consumption of conditions-unmet children is 8% higher, but the number of conditions-unmet children is small.

Third, we calculate counterfactual welfare gains from switching from the implemented CCT to a hypothetical unconditional cash transfer (UCT) that removes conditionality in a budget-neutral manner, providing payments to all eligible households.[2] Following Alatas et al. (2019) and Hanna and Olken (2018), we adopt a social welfare function with a constant relative risk aversion (CRRA) utility. Our baseline results isolate the role of consumption inequality, abstracting from human capital gains. While we address these later, presenting results without them clarifies the paper's main mechanism. Moreover, they may be empirically relevant since evidence on whether CCTs boost human capital gains relative to UCTs is unclear.[3]

Welfare consequences are substantial in some contexts but more muted in others. In Malawi, we estimate the CCT delivers only 67.5% of the welfare gains that an equally budgeted UCT program would deliver. In contrast, in the four other contexts, we find CCTs deliver between 93% to 100% of the UCT benefit. We discuss the drivers of variation across contexts (transfer sizes, more conditions-unmet children, and whether they have relatively lower consumption, all of which are sizable in Malawi).

Fourth, we highlight a key mechanism: even when conditions-met and -unmet children have similar consumption distributions, concave utility implies that UCTs, by reaching all eligibles, can yield higher welfare than CCTs, which do not reach all eligibles. We quantify this "unequal-payments effect" by providing counterfactual welfare ratios as-

[2] Section 3 discusses why we focus on this particular counterfactual.

[3] A review of 35 studies (Baird et al., 2013) finds no significant difference between the effects of CCTs and UCTs on enrollment (both UCTs and CCTs have positive effects relative to pure controls). To our knowledge, three randomized trials include a CCT and UCT arm. Two find that the CCT effect on school enrollment is 40% larger than the UCT's effect (Baird et al., 2011; Akresh et al., 2020), and one finds no difference in their effects (Robertson et al., 2013). Benhassine et al. (2015) compares a CCT to a labeled cash transfer (with no conditionality, but labeled as supportting education) and finds that conditionality "*made almost no difference*."

suming no consumption differences by condition adherence.

Fifth, we incorporate human capital into welfare comparisons using two approaches. The first assumes households make individually and socially optimal choices, applying a revealed-preference argument to bound the value of marginal investments induced by conditionality. Given modest shares of marginals, adding human capital has a limited impact on welfare comparisons even under assumptions favoring CCTs.

We also adopt a return-on-investment (ROI) approach, following Bergstrom and Dodds (2021), to compute the present value of human capital gains induced by conditionality. Even assuming low returns to schooling, this approach suggests CCTs yield equal or higher welfare than UCTs, except in Malawi. It also accommodates a social planner who values education more than households (e.g., if households underestimate returns or over-discount future generations' utility). We also discuss how to incorporate externalities under both approaches.[4]

While taking a position on whether the revealed-preference or ROI approach is preferable is beyond the scope of this paper, the results underscore that evaluating whether households make socially optimal decisions is central to welfare considerations in cash transfer design.

Sixth, we focus on subsamples of older children with lower condition adherence, approximating a policy shift away from younger children who often meet conditions without conditionality. Our results show that the inequality-increasing effect of conditionality can be larger in such contexts, and whether these welfare losses are offset by human capital gains depends on assumptions and context.

We contribute to the literature on cash transfer programs—see Banerjee et al. (2024) for a review. To our knowledge, the only other study comparing CCTs and UCTs within an optimal redistribution framework is Bergstrom and Dodds (2021). Using Progresa data, it focuses on the "targeting *benefit*" of CCTs and examines an internal solution involving

[4] Banerjee et al. (2024) discusses the welfare rationale for conditionality and how it assumes a social planner values human capital more than households do. Another possibility is CCTs providing "tagging" benefits. Even if CCTs have no effects on human capital, a policymaker may prefer to make transfers to conditions-meeting households (e.g., they may direct more resources to children). While household credit constraints provide a rationale for implementing CCTs rather than no transfers, it does not provide a rationale for conditionality itself, since unconditional transfers also relax such constraints. A separate issue, beyond the scope of this paper, is that conditionality is often motivated by political factors.

a mix of CCTs and UCTs. Our paper addresses a different policy choice between only a CCT and only a UCT, and documents the potential "targeting *cost*" of screening out households not meeting conditions.

While the idea that conditions-met and conditions-unmet households differ is not new, it receives relatively little attention in the literature. CCT evaluations rarely, if ever, report consumption differences by condition adherence. Some studies examine how removing conditionality affects conditions-unmet households, but usually focus on outcomes other than consumption (e.g., Baird et al., 2011, studies early marriage). We hope that, by systematically providing consumption comparisons, offering an approach to welfare assessments, and applying it to multiple contexts under different assumptions, this paper highlights the importance of consumption inequality and informs the design of transfer programs.

In developed countries, a public finance literature examines the distributional and welfare effects of behavior-targeted policies, such as early retirement incentives (Kolsrud et al., 2024) and sin taxes (Allcott et al., 2019). Another strand studies how administrative burdens limit enrollment in transfer programs (Herd and Moynihan, 2025). We focus on a separate obstacle (conditionality) amongst those enrolled.

2 Contexts and Data

We focus on randomized evaluations of CCTs. Although our analysis does not directly use randomization, we focus on experiments because they provide both baseline (pre-treatment) consumption and endline (post-treatment) enrollment for the same household, which is not available in cross-sectional surveys.

To select the experiments, we began with the 18 randomized trials considered by Banerjee et al. (2017) study of transfers' effects on labor supply. We added Robertson et al. (2013), one of the three experiments we are aware of that include both CCT and UCT treatment arms. We analyze the experiments meeting three criteria: they evaluate a CCT (e.g., not include only UCT and a control), provide baseline household-level data on eligibility and consumption, and have post-treatment data on school enrollment (the condition). Table A.1 lists the 19 experiments considered. Seven only evaluate UCTs, three lack consumption data, one lacks eligibility information, and one lacks enrollment data. Data was not available for two other studies. Appendix A provides further information and discusses how included and non-included experiments compare.

The five included experiments are Malawi's SIHR (Baird et al., 2011), Mexico's Progresa (Parker and Todd, 2017), Morocco's Tayssir (Benhassine et al., 2015), Nicaragua's RPS (Maluccio and Flores, 2005), and Tanzania's TASAF (Evans et al., 2014). We refer to these as the "source papers" for each context.

Program characteristics. All programs are operated by governments and include girls and boys. The exception is Malawi, which is not government-run and targets teenage girls who were already enrolled at baseline. A review (Baird et al., 2013) classifies the Malawi, Morocco, and Nicaragua programs as having "*explicit schooling conditions monitored and enforced*" and the Mexican Progresa as having "*explicit conditions, (imperfectly) monitored, with minimal enforcement*." It does not consider the Tanzanian program, but Evans et al. (2014) reports that conditions were monitored and enforced.

Column (2) of Table 1 reports average transfer sizes, expressed as shares of per capita household consumption. The Malawi and Mexico programs have large transfer sizes (25.2% and 18.4%, respectively). Nicaragua has a moderate transfer size at 11.2%, while Morocco and Tanzania have substantially smaller transfer sizes, at 3.0% and 6.2% respectively. In some contexts, transfer sizes depend on child age and gender. Appendix B provides a detailed discussion of each context and its data.

Data harmonization. Our analysis samples consist only of children eligible for the CCT. We exclude those in non-CCT and control arms and those ineligible (e.g., age below or above eligibility). We measure consumption *before* payments start (i.e., in a baseline survey before treatment). We adjust consumption so it measures total household consumption per capita.[5]

Our measure of meeting program conditions is a dummy for whether the child is

[5] The adjustment to per capita terms uses as household size the number of adults plus 3/5 the number of children under 12 years old (Deaton, 1997).

enrolled in school.[6] We always use the variable that the source papers use to estimate effects on human capital in the period they label as the "endline." We exclude observations missing data on consumption and/or school enrollment and and exclude households with consumption below the 5th and above the 95th percentile of their context-specific distributions (Appendix B).

Measuring conditions adherence. While school enrollment may be part of a larger set of conditions, including attendance and health components such as vaccinations, we focus on enrollment for four reasons. First, school enrollment is the condition most likely to be enforced in practice (Baird et al., 2013). Second, it provides a comparable definition of conditions-adherence across contexts. Third, data on meeting other conditions is often unavailable. Fourth, in all our contexts, microdata detailing which households were excluded from payments for non-adherence is unavailable. Even if it were available, enforcement of conditions in a pilot evaluation may differ from when the program is operating at scale, and enrollment can potentially be a better predictor of not receiving transfers at this later stage.

Only data reported by households in a survey are available to measure enrollment, raising potential measurement error concerns. We highlight three points here. First, the source papers (and several others, in the Mexican case) use the same enrollment variables as their main outcomes to evaluate CCT effects on human capital. Second, analysis in the source papers for Malawi and Morocco suggests little measurement error on average (Appendix B). Third, if measurement error in the enrollment variable is independent of baseline consumption, our analysis would be unaffected.

Condition adherence. Figure 1 shows the share of children *not* meeting conditions by age; Table 1 reports context-specific averages. In all contexts, shares are low for children under 11 and rise with age above that. Tanzania and Nicaragua are exceptions, with higher shares at younger ages, reflecting delayed school entry. In Nicaragua, shares are low for ages 11–13 because eligibility was limited to children who had not completed

[6] Except in Morocco, where it is enrollment and attending at least once in May 2010.

primary school (Appendix B).

Variation in shares meeting conditions across contexts reflects several factors. The first is the age of eligibles: Malawi has the highest share as it targets 14–20-year-olds, while Nicaragua has the lowest, targeting 7–13-year-olds. The second is the eligibility criteria: in Malawi, only children enrolled at baseline were eligible. Thus, conditional on age, it has fewer conditions-unmet than Morocco or Mexico. Low shares for 12–13-year-olds in Nicaragua stem from eligibility being limited to those who had not completed primary school. The third is baseline national differences, e.g., Mexico's countrywide secondary school enrollment was higher than Morocco's at the time of the experiments.

Subsamples. We also report results for age-based subsamples with higher shares of children not meeting conditions. As Figure 1 shows—and as noted in prior work (e.g., De Janvry and Sadoulet, 2006)—CCTs often target young children who already have high enrollment, raising the possibility that targeting older children could yield larger enrollment effects. We examine such subgroups to shed light on the implications of this potential reform.

We define subsamples as age groups where more than 10% of children have unmet conditions, excluding the earliest school-eligible ages where late entry may explain the pattern. This criterion yields three subsamples: Mexico (ages 13+), Morocco (ages 11+), and Tanzania (ages 13+), and no subsamples for Malawi and Nicaragua.

3 Results

Consumption differences by condition adherence. Figure 2 compares the consumption distributions for conditions-met and conditions-unmet children. Specifically, it provides histograms of the logarithm of total household baseline consumption per capita (before payments start).

In Malawi, the conditions-unmet have a distribution clearly shifted to the left of the conditions-met distribution (and mean consumption is 25% lower). In Morocco, Mexico, and Tanzania, there are no clearly visible differences between the distributions, but mean consumption for conditions-unmet children is 4%-6% lower. In Nicaragua, it is 8% *higher* on average, with the caveat that the conditions-unmet sample is small (64 children). Figure 3 shows these distributions for the subsamples. For Mexico and Tanzania, they are not visibly different than for the full sample. For Morocco, the conditions-unmet distribution is modestly shifted to the left (difference in averages is 4%).

Welfare comparisons: empirical strategy. We compare welfare gains from the implemented CCT to a hypothetical UCT that removes conditionality and pays all eligibles. This is one of many possible comparisons; alternatives include expanding or restricting eligibility, targeting households likely at the margin of human capital investment (De Janvry and Sadoulet, 2006) or with higher expected welfare gains (Haushofer et al., 2022), or combining unconditional transfers with conditional "top-ups" (Bergstrom and Dodds, 2021). We focus on this comparison because it most directly relates to our key mechanism—conditionality's effect on inequality—and its policy relevance.[7]

We take the set of *eligible* recipients as predetermined and study the welfare within this group, similar to Bergstrom and Dodds (2021). The counterfactual UCT has the same aggregate transfer budget: if t_i^{cct} is the transfer paid to a conditions-met child i in a CCT and s is the share meeting conditions, we set the UCT transfer t_i^{uct} equal to $s \cdot t_i^{cct}$.[8]

Let h denote households and i denote (eligible) children. Each child belongs to a single household and $\mathcal{I}_h$ denotes the set of children in household h. $a_i \in \{0,1\}$ indicates if i meets conditions ($=1$) or not ($=0$). Empirically, we set household h's consumption to $c_h = c_h^{baseline} + t_h$, where $c_h^{baseline}$ is the per capita household consumption we directly observe. Under a CCT, $t_h = \sum_{i \in \mathcal{I}_h} a_i t_i^{cct}$. Under a UCT, $t_h = \sum_{i \in \mathcal{I}_h} t_i^{uct}$, with t_i^{cct} and t_i^{uct} denoting the CCT and UCT transfers divided by i's household size.[9]

We measure the ratio between the CCT welfare gain (relative to no transfers) and the

[7] Most programs in Table A.1 fit into a purely conditional and unconditional transfer classification.

[8] Transfers have an i subscript since, in some contexts, payments vary with age and gender (Appendix B). We calculate UCT counterfactual payments that maintain these proportional differences. When analyzing subsamples, we follow the same procedure and do not assume that transfers for younger children outside the subsample are reallocated to older groups.

[9] In Nicaragua, conditionality applies at the household level (all children must meet conditions) and we adapt our analysis to reflect this (Appendix B).

UCT utility gain (relative to no transfers):

$$R = \frac{\sum_{h=1}^{H} size_h [u(c_h^{baseline} + \sum_{i \in \mathcal{I}_h} a_i t_i^{cct}) - u(c_h^{baseline})]}{\sum_{h=1}^{H} size_h [u(c_h^{baseline} + \sum_{i \in \mathcal{I}_h} t_i^{uct}) - u(c_h^{baseline})]} \tag{1}$$

where H is the number of households in a sample and $size_h$ is household size, including members who are not eligible children (e.g., adults).[10]

Equation (1) has five desirable properties. First, it is intuitive: a value of one indicates the CCT delivers the same aggregate utility gains as the UCT. A value of 0.5 indicates that the CCT delivers only half the welfare benefit of a UCT. Second, it provides a comparable metric across contexts. Third, it is invariant to positive affine transformations of $u(\cdot)$. Fourth, it incorporates the entire consumption distribution (not only mean differences) without distributional assumptions. Fifth, it flexibly accounts for household composition, allowing for multiple eligible children, who may or may not meet conditions, within a household.

Our approach is static and abstracts from borrowing and saving (i.e., the marginal propensity to consume is one),[11] from labor supply effects,[12] and from the costs of monitoring and enforcing conditionality. The latter can be sizable and likely to "tilt the scales" towards UCTs' effectiveness.[13] At this stage, we abstract from conditionality incentivizing human capital gains. While we address them later, presenting results without them clarifies the paper's main mechanism and may be empirically relevant since evidence on whether CCTs generate human capital gains over UCTs remains unclear (Footnote 3).

Our approach also implicitly equates meeting conditions with receiving a CCT transfer. If the program is less strict and pays conditions-unmet children, our results overstate UCT benefits if conditions-unmet children are poorer. This discrepancy also affects the interpretation of the baseline experiments, as households may initially comply, expecting enforcement, then adjust once they learn conditions are weakly applied.

[10] $size_h$ is included to weight individual household welfare gains by the size of the household.

[11] The evaluations we study provide limited data on borrowing and saving.

[12] Crosta et al. (2024) finds that UCTs (compared to no transfers) increase labor supply. However, this does not speak directly to our CCT versus UCT comparison.

[13] Caldés et al. (2006) estimate that conditionality alone accounts for 20% of Progresa's administrative costs and 2% of its transfer budget. Benhassine et al. (2015) report that Tayssir's CCT administrative costs are 30% higher than its labeled cash transfer and discuss how such costs make standard CCTs expensive.

Welfare comparisons: results. Table 1 presents R from equation (1) for different contexts. Our baseline estimates follow the parametrization of $u(\cdot)$ from Alatas et al. (2019), Hanna and Olken (2018), and Finkelstein et al. (2019): a constant relative risk aversion (CRRA) utility with a coefficient of risk aversion $\rho = 3$.

Our baseline estimates indicate the Malawi CCT only delivers 67.5% of the utility gain of an equally budgeted UCT program, suggesting substantial welfare loss from introducing conditionality. For Mexico, Morocco, and Tanzania, the same figures are in the 93.2% - 96.1% range. The ratio is essentially one in Nicaragua.

In all subsamples that focus on older age groups with higher conditions-unmet shares, the ratios are lower than for the full sample, and a more substantial welfare loss from conditionality ($R = 0.837$) is observed for Mexico. This highlights that targeting groups with lower adherence can reduce the relative welfare gains from CCTs. We return to this issue and discuss drivers of variation across contexts after highlighting an important mechanism.[14]

Understanding mechanisms. Two mechanisms can affect the CCT-to-UCT welfare gains ratio (R from equation 1). The first are differences in the consumption distribution of conditions-met and conditions-unmet children. If conditions-met have higher baseline consumption, CCT payments target richer households with lower marginal utility. Second, even if the baseline consumption distribution for the conditions-met and conditions-unmet is exactly the same, the R would be smaller than one. This is because the CCT only targets a share of eligibles (the conditions-met) while the UCT provides payments to all of them. Given a concave utility function, the gains using the same budget to provide transfers to all households are thus larger. We label the latter as the "unequal-payments effect" and measure it as the R that would occur if the distribution of baseline consumption for conditions-met were the same as for conditions-unmet.[15]

[14] Table A.3 probes the sensitivity of results to $u(\cdot)$'s concavity, reporting Rs based on ρs between one and four (the range found in the social insurance literature, Chetty and Finkelstein, 2013). As expected, ratios become closer to one as u is less concave (lower ρ). The ratios remain substantially smaller than one in Malawi and the Mexican subsample.

[15] A simple example illustrates the unequal-payments effect. Assume two households have the same consumption c. A transfer of t to one increases welfare by less than providing $t/2$ to both under diminishing marginal utility: $u(c + t) - u(c) < 2[u(c + t/2) - u(c)]$.

We gauge the importance of the unequal-payments effects as follows. Let N be the number of eligible children and s the share not meeting conditions. We randomly draw sN children, "pretend" they are conditions-unmet, and recalculate R using equation (1). We repeat this 1,000 times. In *expectation*, the consumption distributions of those randomly labeled as conditions-unmet are the same as those randomly labeled as conditions-met. Thus, the *average* of the 1,000 simulated Rs can be attributed to the unequal-payments effect.

Column (4) of Table 1 presents the average simulated ratio. In Malawi, it is 0.817. This indicates that, even if there was no difference between the consumption distribution of conditions-met and conditions-unmet, a CCT would only deliver 81.7% of UCT welfare gains. For the Mexico subsample, it almost matches baseline R , implying the unequal-payments effect dominates. In other contexts (except Nicaragua), the simulated average modestly moves a baseline $R < 1$ closer to one.

What explains variation in R? Intuitively, three factors shape the CCT-to-UCT welfare ratio R: i) transfer size relative to consumption, ii) share of conditions-unmet households, and iii) how far their consumption distribution lies left of conditions-met.

Malawi combines all three: the largest transfer, highest unmet share, and clearest consumption differences. In other contexts, full-sample Rs are closer to one, reflecting smaller consumption gaps by condition adherence, smaller transfers (notably Morocco and Tanzania), and fewer conditions-unmet (especially Nicaragua).

All subsamples have a lower R than their respective full samples. As subsamples were deliberately chosen for lower conditions-meeting shares, they highlight a potential tradeoff in targeting groups with low baseline adherence. Precisely because fewer children meet the conditions, the welfare costs of conditionality increasing inequality are larger. This effect can be substantial even with modest consumption differences between conditions-met and -unmet households. In the Mexican subsample, small consumption gaps (Figure 3) still yield a relatively low R (≈ 0.83), due to the unequal-payments effect. We next examine whether such inequality costs are offset by human capital gains.

Incorporating human capital gains. We now incorporate conditionality's potential effects on human capital. Appendix C provides further exposition and discussion. We augment our framework so household h's utility is:

$$u(c_h^* + t_h - \sum_{i \in \mathcal{I}_h} a_i f_i) + \sum_{i \in \mathcal{I}_h} a_i D_i \quad (2)$$

where $f_i \geq 0$ is the consumption opportunity cost from meeting conditions (e.g., foregone child labor income) and c_h^* is per capita consumption absent transfers and such costs. As before, $t_h = \sum_{i \in \mathcal{I}_h} a_i t_i^{cct}$ under a CCT and $t_h = \sum_{i \in \mathcal{I}_h} t_i^{uct}$ under a UCT.

D_i is the value of the human capital gains from meeting conditions (in utils). In theory, it can be interpreted broadly, encompassing future income from schooling, health improvements, delayed marriage and fertility, psychological benefits, and the child's preferences over meeting conditions. In practice, we use two measurement approaches, discussed later.

In our welfare comparison, only investments in *marginal* children are relevant. We define a child as marginal if it meets conditions under a CCT but not under a UCT. Other children are *inframarginal*. We assume that any child meeting conditions under a UCT would do so under no transfers (i.e., UCTs do not affect enrollment). This assumption tilts the comparison toward higher relative gains from CCTs (Appendix C).

Under this framework, knowledge of f_i, D_i, and which children are marginal can be used to incorporate human capital gains in our CCT-to-UCT welfare ratio (equation 1). Its denominator (UCT welfare gains) remains the same. The numerator only changes for households with marginal children: their welfare gains under the CCT are now the difference between equation (2) and utility under no transfers.[16]

We do not observe which children are marginal, but have estimates of the *share* of marginal children (S^{marg}), which equal the CCT effect on meeting conditions (relative to no transfers). For each of our contexts, we obtain S^{marg} from the source papers. We then take a Monte Carlo approach. We randomly assign $S^{marg} \cdot N$ conditions-met children to be marginals. We repeat this 1,000 times, calculating 1,000 CCT-to-UCT welfare ratios for

[16] Appendix C discusses the mapping between c_h^* and $c_h^{baseline}$ and that they coincide for inframarginals.

each context, and report the 1st and 99th percentiles of this distribution.[17]

The final components are f_i and D_i. We set $f_i = 0.15 \cdot c_h^{baseline}$, meaning meeting conditions reduces household consumption by 15%, based on Morocco's source paper—the only one reporting the opportunity cost of foregone child labor income. For D_i, we use two distinct approaches, discussed below.

A revealed-preference approach. If child i is marginal, her household chooses not to enroll her under a UCT. This revealed preference implies:

$$D_i \leq u(c_h^* + t_h^{uct} - \sum_{j \in \mathcal{I}_h^{-i}} a_j f_j) - u(c_h^* + t_h^{uct} - \sum_{j \in \mathcal{I}_h^{-i}} a_j f_j - f_i) \equiv D_i^{RP} \tag{3}$$

where $\mathcal{I}_h^{-i}$ denotes the set of children in household h excluding child i. D_i^{RP} is an upper bound on D_i. If i is marginal, by definition, it does not meet conditions under the UCT, which bounds the utility gain from meeting conditions. Importantly, D_i^{RP} is observable in our data given an assumption on f_i.

A return on investment (ROI) approach. We follow a formulation from Bergstrom and Dodds (2021):

$$D_i^{ROI} = \sum_{t=0}^{T} \beta^t [\frac{c(1)_{it}^{1-\rho}}{1-\rho} - \frac{c(0)_{it}^{1-\rho}}{1-\rho}] \tag{4}$$

D_i^{ROI} assumes i lives T more years and a discount factor β. $c(1)_{it}$ and $c(0)_{it}$ are, respectively, the trajectories of consumption if i meets the conditions or not. As in Bergstrom and Dodds (2021), we set $T = 50$, $c(0)_{it} = c_h^{baseline}$ and $c(1)_{it} = (1 + \phi) \cdot c_h^{baseline}$ for all t.

We set $\beta = 0.90$ and $\phi = 0.03$, where ϕ reflects schooling's effect on lifetime household consumption (which is distinct from labor income). Both parameters are chosen at the low end of "plausible" ranges. D_i^{ROI} rises with each, making these assumptions "pro-UCT." Appendix C provides further discussion.

Comparing the approaches. We aim to provide a broader view of conditionality's welfare effects under varying assumptions in addition to the baseline approach, which ab-

[17] For the subsamples, we use S^{marg} that incorporate larger effects on older children (Appendix C).

stracts from human capital but may be empirically relevant given unclear evidence on whether CCTs versus UCTs affect human capital (Footnote 3).

The revealed-preference and ROI approaches rest on different assumptions about the social welfare function and rationales for conditionality. The revealed-preference approach assumes a social welfare function that coincides with households' utility. In doing so, it abstracts from common rationales for conditionality, as discussed in Section 1. The ROI approach can account for such rationales (e.g., households underestimating the returns to schooling or over-discounting children's future consumption). However, it requires making assumptions about returns to schooling, while the revealed-preference approach allows for a broader interpretation of human capital gains with fewer assumptions (Appendix C).

Results incorporating human capital gains. Column (1) of Table 2 shows the share of marginals; column (2) replicates the baseline ratio from Table 1. Columns (3)–(4) report CCT-to-UCT welfare ratios (R) using the revealed-preference approach. Differences between ratios based on the 1st and 99th percentiles are small. For full samples, the revealed-preference ratios are modestly larger than the baseline ones and do not change qualitative conclusions. Exceptions are Morocco and Tanzania, where the 99th percentile is 0.069 and 0.049 higher.

Incorporating human capital raises R more in subsamples than in full samples, as expected given subsamples have higher S^{marg}. Conditional on S^{marg}, the gap between the baseline and revealed-preference ratios depends only on f_i (cost of meeting conditions). Appendix C, in particular Table A.4, discusses robustness to alternative f_i values, concluding these do not materially affect results.[18]

Columns (5)–(6) report R using the ROI approach, which yields larger ratios than both baseline and revealed-preference cases. Malawi is the only context where R remains well below one. In Mexico, the ROI approach suggests CCT and UCT gains are similar, whereas baseline and revealed-preference indicate UCTs dominate (especially in the subsample). In all other cases, the ROI approach indicates CCTs yield greater gains, often

[18] For the Tanzania subsample, the 1st percentile is lower than the baseline due to randomization draws where multiple marginals are assigned to the same household (Appendix C).

by sizable amounts.

If the ROI approach reflects true schooling gains, its larger Rs relative to revealed-preference ratios imply households underestimate returns and/or over-discount children's future consumption. While assumptions about equation (4) may not perfectly capture returns, they were set at the low end of "plausible" ranges, so alternative parameterizations would likely widen the gap.

Section 4 further discusses the interpretation of the baseline, revealed-preference, and ROI approaches and potential policy implications.

Externalities. A common rationale for CCTs is that the social planner values induced human capital gains more than households do. While the ROI approach captures some reasons (e.g., households underestimate returns), it abstracts from another: externalities (social returns to human capital).

A micro-founded treatment of externalities is beyond this paper's scope, as it would require numerous assumptions about hard-to-observe parameters. Instead, we adopt a "reduced-form" approach and adapt a formulation from Banerjee et al. (2024). We assume a social planner whose welfare from household h is given by:

$$u(c_h^* + t_h - \sum_{i \in \mathcal{I}_h} a_i f_i) + \theta \sum_{i \in \mathcal{I}_h} a_i D_i \tag{5}$$

where θ measures the extent the social planner values human capital gains D_i more than households. If $\theta = 1$ it values it equally (equations 2 and 5 coincide), if $\theta > 1$, there are positive externalities.

A CCT-to-UCT welfare ratio incorporating externalities (R^{EXT}) can be obtained by calculating $R^{EXT} = R^{HC} + (\theta - 1)GHC$, where R^{HC} is a ratio incorporating human capital gains (Table 2) and GHC is the sum of D_i for all marginal children divided by the UCT welfare gains (the denominator in all our welfare ratios). Appendix C provides details and Tables A.5 and A.6 report GHC across contexts and approaches.[19]

These figures, coupled with a choice of parameter θ, allow one to incorporate exter-

[19] In contexts where the share of marginals is large and/or transfers are small (Morocco, Nicaragua, and Tanzania), the GHC is more sizable, but so are the utility costs of enrollment f_i (Appendix C)

nalities in our welfare ratio. For example, in the case of the revealed-preference approach for the Mexican full sample, $R^{HC} = 0.954$ and $GHC = 0.044$. Thus, a social welfare function that values human capital gains three times more than households do ($\theta = 3$) yields $R^{EXT} = 0.954 + (3-1) \cdot 0.044 = 1.04$.[20]

Moreover, one can compute the θ that equates the planner's valuation of CCT and UCT (i.e., solves $R^{EXT} = 1$). For Malawi under the ROI approach, with $R^{HC} = 0.823$ and $GHC = 0.172$, this yields $\theta = 1 + \frac{1-R^{HC}}{GHC} = 2.03$ (planner that values D_i at twice the ROI estimate). In other contexts, ROI-based R^{HC} is near or above one, suggesting conditionality can be "justified" without externalities. Under the revealed-preference approach, $\theta = 9.2$ is needed for $R^{EXT} = 1$ in Malawi. In Mexico's full and subsamples, such θs are roughly two and three, respectively. In other cases, a $\theta \approx 1$ suffices.

We remain agnostic about the appropriate value of θ, as the literature offers little guidance on the quantitative magnitude of social returns to human capital induced by CCTs.

Other forms of human capital. Under the reduced-form approach, D_i^{RP} can be interpreted broadly to include not only the effect of school enrolment but also effects on attendance, test scores, cognitive ability, and even health.[21] The ROI approach, on the other hand, involves specific assumptions tied to returns to schooling.

Explicitly incorporating effects on other forms of human capital (e.g., health) is beyond this paper's scope, as it requires strong assumptions about hard-to-observe outcomes. However, health-related conditions in CCTs often apply to children under five who are outside our samples. In principle, the externalities analysis could be adapted to include health effects—for example, assuming gains from health equal 50% of those from schooling would correspond to a $\theta = 1.5$.

4 Conclusion

We conclude by discussing three potential policy implications. First, our results highlight a potentially underappreciated mechanism: conditionality can exacerbate inequality

[20] We use the 99th-percentile draws for all discussions here.

[21] A caveat is that such returns must accrue to marginal children only (Appendix C).

among eligible households, with quantitatively meaningful effects on overall welfare.

Second, this mechanism can be amplified by efforts to improve CCT targeting. While focusing on groups with low baseline compliance (e.g., older children) may enhance human capital impacts, it can also intensify the inequality-increasing effect of conditionality—precisely because fewer children meet the conditions. This occurs even without consumption differences between conditions-met and -unmet, due to the "unequal-payments effect."

Third, our three approaches to human capital—baseline, revealed preference, and ROI—underscore the importance of clarifying the rationales behind conditionality. If it has no discernible effect on human capital, the baseline results suggest welfare gains from switching to UCTs. The divergence between revealed-preference and ROI estimates shows that, even if conditionality boosts human capital, justifying it on welfare grounds may require assuming households deviate from optimal decisions (e.g., underestimating returns, over-discounting children's future utility, or failing to internalize externalities).

Lastly, while this paper focuses on one policy comparison, the framework can be extended to study changes in eligibility, targeting, and other design features. Our results suggest that alternatives such as restricting eligibility to poorer households or combining unconditional transfers with conditional top-ups merit further exploration.

References

Akresh, Richard, Damien De Walque, and Harounan Kazianga, "Evidence from a Randomized Evaluation of the Household Welfare Impacts of Conditional and Unconditional Cash Transfers Given to Mothers or Fathers," Policy Research Working Paper 7730, World Bank 2020.

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A Olken, Ririn Purnamasari, and Matthew Wai-Poi, "Does elite capture matter? Local elites and targeted welfare programs in Indonesia," in "AEA Papers and Proceedings," Vol. 109 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2019, pp. 334–339.

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky, "Regressive sin taxes, with an application to the optimal soda tax," *The Quarterly Journal of Economics*, 2019, *134* (3), 1557–1626.

American Institutes for Research, *Zambia's Child Grant Program: 24 Month Impact Report*, American Institutes for Research, 2013.

Asfaw, Solomon, Benjamin Davis, Joe Dewbre, Sudhanshu Handa, and Paul Winters, "Cash Transfer Programme, Productive Activities and Labour Supply: Evidence from Randomized Experiment in Kenya," *Journal of Development Studies*, August 2014, *50* (8), 1172–1196.

Baird, Sarah, Craig McIntosh, and Berk Özler, "Cash or condition? Evidence from a cash transfer experiment," *The Quarterly journal of economics*, 2011, *126* (4), 1709–1753.

_ , Francisco HG Ferreira, Berk Özler, and Michael Woolcock, "Relative effectiveness of conditional and unconditional cash transfers for schooling outcomes in developing countries: a systematic review," *Campbell Systematic Reviews*, 2013, *9* (1), 1–124.

Banerjee, Abhijit, Rema Hanna, Benjamin A Olken, and Diana Sverdlin Lisker, "Social Protection in the Developing World," *Journal of Economic Literature*, 2024.

Banerjee, Abhijit V, Rema Hanna, Gabriel E Kreindler, and Benjamin A Olken, "Debunking the stereotype of the lazy welfare recipient: Evidence from cash transfer programs," *The World Bank Research Observer*, 2017, *32* (2), 155–184.

Barrera-Osorio, Felipe, Marianne Bertrand, Leigh L. Linden, and Francisco Perez-Calle, "Improving the Design of Conditional Transfer Programs: Evidence from a Randomized Education Experiment in Colombia," *American Economic Journal: Applied Economics*, 2011, *3* (2), 167–195.

Behrman, Jere R, Susan W Parker, and Petra E Todd, "Schooling impacts of conditional cash transfers on young children: Evidence from Mexico," *Economic development and cultural change*, 2009, *57* (3), 439–477.

Benhassine, Najy, Florencia Devoto, Esther Duflo, Pascaline Dupas, and Victor Pouliquen, "Turning a shove into a nudge? A "labeled cash transfer" for education," *American Economic Journal: Economic Policy*, 2015, *7* (3), 86–125.

Bergstrom, Katy and William Dodds, "Optimal taxation with multiple dimensions of heterogeneity," *Journal of Public Economics*, 2021, *200*, 104442.

Caldés, Natàlia, David Coady, and John A Maluccio, "The cost of poverty alleviation transfer programs: a comparative analysis of three programs in Latin America," *World development*, 2006, *34* (5), 818–837.

Chaudhury, Nazmul, Jed Friedman, and Junko Onishi, "Philippines Conditional Cash Transfer Program Impact Evaluation 2012," Technical Report 75533-PH, World Bank, Manila 2013.

Chetty, Raj and Amy Finkelstein, "Social insurance: Connecting theory to data," in "Handbook of public economics," Vol. 5, Elsevier, 2013, pp. 111–193.

Covarrubias, Katia, Benjamin Davis, and Paul Winters, "From Protection to Production: Productive Impacts of the Malawi Social Cash Transfer Scheme," *Journal of Development Effectiveness*, 2012, *4* (1), 50–77.

Crosta, Tommaso, Dean Karlan, Finley Ong, Julius Rüschenpöhler, and Christopher R Udry, "Unconditional cash transfers: A Bayesian meta-analysis of randomized evaluations in low and middle income countries," Technical Report, National Bureau of Economic Research 2024.

Deaton, Angus, "The analysis of household surveys: a microeconometric approach to development policy," *The World Bank*, 1997.

Edmonds, Eric V. and Norbert Schady, "Poverty Alleviation and Child Labor," *American Economic Journal: Economic Policy*, 2012, *4* (4), 100–124.

Evans, David, Stephanie Hausladen, Katrina Kosec, and Natasha Reese, *Community-Based Conditional Cash Transfers in Tanzania: Results from a Randomized Trial*, Washington, DC: World Bank Publications, 2014.

Finkelstein, Amy, Nathaniel Hendren, and Erzo FP Luttmer, "The value of medicaid: Interpreting results from the oregon health insurance experiment," *Journal of Political Economy*, 2019, *127* (6), 2836–2874.

Galiani, Sebastian and Patrick J. McEwan, "The Heterogeneous Impact of Conditional Cash Transfers," *Journal of Public Economics*, 2013, *103*, 85–96.

Gilligan, Daniel O. and Shalini Roy, "Resources, Stimulation, and Cognition: How Transfer Programs and Preschool Shape Cognitive Development in Uganda," 2013.

Hanna, Rema and Benjamin A Olken, "Universal basic incomes versus targeted transfers: Anti-poverty programs in developing countries," *Journal of Economic Perspectives*, 2018, *32* (4), 201–226.

Haushofer, Johannes and Jeremy Shapiro, "Household Response to Income Changes: Evidence from an Unconditional Cash Transfer Program in Kenya," Technical Report, Massachusetts Institute of Technology 2013.

_ , Paul Niehaus, Carlos Paramo, Edward Miguel, and Michael W Walker, "Targeting impact versus deprivation," Technical Report, National Bureau of Economic Research 2022.

Herd, Pamela and Donald Moynihan, "Administrative Burdens in the Social Safety Net," *Journal of Economic Perspectives*, 2025, *39* (1), 129–150.

Janvry, Alain De and Elisabeth Sadoulet, "Making conditional cash transfer programs more efficient: designing for maximum effect of the conditionality," *The World Bank Economic Review*, 2006, *20* (1), 1–29.

Kolsrud, Jonas, Camille Landais, Daniel Reck, and Johannes Spinnewijn, "Retirement consumption and pension design," *American Economic Review*, 2024, *114* (1), 89–133.

Macours, Karen, Norbert Schady, and Renos Vakis, "Cash Transfers, Behavioral Changes, and Cognitive Development in Early Childhood: Evidence from a Randomized Experiment," *American Economic Journal: Applied Economics*, 2012, *4* (2), 247–273.

Maluccio, John and Rafael Flores, "Impact Evaluation of a Conditional Cash Transfer Program: The Nicaraguan Red de Protección Social," 2005.

Parker, Susan W and Petra E Todd, "Conditional cash transfers: The case of Progresa/Oportunidades," *Journal of Economic Literature*, 2017, *55* (3), 866–915.

Robertson, Laura, Phyllis Mushati, Jeffrey W Eaton, Lovemore Dumba, Gideon Mavise, Jeremiah Makoni, Christina Schumacher, Tom Crea, Roeland Monasch, Lorraine Sherr et al., "Effects of unconditional and conditional cash transfers on child health and development in Zimbabwe: a cluster-randomised trial," *The Lancet*, 2013, *381* (9874), 1283–1292.

Skoufias, Emmanuel, Mishel Unar, and Teresa Gonzalez de Cossio, "The Poverty Impacts of Cash and In-Kind Transfers: Experimental Evidence from Rural Mexico," *Journal of Development Effectiveness*, 2013, *5* (4), 401–429.

UNESCO Institute for Statistics, "Gross enrolment ratio in primary education," Dataset 2025. With minor processing by Our World in Data. Original source: UNESCO Institute for Statistics (UIS) – Education.

World Bank Office Jakarta, *Main Findings from the Impact Evaluation of Indonesia's Pilot Household Conditional Cash Transfer Program*, World Bank, 2011.

Figure 1: Share of Eligible Children Not Meeting Conditions

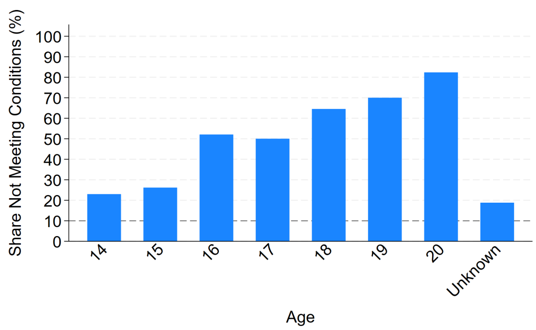

(a) Malawi
(N: 470)

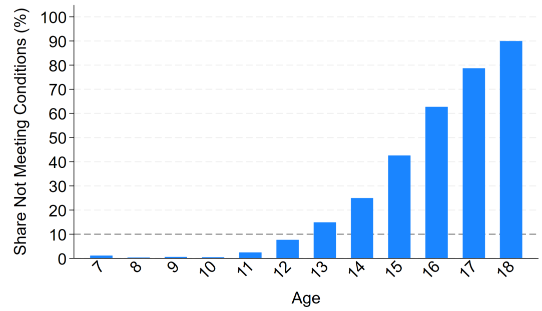

(b) Mexico
(N: 11,093)

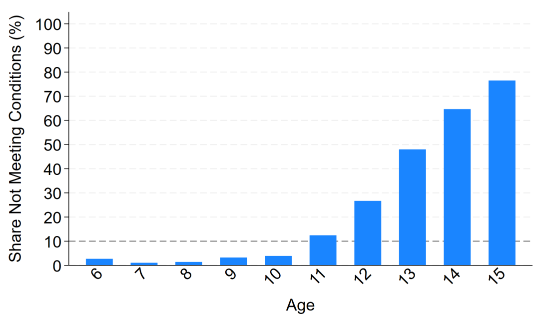

(c) Morocco
(N: 5,988)

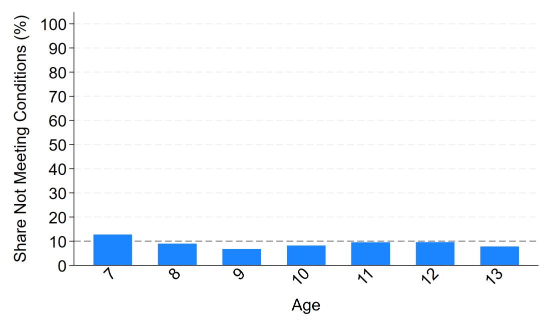

(d) Nicaragua
(N: 712)

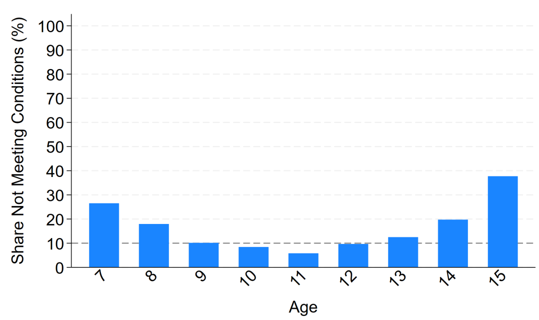

(e) Tanzania
(N: 753)

Each subfigure shows the percentage of children meeting schooling conditions by age. The unit of observation is a child, and N is the number of observations in each sample. Only eligible children from the CCT arm are included. Dashed horizontal lines indicate the 10% level (used to define subsamples, see main text).

Figure 2: Baseline Consumption Distribution by Condition Adherence: Full Samples

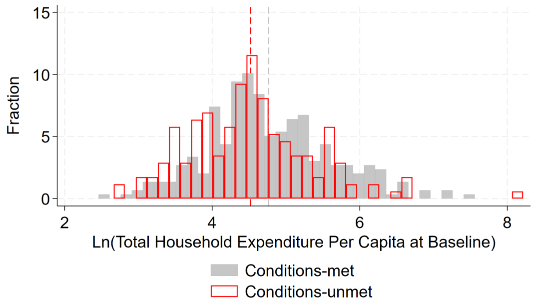

(a) Malawi
(Diff. in Means: 0.25)

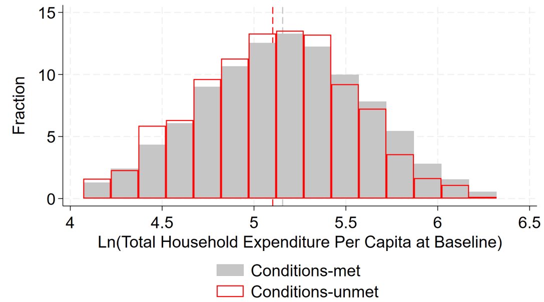

(b) Mexico
(Diff. in Means: 0.06)

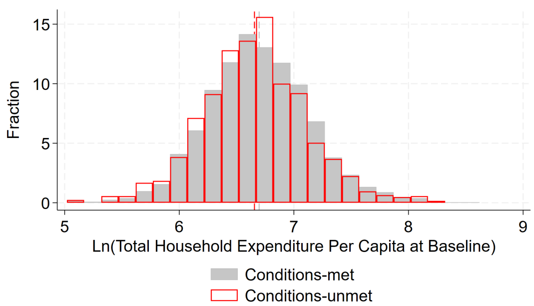

(c) Morocco
(Diff. in Means: 0.04)

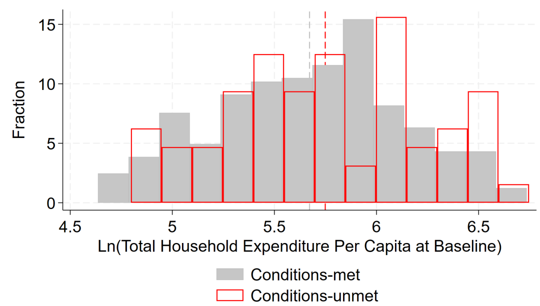

(d) Nicaragua
(Diff. in Means: -0.08)

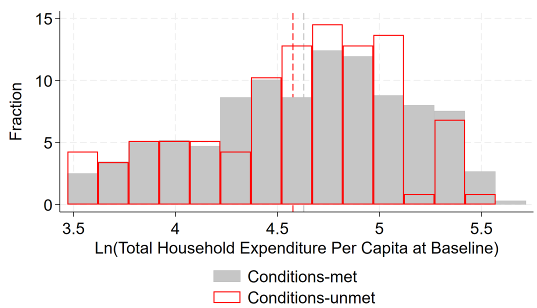

(e) Tanzania
(Diff. in Means: 0.05)

The unit of observation is a child. Each subfigure presents histograms of the logarithm of children's total household consumption per capita at baseline (i.e., before CCT payments start). Only eligible children from the CCT arm are included. Solid gray-colored bars provide the distribution for conditions-met children, and the red-bordered bars do so for conditions-unmet children. Vertical lines indicate the mean consumption for each group. "Diff. in Means" is the conditions-met mean minus the conditions-unmet mean (thus measured in log points).

Figure 3: Baseline Consumption Distribution by Condition Adherence: Subsamples

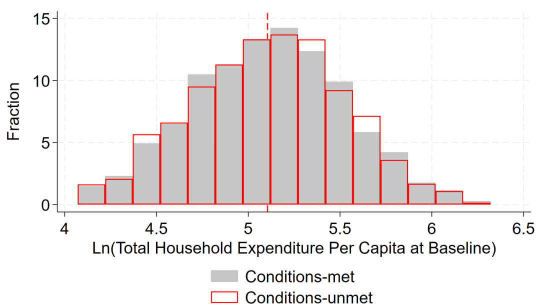

(a) Mexico (ages 13-18)
(Diff. in Means: 0.01)

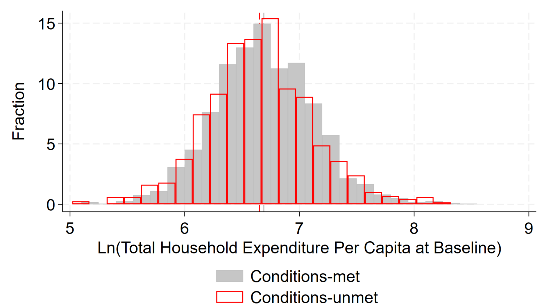

(b) Morocco (ages 11-15)
(Diff. in Means: 0.04)

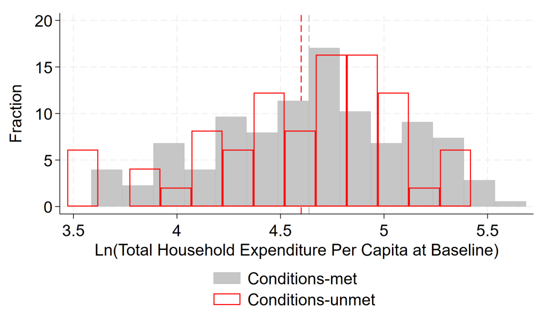

(c) Tanzania (ages 13-15)
(Diff. in Means: 0.04)

The unit of observation is a child. Each subfigure presents, for a given subsample, histograms of the logarithm of children's total household consumption per capita at baseline (i.e., before CCT payments start). Only eligible children from the CCT arm are included. Solid gray-colored bars provide the distribution for conditions-met children, and the red-bordered bars do so for conditions-unmet children. Vertical lines indicate the mean consumption for each group. "Diff. in Means" is the conditions-met mean minus the conditions-unmet mean (thus measured in log points).

Table 1: Baseline Welfare Comparisons

	Share Cond.-Unmet	Avg. Transfer/Cons.	CCT-to-UCT Welfare Ratio (R)	
			Baseline	Avg. Simulated Ratio
	(1)	(2)	(3)	(4)
Panel A: Full Samples				
Malawi	0.368	0.252	0.675	0.817
Mexico	0.197	0.184	0.932	0.989
Morocco	0.208	0.030	0.961	1.017
Nicaragua	0.090	0.112	1.002	0.960
Tanzania	0.156	0.062	0.958	0.981
Panel B: Subsamples				
Mexico (ages 13-18)	0.423	0.241	0.837	0.830
Morocco (ages 11-15)	0.404	0.033	0.898	0.968
Tanzania (ages 13-15)	0.218	0.056	0.912	0.972

Column (1) presents the share of conditions-unmet children in the sample. Column (2) provides the average CCT transfer size expressed as a share of consumption. Specifically, for each child, we take the respective annual CCT payment she would receive if meeting conditions and divide it by her total annual household consumption. We report the sample averages in column (2). Column (3) presents the ratio between the welfare gains from the observed CCT and a counterfactual UCT program that uses the same budget to provide payments to all eligible children (R from equation 1). Column (4), the average simulated ratio, presents the mean simulated R when we randomly assign conditions-unmet status, thus providing the counterfactual R due to the "unequal-payments effect" only. Columns (3)–(4) are based on a CRRA utility with $\rho = 3$.

Table 2: Welfare Comparisons Incorporating Human Capital Gains

		CCT-to-UCT Welfare Ratio (R)				
			Revealed Preference		Return on Investment	
	S^{marg}	Baseline	1st pctile	99th pctile	1st pctile	99th pctile
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Samples						
Malawi	0.058	0.675	0.677	0.689	0.693	0.823
Mexico	0.032	0.932	0.947	0.954	0.997	1.017
Morocco	0.054	0.961	1.003	1.030	1.180	1.285
Nicaragua	0.128	1.002	1.009	1.016	1.191	1.364
Tanzania	0.040	0.958	0.969	1.007	1.024	1.181
Panel B: Subsamples						
Mexico (ages 13–18)	0.065	0.837	0.857	0.865	0.938	0.971
Morocco (ages 11–15)	0.112	0.898	0.955	0.985	1.348	1.522
Tanzania (ages 13–15)	0.134	0.912	0.873	0.982	1.122	1.559

Column (1) presents the share of marginal children (S^{marg}): those that meet conditions (e.g., enroll in school) under the CCT but not under a UCT. For reference, column (2) replicates the baseline CCT-to-UCT welfare ratio (R) from column (3) in Table 1. It provides R assuming no gains from human capital. Columns (3) and (4) present R using the revealed-preference approach, while columns (5) and (6) do so using the return on investment (ROI) approach. Columns (3) and (5) present the 1st percentile of the distribution of simulated R based on randomly assigning which conditions-met children are marginal, while columns (4) and (6) do so for the 99th percentile. Columns (2)–(6) are based on a CRRA utility with $\rho = 3$.

Online Appendix

A Comparing Programs Included and Not Included in the Analysis

Table A.1 lists the programs considered for the analysis and the reasons for excluding others. See also discussion in Section 2. Additional literature searches did not yield other suitable experiments.

Table A.2 reports shares of marginals and conditions-unmet children for our five contexts (replicating main-text tables) and for the five non-included studies where such data are available. Of the 14 programs in Table A.1 that we considered but could not analyze, seven were experiments with a UCT arm and a pure control, so they lacked a CCT arm and we could not infer shares meeting conditions or marginals. Of the remaining seven, the Indonesia (PKH) paper notes that conditions were not enforced, and the Nicaragua (Atención a Crisis) source paper focused on cognitive impacts and did not report enrollment effects, preventing calculation of marginals. This leaves five non-included studies for which we report shares of marginals and conditions-unmet children in the CCT treatment.

On average, the share of marginals is 6.2% in our five contexts, compared to 8.0% in non-included studies. The average share of children not meeting conditions is 19.5% in our contexts versus 20.7% in non-included ones. Distributions are broadly similar. Marginals range from 3.2%–12.8% in our contexts and 1.7%–18.0% in non-included programs. Conditions-unmet range from 9.0%–36.8% in our contexts and 2.5%–32.0% in non-included ones.

This suggests that the share of marginals and the fraction of households not meeting conditions in our contexts are typical of CCT settings. We highlight that the Burkina Faso program has a share of conditions-unmet of 32.0%, which is similar to our context with the highest share of conditions-unmet (Malawi, at 36.8%), indicating such shares may not be unusual.

A caveat is that age ranges differ somewhat across programs, though both groups include one context with older children (Malawi and Colombia) and others covering roughly ages 6–15. Note that we cannot compare programs in Table A.2 to CCTs that were not part of randomized evaluations (e.g., Brazil's Bolsa Família).

B Further Information on the Five Contexts/Experiments

This appendix provides further discussion and information on data construction for the five contexts (experiments) that we analyse. Section 2 details data harmonization. We expand on two aspects of the data here.

Outliers. We exclude observations below the 5th and above the 95th percentile of context-specific consumption distributions to reduce sensitivity to outliers. Some exclusion is inevitable, as the minimums of the distributions likely reflect measurement error (e.g., Mexico's minimum is USD 0.57 per capita per year; Tanzania's is USD 7.96). Under concave utility, such outliers can disproportionately affect welfare calculations.

Measurement error in school enrollment variables. As discussed in Section 2, the source papers for Morocco and Malawi suggest there is little measurement error in survey-based school enrollment data, at least on average. In Malawi, Baird et al. (2011) measures enrollment using school surveys for a subsample of participants. Mean self-reported and teacher-reported enrollment rates are virtually the same in the CCT arm (the focus of our paper). In Morocco, Benhassine et al. (2015) estimates the effects on enrollment and dropout behavior using both household surveys and school records and finds similar results. We are unaware of papers documenting how misreporting of enrollment in surveys correlates with income or consumption.[1]

Malawi (SIHR). The source paper for this experiment is Baird et al. (2011). The experiment occurred in 176 enumeration areas (EAs) in the Zomba district. The study focused

[1] Data from school surveys from Malawi, which only covered 43% of participants, is unavailable. We cannot use school records data from Morocco since the replication files do not link it to the household survey on consumption (understandably, given privacy considerations).

only on never-married women who were enrolled in school at baseline. Randomization was at the EA level: 46 were assigned to a CCT arm, 27 to a UCT arm, and 15 to control. Our sample includes 470 eligible girls across 425 households from the CCT arm sampled by the evaluators.

The CCT arm involved two types of payments. First, parents of girls attending school received 4, 6, 8, or 10 USD per month. The amount was randomly determined at the EA level. Second, girls received an individual payment of 1, 2, 3, 4, or 5 USD per month. The amount was determined by a public lottery within the community. Payments were conditional on school enrollment and attendance greater than 80% of school days. We use the expected payment to define the CCT transfer in our analysis.

We use consumption data from the baseline survey collected between October 2007 and January 2008. Specifically, we use the household's response to the question *"What were your total monthly household expenditures last month?"*

We define meeting conditions as being enrolled in school in Term 1, 2010. The source paper also uses the same variable to define not meeting conditions in their analysis of effects on pregnancy and marriage.

School enrollment rates in the sample are relatively high compared to the broader Malawi population. The source paper reports that net secondary enrollment in Malawi was 24% at the time of the study, while 63% of the sample was enrolled. This likely reflects the experiment's sample selection criteria, which included only girls "*who were in school at baseline*" (Baird et al., 2011) and had never been married.

The transfer sizes is substantial. The average (annualized) CCT transfer for one child divided by the (annualized) household consumption is 25.2%.

Mexico (Progresa). The Progresa CCT has been studied by several papers. See Parker and Todd (2017) for a survey. For the experimental evaluation, 506 rural communities from seven Mexican states were selected in 1997, with 320 randomly assigned to treatment (receiving treatment starting in May 1998) and the other 186 to control (receiving benefits starting in January 2000). Our sample includes 11,093 eligible children across 4,993 households in treatment communities.

Payments were conditional on a child aged 7-18 enrolling in school between 3rd and 9th grade and on attendance greater than 85% of school days. Payments depended on the grade of enrollment.[2]

We use consumption data from the March 1998 ENCEL survey, before payments started. Specifically, we use the household's response to the question "Generally, how much money do you have available for total household expenses per week?"

We define meeting conditions as being enrolled in school in the May 1999 ENCEL survey (which we treat as an endline survey since it is the last one before the control group starts receiving treatment). There is near-universal enrollment by those aged 6-11. The source paper indicates this is representative of the larger Mexican population.

A substantial share of students drop out at the end of primary school (6th grade in Mexico). Thus the share not meeting conditions at age 12 and above is higher. The average enrollment rate for those aged 12-18 in the sample is 65.4%, which is similar to the Mexico-wide enrollment ratio in 1998 of 67.6% (UNESCO Institute for Statistics, 2025). Overall, this suggests the Progresa study context is representative of national school enrollment rates at the time of the study.

The transfer sizes are substantial. The average (annualized) CCT transfer for one child divided by the (annualized) household consumption is 18.4%.

Morocco (Tayssir). The source paper for this experiment is Benhassine et al. (2015). The program targeted children of primary school age (6 to 15). The experiment was implemented in 320 "school sectors" in the five poorest regions of Morocco. Randomization was at the sector level: 60 were assigned to a control group and the remaining 260 randomized to the four following treatments: a CCT group where the payment was made to the father (90 sectors), a CCT group where the payment was made to the mother (90 sectors), a "labeled" cash transfer group where the payment was made to the father (40

[2] For primary school students, it varied from 75 Mexican pesos (8.2 USD) per month for 3rd grade to 150 pesos (16.4 USD) for 6th grade. For secondary school, payments also depended on grades and were larger for girls than boys. It varied from 220 Mexican pesos (24 USD) for boys in 7th grade to 285 pesos (31.1 USD) for girls in 9th grade. Families also received yearly payments of 135 Mexican pesos (15 USD) for each child in primary school and 170 Mexican pesos (18.6 USD) for each child in secondary school. We can assign a grade to each child by adding one to the last completed grade, thus determining the grade a child should be enrolled at.

sectors), and a "labeled" transfer where the payment was made to the mother (40 sectors). The purpose of the "labeled" transfer group is to test the effects of conditionality on school participation. Our sample includes 5,988 children across 2,421 households from both the CCT to fathers and CCT to mothers arms.[3]

Payments were 60 Moroccan dirhams (8 USD) per month for grades 1 and 2 (6 to 7-year-olds), 80 dirhams (10 USD) for grades 3 and 4 (8 to 9-year-olds), and 130 dirhams (13 USD) for children 10 to 15-years-old. Parents in all treatment arms had to enroll their children in the program, at their school, annually to receive payments. The enrollment process in the CCT groups made clear that payment receipt was conditional on enrollment and "*no more than four absences in the month*."

Our consumption variable is taken directly from the source paper's replication files. We use monthly per capita consumption from the June 2008 baseline survey, reported in Table 2 of Benhassine et al. (2015).

We define meeting conditions using the main school participation outcome in the source paper: whether a child was reported as being enrolled and having attended school at least once in "*end of program year 2*" (May 2010).

There is near-universal enrollment by those aged 6-10, which is representative of the Morocco-wide primary school enrollment rates. This study's secondary enrollment rate is .49, which is relatively similar to the national secondary enrollment rates of .577 at the time (UNESCO Institute for Statistics, 2025).

However, the experiment was conducted in the poorest municipalities within Morocco's five poorest regions, making the areas unrepresentative of the country as a whole. Recipients were required to enroll their children in their intended school, and headmasters advertised the program. Although baseline enrollment was not mandatory (unlike in Malawi), the design likely selected households more inclined to enroll their children than others in the same municipalities.

A key feature of the Morocco experiment, also noted in Benhassine et al. (2015), is the relatively small payment size. The average (annualized) CCT transfer for one child

[3] One CCT-to-fathers sector could not be reached due to flooding during the baseline survey, leaving 89 sectors in this arm.

divided by the (annualized) household consumption is 3.0%.

Nicaragua (RPS). The source paper for this experiment is Maluccio and Flores (2005). The Red de Protección Social (RPS) pilot took place in two of Nicaragua's poorest departments (Madriz and Matagalpa). It targeted 42 *comarcas*, randomly assigning 21 to the treatment and 21 to the control group. Our sample includes 712 eligible children across 389 households in the treated *comarcas*.

RPS featured two core components. First, the education component provided an annual C$1,440 (112 USD) per household with children aged 7-13 who had not yet completed the fourth grade of primary school, conditional on enrollment and maintaining at least 85% attendance for all such children. Additionally, there was an annual payment of C$275 (21 USD) per child for school supplies. Thus, unlike other programs, all eligible children must meet the conditions for the household to receive the transfer. Second, a food security, health, and nutrition component offered C$2,880 (224 USD) per household each year on the condition that parents attend bi-monthly health workshops and bring children under five for routine checkups.

Our analysis considers only the first component. To account for the household-level conditionality, we adapt the analysis described in Section 3 as follows. Empirically, we set household h's consumption under a CCT equal to $c_h^{cct} = c_h^{baseline} + t_h^{cct}$ if *all* eligible children in h meet the conditions, where t_h^{cct} is the first component transfer divided by household size. If at least one eligible child in h does not meet conditions, then $c_h^{cct} = c_h^{baseline}$. Household h consumption under a UCT is $c_h^{uct} = c_h^{baseline} + t_h^{uct}$ for all h, where t_h^{cct} is the UCT transfer (see Section 3). In computing Figures 1 and 2, and calculating the share of conditions-unmet, we also define child i from household h as conditions-unmet if all eligible children in h meet the conditions.

We use consumption data from the baseline survey (before treatment) conducted in August and September 2000. The survey included detailed itemized questions on the consumption of several food and non-food items. We use these questions to construct total household consumption.

We define meeting conditions as being enrolled in school in the endline survey conducted in October 2002. Compared to other contexts, the share of conditions-unmet children in the RPS program is small, at only 9.0%. Moreover, our sample only includes 64 children who do not meet the conditions, making comparisons less precise.

The high enrollment rate reflects two factors. First, eligibility was limited to children ages 7–13 who had not completed fourth grade, targeting those likely to attend primary school. While Nicaragua had near-universal primary enrollment at the time (UNESCO Institute for Statistics, 2025), the control group's enrollment was 80%, likely because the experiment took place in a poorer region.

Second, this is where the CCT had its largest impact on enrollment: 12.8 p.p. (see Table 2 and Appendix C). Thus, the study area appears to have lower baseline enrollment than the national average, but enrollment in the CCT arm approached the national level due to the program's large effect.

The average (annualized) CCT transfer for one child divided by the (annualized) household consumption is 11.2%.

Tanzania (TASAF). The source paper for this experiment is Evans et al. (2014). The TASAF experiment took place in three Tanzanian districts (Bagamoyo, Chamwino, and Kibaha). The program targeted 80 villages, randomly assigning 40 to the treatment and 40 to the control group. Eligibility was based on locally defined criteria for vulnerability, such as orphanhood and chronic illness within the family. Our sample includes 753 eligible children across 392 households in the treated villages.

Payments were determined by household composition. Each child aged 7–15 received USD$3 per month (approximately 50% of the food poverty line), conditional on compliance. Children were required to be enrolled in school and maintain at least 80% school attendance.

We use consumption data from the baseline survey (before treatment) conducted in early 2009. The survey included detailed itemized questions on consumption of several food and non-food items, including the value of own crops that were consumed and gifts. We use these questions to construct total household consumption.

We define meeting conditions as being enrolled in school in the endline survey conducted in late 2012. The source paper discusses that the program aimed to increase primary school enrollment. While the official primary school age in Tanzania is 7–13, children aged 14–15 were also eligible, as 14- and 15-year-olds often have not completed primary school due to late entry, grade repetition, or missed years. Thus, the relatively low shares of conditions-unmet children at ages 14-15 (compared to those of similar age in other contexts) may reflect that children remain in primary school at such ages.[4]

The source paper notes that the three districts where the experiment took place were selected for being particularly poor and vulnerable relative to the rest of Tanzania. Within these districts, the program targeted the poorest and most vulnerable populations. Comparing sample enrollment to national rates is difficult due to limited data.[5]

The average (annualized) CCT transfer for one child divided by the (annualized) household consumption is 6.2%.

C Incorporating Human Capital Gains in the CCT-to-UCT Welfare Ratio

This paper presents three approaches to incorporating human capital investments into the CCT-to-UCT welfare ratio. The first is a baseline version that abstracts from human capital (equation 1 and Table 1, discussed in the main text). The other two are the revealed-preference and return on investment (ROI) approaches, also discussed in the main text and further detailed in this appendix.

Our goal is to present three approaches that offer a broader perspective on the welfare consequences of conditionality under different scenarios and assumptions. Determining which approach is preferable is beyond the scope of this paper.

Regarding the baseline approach that abstracts from human capital, we reiterate that the evidence on whether CCTs, compared to UCTs, have effects on human capital remains unclear (see Footnote 3). Thus, this approach may be empirically relevant in some

[4] According to UNESCO Institute for Statistics (2025), the official primary school age in Mexico, Morocco, and Nicaragua is 6-11.

[5] UNESCO Institute for Statistics (2025) and other cross-country databases report only gross primary enrollment rates for Tanzania, which are hard to interpret in contexts where older children often enroll.

contexts. We also reiterate that, in both the revealed preference and ROI approaches, we make a substantial "pro-CCT" assumption: that the effect of a CCT relative to no transfers equals its effect relative to a UCT (see also further discussion below).

The revealed-preference and ROI approaches rest on different assumptions about the social welfare function and the rationale for conditionality. If the social welfare function coincides with households' utility, the revealed-preference approach is appropriate. While this aligns with the "consumer sovereignty" principle in welfare economics, it excludes common rationales for conditionality, such as households underestimating returns to human capital or discounting future generations' utility more than the social planner (see Footnote 4).

The ROI approach can account for such rationales. For example, returns to schooling and discount factors in equation (4) may differ from those governing household decisions. However, this approach requires additional assumptions. Moreover, while gains under the revealed-preference approach can be interpreted broadly—beyond school enrollment to include attendance, cognitive improvements, delayed marriage, and health—gains under the ROI approach are tied to specific assumptions about returns to schooling.

Lastly, one potential interpretation for cases where the two different approaches produce substantially different CCT-to-UCT welfare ratios is that indeed households perceived returns to schooling and its discount factors differ from those a social planner uses to rationalize conditionality. This highlights empirically that the justification for a CCT may rely on households not making optimal decisions.

Framework. We consider N (eligible) children indexed by i. Each child belongs to a single household. There are H households indexed by h. A household may have multiple eligible children. Let $\mathcal{I}_h \subseteq \{1, \ldots, N\}$ denote the set of children in household h. We denote the set of children in household h, excluding child i, to be $\mathcal{I}_h^{-i} \equiv \mathcal{I}_h \setminus \{i\}$.

The household decides whether each of its children will take a binary action $a_i \in \{0, 1\}$. $a_i = 1$ implies meeting the conditions of the CCT program, and $a_i = 0$ implies not meeting such conditions. At this stage, a_i can be interpreted broadly, including complying with health conditions, attendance at school, improved test scores, and more. We later discuss specific interpretations under each approach.

If household h chooses $a_i = 0$ for all its children and does not receive any transfers, it has a per capita consumption of c_h^* . We later discuss how c_h^* relates to observed baseline consumption. For each child i that household h chooses $a_i = 1$, it receives a utility of D_i, but also pays a consumption cost f_i. Its utility is thus:

$$u(c_h^* + t_h - \sum_{i \in \mathcal{I}_h} a_i f_i) + \sum_{i \in \mathcal{I}_h} a_i D_i$$

which is equation (2) from the main text. t_h is the transfers it receives (the definition of which we return to later). $f_i \geq 0$ is the consumption opportunity cost of child i meeting the conditions (e.g. foregone income from child labor).

Interpreting D_i. In the revealed-preference approach, D_i can be interpreted as how much the household values having i meet the conditions (in utils). D_i is a combination of how the household's beliefs about the net present value of the utility of future additional income from more human capital, preferences of the child for working, leisure, and school, and also direct utility (e.g., the pride of investing in the child's human capital). Moreover, D_i can be interpreted broadly beyond school enrollment to include attendance, cognitive improvements, delayed marriage and fertility, and health gains.

In the ROI investment, D_i is determined by specific assumptions on the returns to school enrollment. Moreover, D_i can then be interpreted as how much the *social planner* values D_i, which can be distinct from the household valuation. For example, the household may under-estimate the true returns to schooling or over-discount, relative to the social planner, the future consumption of their children.

In both approaches, D_i enters a social welfare function. Under the revealed-preference approach, the D_i from a social welfare function and from the household's utility function driving their choices must be the same. In the ROI approach, we are agnostic about how households make a_i decisions.

Transfer policy scenarios. We consider three possible scenarios: no transfers (nt), uct, and cct. Whether a child meets the condition may depend on the scenario, and we denote a_i^{nt}, a_i^{uct}, and a_i^{cct} as the dummies indicating if i would meet the conditions in the relevant scenario. The scenarios also define the transfers, with $t_h^{nt} = 0$, $t_h^{uct} = \sum_{i \in \mathcal{I}_h} t_i^{uct}$, and $t_h^{cct} = \sum_{i \in \mathcal{I}_h} a_i^{cct} t_i^{cct}$, where t_i^{cct} is the transfer child i would receive if meeting the conditions of the CCT, divided by household size. t_i^{uct} is the UCT transfer, which is rescaled from the CCT transfer to keep comparisons budget neutral (see Section 3). As previously discussed, the framework is adapted for Nicaragua to account for the household-level conditionality (all children must meet conditions).

Children types. We assume that $a_i^{nt} = a_i^{uct}$ for all i. This is a key assumption discussed in Section 3. It implies that the UCT has no effect on human capital investment relative to no transfers. We make this assumption for tractability (as discussed below), but note that it works "against" a UCT and "in favor" of a CCT in our welfare ratio calculations. [6]

Given this, there are three relevant types of children:

1. Marginal: $a_i^{nt} = a_i^{uct} = 0$, but $a_i^{cct} = 1$.

2. Inframarginal conditions-unmet: $a_i^{nt} = a_i^{uct} = a_i^{cct} = 0$.

3. Inframarginal conditions-met: $a_i^{nt} = a_i^{uct} = a_i^{cct} = 1$.

Mapping c_h^* to $c_h^{baseline}$. Let m_i be a dummy indicator that is equal to one if child i is marginal. We assume that households' a_i decisions under the baseline are the same as under the no-transfer scenarios (since indeed there are no transfers in the baseline). Thus, under a CCT, $c_h^* - \sum_{i \in \mathcal{I}_h} a_i^{cct} f_i = c_h^{baseline} - \sum_{i \in \mathcal{I}_h} m_i f_i$. In the baseline period, there are no transfers, and thus it is already incurring f_i for the inframarginal conditions-met, and

[6] Some evidence suggests that UCTs, relative to a pure control, increase school enrollment. See Baird et al. (2011) for the Malawi case and Benhassine et al. (2015) for Morocco (with the caveat that the latter studies a "labeled cash transfer" which is similar to a UCT in that the conditionality is not enforced). Our framework allows for the possibility of a UCT increasing enrollment via an income effect. As a household's consumption rises with the UCT payment, their marginal utility becomes lower, and so does the utility cost of incurring f_i. The household compares this cost to D_i, which does not depend on income. Thus, the assumption implicitly involves assuming that the distributions of f_i and D_i are such that the income effect of a UCT is not large enough to create any child with $a_i^{nt} \neq a_i^{uct}$.

$c_h^{baseline}$ already reflects that. Under a UCT and no transfers, $c_h^* - \sum_{i \in \mathcal{I}_h} a_i^{uct} f_i = c_h^* - \sum_{i \in \mathcal{I}_h} a_i^{nt} f_i = c_h^{baseline}$ given a similar logic.

CCT-to-UCT welfare ratios incorporating human capital gains. The CCT-to-UCT ratio, including gains from human capital investments is given by:

$$R^{HC} = \frac{\sum_{h=1}^{H} size_h \cdot [u(c_h^{baseline} + t_h^{CCT} - \sum_{i \in \mathcal{I}_j} m_i f_i) + \sum_{i \in \mathcal{I}_j} m_i D_i - u(c_h^{baseline})]}{\sum_{h=1}^{H} size_h \cdot [u(c_h^{baseline} + t_h^{UCT}) - u(c_h^{baseline})]} \tag{6}$$

Equation (6) highlights a key component of our revealed-preference approach: only the human capital gains for *marginal* children matter for a CCT versus UCT welfare comparison. The denominator of equation (6) is the same as in the baseline CCT-to-UCT welfare ratio (R from equation 1). The numerator is only different due to the households with marginal children ($m_i = 1$). This follows from our assumptions on children's types discussed above. Intuitively, the gains and costs from human capital investments by inframarginals are "differenced out" in both the numerator and denominator of the formula.

Computing R^{HC} requires knowledge of m_i, f_i, and D_i, which are not directly observable. All other terms are directly observable in our data (baseline consumption, household size, and transfer amounts). Indeed, the discussion that follows below focuses on how we measure m_i, f_i, and D_i.

If there are no marginal children ($m_i = 0$ for all i), then $R^{HC} = R$: equation (6) simplifies to equation (1), so the ratio with human capital R^{HC} equals the baseline R. Thus, the baseline R can be interpreted as the ratio observed if no children are induced by the CCT (compared to a UCT) to invest in human capital. The same holds if $D_i = 0$ and $f_i = 0$ for all i (no gains or costs from meeting conditions).

Estimating the share of marginals. While we do not observe *which* children are marginal, we can exploit that the *share* of marginal children (S^{marg}) equals the treatment effect of a CCT, compared to a pure control. Note the assumption that UCTs, relative to no controls, do not affect a_i plays a role here.

For each of our contexts, we locate in the source papers their "headline" effect of a CCT versus a pure control on enrollment. All source papers use random assignment to estimate such effects. For Malawi, we set $S^{marg} = .058$, based on Table III, column (8) of Baird et al. (2011).

For Mexico, we use the results from Table 6, row (3) from Parker and Todd (2017), which are based on Skoufias et al. (2013). For boys ages 8–11 (grades 3-6) the enrollment impact is 1.3–1.8 percentage points; for girls 8–11 there are no significant impacts. For boys 12–17 (grades 7-9) they report 3.2–5.8 percentage point impacts, while girls 12–17 have 7.5–9.5 percentage point impacts. For the entire Mexico sample we use .0323 which equals the mid-points of the within gender results, averaged across the genders, and then averaged across the grades using the number of grades in each group: $.0323 = \frac{1}{7}[(.0155 + 0)/.02] * 4 + [(.045 + .085)/2] * 3]$. For the Mexico sub-sample of 13-18 year olds, we use the average of the effects across the genders for the 7th-9th grade groups (which is the closest easily available treatment effects for older ages): $.065 = (.045 + .085)/2$.

For Morocco, we set $S^{marg} = 0.054$, based on Table 5, panel A, row 1 of Benhassine et al. (2015). We take the average of the "CCT to Mothers" and "CCT to Fathers" effects. The paper does not separately report effects for the ages 11-15 group. We calculate an effect for this subsample by assuming the full-sample effect is entirely driven by those aged 11–15 (i.e., zero effect for younger children) and rescaling based on their share of the sample. We prefer this approach, which "ties our hands," to estimating effects not reported in the source papers ourselves. This leads to a $S^{marg} = .112$ for the Morocco ages 11-15 subsample.

For Nicaragua, we set $S^{marg} = 0.128$ based on Table 4.8 of Maluccio and Flores (2005). For Tanzania, we set $S^{marg} = 0.040$ using Table 7.12, column 3, row 2 in Evans et al. (2014). The paper reports this as an effect for ages 0–18, which may include children too young to be eligible (under 5). Evans et al. (2014) also reports an effect of 0.030 for ages 7–14, so $S^{marg} = 0.040$ appears to be a reasonable approximation. We calculate the effect for the Tanzania subsample (ages 13–15) by rescaling the effect as in the Moroccan case, obtaining a $S^{marg} = .134$

Assigning marginals (m_i). While we know the *share* of marginal children (S^{marg}), we do not know which specific children are marginal. By definition, only children who meet conditions under the CCT can be marginals.

We thus take the Monte Carlo approach described in the main text: randomly assigning $S^{marg} \cdot N$ of the conditions-met children to be assigned as marginal. The draws are independent across children, including those of the same households (e.g., two or more children in the same household may be selected randomly to be marginals in the same household). We take 1,000 draws for each context, recalculate the ratio R^{HC} for each draw, and our tables report the 1st and 99th percentiles of this R^{HC} distribution.

Under the revealed-preference approach to calculating D_i, we also ensure all analyzed draws of marginal children satisfy a revealed-preference criterion. We revisit this issue below when discussing that approach.

The cost of meeting conditions (f_i). We set f_i equal to 15% of i's baseline consumption ($f_i = 0.15 \cdot c_h^{baseline}$) in both our revealed preference and ROI approaches. We base this on the figures provided for Morocco by Benhassine et al. (2015), the only source paper providing figures on foregone child labor. It reports that CCT transfers equal "*a fifth of mean*" child labor earnings. Since transfers are 3% of household consumption (Table 1), we derive $0.15 = 0.03 \cdot 5$.

Benhassine et al. (2015) note their child labor data is from a small sample, and many conditions-unmet children do not work. However, this may not be directly relevant, as we are interested in determining f_i for the marginals, whom we only observe as conditions-met under a CCT and do not know if they would have worked under a counterfactual scenario. It is also plausible that the marginal children are the ones more likely to work if not enrolled (as they are the ones affected by the conditionality of the cash transfer).[7]

[7] Another possibility would be to set $f_i = x$, where x is a constant in dollars that is the same for all marginal children, and not a fraction of i's baseline consumption. However, setting $f_i = x$ runs into the issue that, if x is set to be substantial (say, 15% of average household consumption), it may be larger than the poorest households' baseline consumption, implying they have negative consumption if they have marginal children, and making it impossible to compute their value of utility.

A revealed preference bound on human capital gains (D_i). If child i is marginal, then, by definition, its household under a UCT obtains higher utility from choosing $a_i = 0$ than $a_i = 1$ (i.e., not meeting the conditions provides higher utility). Therefore, the following holds if i is marginal:

$$u(c_h^* + t_h^{uct} - \sum_{j \in \mathcal{I}_h^{-i}} a_j f_j - f_i) + \sum_{j \in \mathcal{I}_h^{-i}} a_j D_j + D_i \leq u(c_h^* + t_h^{uct} - \sum_{j \in \mathcal{I}_h^{-i}} a_j f_j) + \sum_{j \in \mathcal{I}_h^{-i}} a_j D_j$$

which can be rearranged to

$$D_i \leq u(c_h^* + t_h^{uct} - \sum_{j \in \mathcal{I}_h^{-i}} a_j f_j) - u(c_h^* + t_h^{uct} - \sum_{j \in \mathcal{I}_h^{-i}} a_j f_j - f_i) \equiv D_i^{RP}$$

which is equation (3) in the main text. The equation above provides D_i^{RP}, which is our revealed-preference upper bound on D_i. Importantly, D_i^{RP} is directly observable in the data under our previously discussed assumptions for f_i and the mapping between c_h^* and $c_h^{baseline}$. By substituting D_i^{RP} in the place of D_i in equation (6), we can estimate R^{HC} for a given draw of marginals (the m_i).

Revealed-preference approach: discussion. The key component of the approach is that, by definition, household decision-makers choose not to meet conditions under a UCT.

Intuitively, the CCT, compared to a UCT, is unlikely to change the behavior of households whose children have large benefits and low costs of human capital investments: these children would be enrolled under the UCT as well. Similarly, if marginal children have large benefits, they also must face high costs. This is also the intuition for why D_i^{RP} is increasing in f_i (which can be seen by inspecting equation 3). Moreover, assuming $f_i = 0$ and $D_i^{RP} > 0$ is inconsistent with equation (3). This highlights that both the costs and benefits of human capital investment must be incorporated in a framework that relies on revealed preference.

The "mechanics" behind equation (3) is that the concavity of u places a bound on how large the utility gain D_i can be relative to the cost f_i. Moreover, the upper bound D_i^{RP} on D_i, for a given f_i, is obtained when the household with marginal children i is indifferent between losing f_i of consumption and receiving D_i of utility under a UCT. However, under a CCT, the utility loss from incurring f_i is lower than under a UCT (because CCT payments are larger than UCT payments and u is concave). However, the benefit D_i is the same under both programs. To see the role of concavity, note that, if utility is linear and $u(x) = x$, equation (3) implies $f_i = D_i^{RP}$ and R^{HC} (equation 6) simplifies to become equal to R (equation 1).

As noted above, D_i^{RP} can be interpreted broadly to include effects beyond school enrollment, such as attendance, test scores, cognitive improvements, and health gains. A caveat is that, since we use S^{marg} based on enrollment effects, the method implicitly assumes all such gains accrue only to children marginal in the enrollment decision.

At first pass, this approach may appear at odds with the common intuition that "if households make optimal decisions, imposing conditionality can only lower their welfare." However, such intuition implicitly assumes that imposing conditionality does not affect transfer sizes. In our comparisons, which are "budget neutral," removing conditionality (switching from a CCT to a UCT) lowers the size of the transfers to allow conditions-unmet children to receive them. In other words, breaking budget neutrality and assuming $t_i^{cct} = t_i^{uct}$ would make the CCT gains from equation (6) be the same as the numerator of our baseline welfare ratio (equation 1). This highlights a key aspect of our counterfactuals, which is that conditionality has three consequences: i) inducing human capital investments, ii) determining who receives the transfers, and iii) the size of the transfers, given a fixed budget and the number meeting conditions. The "common intuition" abstracts away from the latter consequence.

Revealed-preference approach: assigning marginals. In our random assignment procedure, a household with multiple conditions-met children may be assigned more than one marginal child. If assigned only one, equation (6) implies the household has higher utility under the CCT than under no transfers. This occurs since D_i^{RP} is always larger than the utility loss from losing f_i of consumption under CCT consumption levels, as discussed above. However, if a household is assigned two or more marginal children, its utility under the CCT can *potentially* be lower than under no transfers. This is due to the concavity of u. For example, the utility loss of a $2f_i$ reduction in consumption is greater than twice the loss from a f_i reduction, and thus may exceed $2D_i^{RP}$.

However, a household making optimal choices cannot have lower utility under a CCT relative to no transfers in our framework (i.e., it can choose to take the same actions under both scenarios, which cannot provide lower utility under the CCT). To account for this in our assignment of marginals, when obtaining the 1,000 draws, we include only draws where *every* household in the sample has utility under a CCT that is equal to or higher than under no transfers (i.e., its contribution to the numerator of equation (6) is nonnegative). We implement this by discarding draws that do not meet this criterion and continuing draws until we obtain 1,000 valid ones.[8]

Revealed-preference approach: sensitivity to choice of f_i. Given observed consumption, marginal assignment, and transfer size, D_i^{RP} and R^{HC} under the revealed-preference approach depend solely on f_i, making sensitivity to it worth exploring.

Table A.4 reports revealed-preference R^{HC} results using $f_i = 0.30 \cdot c_h^{baseline}$, doubling the value used in Table 2. Baseline results abstracting from human capital gains (R in Table 1) correspond to $f_i = 0$. Together, Tables 1, 2, and A.4 span a range of f_i values. Within a context, differences in welfare ratios are modest and rarely alter qualitative conclusions: ratios well below or near one tend to remain so under alternative f_i assumptions.

The ratio R^{HC} increases with f_i in most contexts. However, this is not generally true. For a case where each household h has no more than one marginal child, R^{HC} is increasing in f_i. This occurs since D_i^{RP} is larger than the utility loss from consuming f_i less when evaluated at the $c_h^{baseline} + t_h^{CCT}$ level. This holds since u is concave, $t_h^{CCT} > t_h^{UCT}$, and D_i^{RP} is (by definition) the absolute value of the utility loss from consuming f_i less when evaluated at the $c_h^{baseline} + t_h^{UCT}$ level. This implies that a household h contribution to the numerator of R^{HC} is larger when it is assigned to have exactly one marginal child than

[8] For the Nicaragua and Morocco samples, this approach is computationally cumbersome because most marginal assignment draws include at least one household with lower utility under the CCT than under no transfers. This arises because transfers are small relative to average consumption and many households have two or more eligible children. In these cases, we instead use a random assignment algorithm that limits each household to one marginal. This likely has little effect on results: in our inspection of realized draws, nearly all cases with multiple marginals led to lower utility under the CCT than under no transfers.

when it is assigned to have zero marginal children.

However, since u is concave, this logic may not hold. The utility loss from consuming $k \cdot f_i$ less is larger than k times the utility loss from consuming f_i less. Thus, being assigned k equals two or more marginal children may lower or increase household h's contribution, depending on baseline consumption, transfer sizes, and f_i.

For example, in the Tanzania subsample, the 1st-percentile R^{HC} in Table 2 is lower than the case with $f_i = 0$ (Table 1) due to randomization draws where multiple marginals are assigned to the same household.

A return on investment (ROI) approach to measuring human capital gains (D_i). Under this approach, we make explicit assumptions on rates of return to human capital and discount rates to calculate D_i. In particular:

$$D_i^{ROI} = \sum_{t=0}^{T} \beta^t [\frac{c(1)_{it}^{1-\rho}}{1-\rho} - \frac{c(0)_{it}^{1-\rho}}{1-\rho}]$$

which is equation (4) in the main text. It involves the child living for T more years and a discount factor of β. $c(1)_{it}$ and $c(0)_{it}$ are, respectively, the trajectories of consumption if the child meets the conditions or not. Following Bergstrom and Dodds (2021), we set $c(0)_{it}$ to be equal to $c_h^{baseline}$ in all T, and $c(1)_{it}$ equal to $(1 + \phi) \cdot c_h^{baseline}$, where ϕ is the returns to schooling.

Under an assumption on the value of β and ϕ, D_i^{ROI} is directly observable in the data. By substituting D_i^{ROI} in the place of D_i in equation (6), we can estimate R^{HC} for a given draw of marginals (the m_i).

Return on investment (ROI) approach: discussion. Our formulation follows Bergstrom and Dodds (2021), the only study we are aware of that provides an ROI calculation based on *utility from consumption*, and thus consistent with our framework. For comparison, using Progresa data, Behrman et al. (2009) calculates ROI in monetary terms (i.e., program benefits as the net present value of future wage gains, measured in dollars).[9]

The Bergstrom and Dodds (2021) formulation also assumes consumption is constant in all future periods. This implies perfect consumption smoothing over time. While full smoothing is unlikely for relatively poor households in developing countries, this assumption avoids introducing additional assumptions about consumption dynamics.

Our formulation focuses on consumption and its utility, not wages or income, which has two implications for interpreting equation (4). First, ϕ may differ from returns to schooling estimated via a Mincerian wage regression, since i may not work every period (so schooling's effect on future consumption can differ from its effect on wages conditional on working). Second, assuming a child who fails to meet conditions has the same per capita consumption as their current household, $c_h^{baseline}$, does not imply lifelong child-labor earnings. It assumes no overall household consumption growth—an unlikely scenario, but one that avoids further parameter choices.

Moreover, $c_h^{baseline}$ is per capita *household* consumption, which may reflect income from multiple members. Thus, ϕ captures the effect of schooling on child i's future household consumption. This effect may differ from its impact on i's own income due to factors such as dilution (schooling affects only i among several earners), assortative mating (schooling leads to marrying a higher human capital spouse), and fertility (schooling influences household size).

Fully accounting for such factors is beyond the scope of this paper, as it would require additional assumptions or data not available in the experiments we analyze. We note here that, in our formulation, following Bergstrom and Dodds (2021), ϕ matches the effect on i's wages (e.g., from a Mincerian regression) under a scenario where the household remains constant in size, child i becomes the sole earner, and works every period until T.

Choice of parameters in the ROI approach. Beyond the parameter governing the concavity of the utility function (ρ), which we set equal to three in all baseline estimations,

[9] These are not directly comparable, as discounted utility flows cannot be directly inferred from the utility of net present monetary flows. Moreover, we are not aware of other papers applying an ROI approach in a CCT context. Most of the literature, surveyed in Banerjee et al. (2024), focuses on effectiveness metrics such as "school enrollment per dollar" or "additional years of schooling per dollar," which do not require valuing the future benefits of schooling, not to mention assigning utility values to them.

equation (4) requires setting three other parameters. We set $T = 50$, following Bergstrom and Dodds (2021) and Behrman et al. (2009).

We set $\beta = 0.90$, the lowest discount factor considered by Behrman et al. (2009) in their analysis of Progresa data. We set $\phi = 0.03$, half the lowest return to schooling in Behrman et al. (2009). We halve it because their estimate reflects the effect of an additional year of schooling on wages conditional on working (e.g., from a Mincerian regression), whereas ours concerns consumption, which is likely smaller since not all individuals work continuously and household consumption depends on other members' income.[10]

Higher values of β and ϕ increase D_i^{ROI} and make the CCT-to-UCT welfare ratio larger (more "pro-CCT"). As discussed in Section 3, ratios are near or above one in all contexts under such assumptions (except in Malawi). Thus, even with relatively "pro-UCT" assumptions on discount factors and returns to schooling, CCTs can yield higher welfare than UCTs. A full quantification of how relative welfare changes under alternative assumptions is beyond the scope of this paper, especially given the difficulty of pinning down the "correct" discount factor or schooling effect on consumption in each context.[11]

Externalities. Beyond households underestimating returns or over-discounting future consumption (which the ROI approach can account for), another possible rationale for CCTs is externalities: if social returns to human capital investments exceed private returns, conditionality can improve welfare.

We address externalities using a formulation adapted from the one used by Banerjee et al. (2024) to conceptually discuss social returns to human capital. So far, we assumed household h's utility follows equation (2) and evaluated welfare as the sum of these utilities across all h.

To incorporate externalities, we let the utility functions that underlie the social welfare function differ from households' utility functions. In particular, we assume a social

[10] Behrman et al. (2009) only includes men in their analysis, whereas all our contexts include women too.
[11] Consistent with this discussion, Bergstrom and Dodds (2021) assumes a $\beta = 0.95$ and $\phi = 0.08$.

planner whose welfare from household h is given by:

$$u(c_h^* + t_h - \sum_{i \in \mathcal{I}_h} a_i f_i) + \theta \sum_{i \in \mathcal{I}_h} a_i D_i$$

which is equation (5) from the main text. Equation (5) differs from (2) only by the parameter θ. If $\theta = 1$, they coincide. Thus, θ measures how much more the social planner values human capital than households. When $\theta > 1$, social returns exceed private returns, implying positive externalities.

We adopt this stylized approach because a fully microfounded treatment of externalities is beyond this paper's scope. It would require numerous assumptions about hard-to-observe parameters and data that are not available (e.g., which agents beyond h benefit from h's children's human capital, by how much, and whether they are transfer-eligible). To our knowledge, no study quantifies welfare gains from externalities in CCTs.

A CCT-to-UCT welfare ratio that incorporates externalities according to equation (5) would thus be:

$$R^{EXT} = \frac{\sum_{h=1}^{H} size_h \cdot [u(c_h^{baseline} + t_h^{CCT} - \sum_{i \in \mathcal{I}_j} m_i f_i) + \theta \sum_{i \in \mathcal{I}_j} m_i D_i - u(c_h^{baseline})]}{\sum_{h=1}^{H} size_h \cdot [u(c_h^{baseline} + t_h^{UCT}) - u(c_h^{baseline})]} \tag{7}$$

R^{EXT} differs from the ratio incorporating human capital gains (R^{HC} from equation 6) only by including the parameter θ. Moreover, comparing equations equations (6) and (7), we obtain:

$$R^{EXT} = R^{HC} + (\theta - 1)GHC \tag{8}$$

where GHC is the sum of **g**ain from **h**uman **c**apital (the D_i) for marginal children, divided by the denominator of the welfare ratios (the welfare gains from a UCT):

$$GHC = \frac{\sum_{h=1}^{H} size_h \sum_{i \in \mathcal{I}_j} m_i D_i}{\sum_{h=1}^{H} size_h \cdot [u(c_h^{baseline} + t_h^{UCT}) - u(c_h^{baseline})]} \tag{9}$$

Importantly, GHC can be calculated in our data under the revealed preference or ROI approaches (i.e., substituting D_i^{RP} or D_i^{ROI} in the place of D_i in equation 9).

For an observed R^{HC}, GHC, and value of θ, one can determine a R^{EXT} ratio that includes externalities. For example, in the case of the revealed-preference approach for Mexican full sample, $R^{HC} = 0.954$ and $GHC = 0.044$.[12] Thus, a social welfare function that values human capital gains three times more than households do ($\theta = 3$) would have a $R^{EXT} = 0.954 + (3-1) \cdot 0.044 = 1.04$. It would also allow one to calculate the value of θ that would make the social planner indifferent between a CCT and UCT: $1 + \frac{1-R^{HC}}{GHC}$, which in this example is $1 + \frac{1-0.954}{0.044} = 2.05$

We remain agnostic about the appropriate value of θ, as the literature offers little guidance on the quantitative magnitude of social returns to human capital—whether induced by CCTs or more generally. Instead, we provide GHC across contexts for both our revealed-preference and ROI approaches, which allows the reader to calculate R^{EXT} for any particular of θ.

For completeness, we also report a measure of the aggregate costs of human capital (CHC) that is analogous to the gains from human capital (GHC). Algebraic manipulation of equations (1), (6), and (9) yields $R^{HC} = R + GHC + CHC$, where R is the baseline ratio (equation 1) and

(10)

Thus, CHC is the sum of **c**osts of **h**uman **c**apital (loss utility from f_i) from marginal children, divided by the denominator of the welfare ratios (the welfare gains from a UCT).[13]

Tables A.5 and A.6 report the CHC and GHC across contexts, for both the revealed preference and ROI approaches. It also provides the baseline ratio R and the R^{HC} ratios for context. In some cases, GHC and CHC are large in absolute value—positive and negative, respectively—offsetting each other in the gap between R and R^{HC}. This typically occurs when the share of marginals is high and/or transfers are small: the numerators in GHC and CHC grow with marginal share, while smaller transfers reduce the denominator.

[12] We use the 99th-percentile draws for this example.

[13] CHC could also be used in combination with GHC in an additional formulation accounting for externalities in the cost side (i.e., the private and social costs of f_i differ).

Table A.1: Programs Considered for Analysis

Country	Program Name	Reason for Non-Inclusion	Source Paper
Included in our analysis:			
Malawi	SIHR	-	Baird et al. (2011)
Mexico	Progresa	-	Parker and Todd (2017)
Morocco	Tayssir	-	Benhassine et al. (2015)
Nicaragua	Red de Protección Social (RPS)	-	Maluccio and Flores (2005)
Tanzania	TASAF	-	Evans et al. (2014)
Not included in our analysis:			
Burkina Faso	NCTPP	Data not available	Akresh et al. (2020)
Indonesia	Program Keluarga Harapan (PKH)	No baseline eligibility data	World Bank Office Jakarta (2011)
Nicaragua	Atención a Crisis	Enrollment data not available	Macours et al. (2012)
Honduras	Programa de Asignación Familiar II (PRAF II)	No consumption data	Galiani and McEwan (2013)
Philippines	Pantawid Pamilyang Pilipino Program (PPPP)	No consumption data	Chaudhury et al. (2013)
Colombia	SCAE	No consumption data	Barrera-Osorio et al. (2011)
Uganda	Cash transfer for pre-school	Unconditional	Gilligan and Roy (2013)
Kenya	GiveDirectly	Unconditional	Haushofer and Shapiro (2013)
Mexico	Programa de Apoyo Alimentario (PAL)	Unconditional	Skoufias et al. (2013)
Ecuador	BDH	Unconditional	Edmonds and Schady (2012)
Kenya	CT-OVC	Unconditional	Asfaw et al. (2014)
Malawi	SCT program	Unconditional	Covarrubias et al. (2012)
Zambia	Child Program	Unconditional	American Institutes for Research (2013)
Zimbabwe	Manicaland CCT/UCT Experiment	Data not available	Robertson et al. (2013)

The papers listed above are the 18 trials considered by Banerjee et al. (2017) study of cash transfers' effects on labor supply with the addition of Robertson et al. (2013), as it is one of the experiments we are aware of that includes both CCT and UCT treatment arms. See Section 2 and Appendix A for further discussion.

Table A.2: Share of Marginals and Conditions-Unmet Children Across Programs

Country	Age Range (1)	Share Marginal (2)	Share Cond.-unmet (3)
Panel A: CCT programs included in our analysis			
Malawi	14–20	0.058	0.368
Mexico	7–18	0.032	0.197
Morocco	6–15	0.053	0.208
Nicaragua	7–13	0.128	0.090
Tanzania	7–15	0.040	0.156
Included avg.:		*0.062*	*0.195*
Panel B: CCT programs not included in our analysis			
Burkina Faso	7–15	0.180	0.320
Honduras	6–12	0.080	0.260
Philippines	6–11	0.045	0.025
Colombia	Grades 6–10	0.017	0.285
Zimbabwe	6–12	0.076	0.144
Not-included avg.:		*0.080*	*0.207*

Panel A includes the five contexts analyzed in the paper. Panel B includes programs considered but not analyzed in the paper, except those that (1) only evaluate unconditional transfers against a pure control (see Table A.1 for further information), or (2) conditions were stated but not enforced (Indonesia PKH), or (3) where data on the share of marginals is not available in the source paper (Nicaragua *Atención a Crisis*). Column (1) provide the age (or grade) range of eligible participants. Column (2) is the share of children who are marginal (induced by the CCT to meet conditions). Column (3) provides the share of conditions-unmet children. Numbers are obtained from source papers listed in Table A.1. See main text for further information and definition of how marginals and conditions-unmet are measured.

Table A.3: CCT-to-UCT Welfare Ratios: Sensitivity to Choice of ρ

Sample	CCT-to-UCT Welfare Ratio (R)			
	Ln Utility	CRRA ($\rho = 2$)	CRRA ($\rho = 3$)	CRRA ($\rho = 4$)
	(1)	(2)	(3)	(4)
Panel A: Full Samples				
Malawi	0.878	0.762	0.675	0.629
Mexico	0.978	0.952	0.932	0.918
Morocco	0.990	0.978	0.961	0.938
Nicaragua	0.995	1.004	1.002	1.001
Tanzania	0.989	0.987	0.958	0.946
Panel B: Subsamples				
Mexico (ages 13-18)	0.918	0.875	0.837	0.804
Morocco (ages 11-15)	0.976	0.944	0.898	0.835
Tanzania (ages 13-15)	0.982	0.954	0.912	0.861

Table presents the ratio between the welfare gains from the observed CCT and a counterfactual UCT program that uses the same budget to provide payments to all eligible children (R from equation 1). All ratios are based on assuming a CRRA utility function with a coefficient of risk aversion ρ. Column (1) assumes log utility (equivalent of $\rho = 1$), and columns (2), (3), and (4) assume ρ equal to two, three, or four, respectively. Column (3) replicates the baseline result in column (3) of Table 1, and is replicated here for reference.

Table A.4: Revealed Preference Welfare Ratios with $f_i = 0.3 \cdot c_h^{baseline}$

		CCT-to-UCT Welfare Ratio (R)		
			Revealed Preference	
	S^{marg}	Baseline	1st pctile	99th pctile
	(1)	(2)	(3)	(4)
Panel A: Full Samples				
Malawi	0.058	0.675	0.681	0.711
Mexico	0.032	0.932	0.974	0.990
Morocco	0.054	0.961	1.090	1.162
Nicaragua	0.128	1.002	1.021	1.042
Tanzania	0.04	0.958	0.988	1.098
Panel B: Subsamples				
Mexico (ages 13-18)	0.065	0.837	0.890	0.911
Morocco (ages 11-15)	0.112	0.898	1.070	1.153
Tanzania (ages 13-15)	0.134	0.912	0.964	1.117

Column (1) presents the share of marginal children (S^{marg}): those that meet conditions (e.g., enroll in school) under the CCT but not under a UCT. For reference, column (2) replicates the baseline CCT-to-UCT welfare ratio (R) from column (3) in Table 1. It provides R assuming no gains from human capital. Columns (3)-(4) present R using the revealed-preference approach. Column (3) provides the 1st percentile of the distribution of simulated R based on randomly assigning which conditions-met children are marginal, while column (4) does so for the 99th percentile. They differ from columns (3) and (4) in Table 2 only in that they double the value of f_i so it equals 30% of baseline consumption. Columns (2)–(4) are based on a CRRA utility with $\rho = 3$.

Table A.5: Gains and Costs from Human Capital - Revealed-Preference Approach

	Baseline	1st percentile			99th percentile		
	R	GHC	CHC	R^{HC}	GHC	CHC	R^{HC}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Full Samples							
Malawi	0.675	0.010	-0.008	0.677	0.038	-0.024	0.689
Mexico	0.932	0.034	-0.019	0.947	0.044	-0.022	0.954
Morocco	0.961	0.297	-0.255	1.003	0.404	-0.335	1.03
Nicaragua	1.002	0.214	-0.207	1.009	0.292	-0.278	1.016
Tanzania	0.958	0.074	-0.063	0.969	0.217	-0.168	1.007
Panel B: Subsamples							
Mexico (ages 13-18)	0.837	0.054	-0.034	0.857	0.069	-0.041	0.865
Morocco (ages 11-15)	0.898	0.66	-0.603	0.955	0.877	-0.79	0.985
Tanzania (ages 13-15)	0.912	0.271	-0.31	0.873	0.901	-0.831	0.982

For reference, column (1) presents R, the baseline CCT-to-UCT welfare ratio from Table 1, while columns (4) and (7) present R^{HC}, the welfare ratio incorporating gains from human capital from Table 2. Columns (2) and (5) present GHC: the sum of **g**ains from **h**uman **c**apital (the D_i) for marginal children, divided by the denominator of the welfare ratios (equation 9). Columns (3) and (6) present CHC: the sum of **c**osts of **h**uman **c**apital (loss utility from f_i) from marginal children, divided by the denominator of the welfare ratios (equation 10). Results are presented separately for the 1st percentile and the 99th percentile of the distribution of simulated ratios based on randomly assigning which conditions-met children are marginal. All columns are based on a CRRA utility with $\rho = 3$, and columns (2)-(7) use the revealed-preference approach to incorporate human capital gains. For a given percentile, $R + GHC + CHC = R^{HC}$. See discussion in Section 3 and Appendix C.

Table A.6: Gains and Costs from Human Capital Investments - ROI Approach

	Baseline	1st percentile			99th percentile		
	R	GHC	CHC	R^{HC}	GHC	CHC	R^{HC}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Full Samples							
Malawi	0.675	0.027	-0.009	0.693	0.172	-0.024	0.823
Mexico	0.932	0.083	-0.018	0.997	0.107	-0.022	1.017
Morocco	0.961	0.48	-0.261	1.18	0.679	-0.355	1.285
Nicaragua	1.002	3.965	-3.776	1.191	6.158	-5.796	1.364
Tanzania	0.958	0.119	-0.053	1.024	0.402	-0.179	1.181
Panel B: Subsamples							
Mexico (ages 13-18)	0.837	0.135	-0.034	0.938	0.175	-0.039	0.973
Morocco (ages 11-15)	0.898	1.072	-0.634	1.348	1.433	-0.819	1.522
Tanzania (ages 13-15)	0.912	0.466	-0.256	1.122	1.569	-0.922	1.559

For reference, column (1) presents R, the baseline CCT-to-UCT welfare ratio from Table 1, while columns (4) and (7) present R^{HC}, the welfare ratio incorporating gains from human capital from Table 2. Columns (2) and (5) present GHC: the sum of **g**ains from **h**uman **c**apital (the D_i) for marginal children, divided by the denominator of the welfare ratios (equation 9). Columns (3) and (6) present CHC: the sum of **c**osts of **h**uman **c**apital (loss utility from f_i) from marginal children, divided by the denominator of the welfare ratios (equation 10). Results are presented separately for the 1st percentile and the 99th percentile of the distribution of simulated ratios based on randomly assigning which conditions-met children are marginal. All columns are based on a CRRA utility with $\rho = 3$, and columns (2)-(7) use the return on investment (ROI) approach to incorporate human capital gains. For a given percentile, $R + GHC + CHC = R^{HC}$. See discussion in Section 3 and Appendix C.