Note from the authors (in case you miss the comment we added to our submission): this paper is currently conditionally accepted at the American Economic Review: Insights. The editor requested we resubmit by early October, so the paper may be "finally accepted" at the time of the meeting. We also believe this would be a good fit for a shorter presentation.

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

Santosh Anagol† Thomas Fujiwara‡ Martin Navarrete§

May 2025

Abstract

Conditionality can prevent poor households from receiving cash transfers. Re-analyzing five randomized evaluations of conditional cash transfers, we find: (1) households not meeting conditions are common, representing 4.6% to 37% of eligible households; (2) these households often have lower baseline consumption than households that meet conditions; (3) under standard social welfare function assumptions, a budget-neutral switch to unconditional cash transfers can raise the welfare gains from increased consumption by up to 46%, although we find important heterogeneity by context. Our results suggest conditionality exacerbating consumption inequality can be quantitatively important for the welfare impacts of cash transfer programs.

*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

More than 120 low- and middle-income countries run cash transfer programs for poor families, reaching hundreds of millions, if not billions, of households (Banerjee et al., 2024). A crucial design choice for these programs is whether to condition transfers on recipient behavior. Conditional cash transfers (CCTs), which are "*perhaps the single largest innovation in social protection programs in the developing world in the past 30 years*" (Banerjee et al., 2024), usually have two requirements: recipient households must i) have income below a threshold, and ii) meet conditions such as enrolling children in school.

If excluded households are poorer than those meeting the conditions, the transfer program effectively targets richer households and increases inequality among those eligible. Under standard social welfare functions that value progressive redistribution, this diminishes the program's welfare gains. To our knowledge, no existing work quantifies the welfare consequences of conditionality that exacerbates consumption inequality; this is the focus of our paper.[1]

Theoretically, it is unclear whether poorer or richer households are more likely to meet conditions. While poorer households derive higher marginal utility from transfers, they may also face higher opportunity costs, be more reliant on child labor, or struggle more with administrative hurdles. They may also, correctly or incorrectly, perceive lower returns to schooling for their children.

We analyze data from five randomized evaluations of CCTs in Malawi (SIHR), Mexico (Progresa), Morocco (Tayssir), Nicaragua (RPS), and Tanzania (TASAF). We present five results for each context. First, we report the fraction of condition-met and condition-unmet households by child age and/or grade. Condition-met households are eligible households that do not meet their program's most basic educational condition: school enrollment. Among eligible households assigned to a CCT treatment, the shares of condition-unmet households are 36.8% in Malawi, 19.7% in Mexico, 20.8% in Morocco, 15.6% in Tanzania, and 4.6% in Nicaragua. These shares are higher in certain groups. For example, children in Mexico who completed only primary school (6th grade) have a 56.4% share of

[1] A large literature estimates the effects of CCTs on education, health, labor supply, and other outcomes. See Banerjee et al. (2024) for a survey.

condition-unmet households. Overall, these sizable rates of households not meeting conditions suggest that conditionality may have substantial welfare consequences through our inequality mechanism.[2]

Second, we compare the baseline (before transfers start) consumption distributions for condition-met and condition-unmet households. Although straightforward, this comparison has not been systematically reviewed and is not usually reported in individual evaluations of CCTs. In Malawi, condition-unmet households have, on average, 26% lower consumption. In Mexico, Tanzania, and Morocco, condition-unmet households are also poorer than condition-met households, but differences in averages are more modest (between 4%-5%). In Nicaragua, the average consumption of condition-unmet households is 16% higher, but this number should be interpreted with caution given the small sample (33). In no context do we find clear evidence of condition-unmet households having substantially larger consumption. Overall, conditionality may lead to ex-ante poorer households not receiving transfers, but magnitudes depend on the context.

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. We follow a social welfare formulation from Alatas et al. (2019) and Hanna and Olken (2018), using a Constant Relative Risk Aversion (CRRA) utility function with a coefficient of relative risk aversion $\rho = 3$ for our baseline calculations.[3] We find welfare consequences are substantial in some contexts but more muted in others. In Malawi, we estimate the CCT delivers only 68.4% of the welfare gains that an equally budgeted UCT program would deliver. In other words, it suggests a change to a UCT would increase welfare by 46%. In Mexico, the CCT delivers 71.3% of the benefit of a UCT when focusing on children who are likely to be non-compliers (those who only finished 6th grade). In Morocco, the CCT delivers 90.8% of the UCT's welfare gains (when focusing on children ages 11-15, which are more likely the non-compliers). In Tanzania, the same figure is 95%, and in Nicaragua, the welfare

[2] We considered a larger set of 19 programs, but only these five met our requirements for program design and data availability (see Section 2).

[3] We follow these papers since they study transfer recipients in a developing country (Indonesia). Studies focusing on developed countries use similar formulations (see Section 3).

gains of the CCT and a UCT are similar.

What explains variation in the CCT-relative-to-UCT welfare gains? Intuitively, three components reduce relative CCT gains: larger transfer sizes, more non-compliance, and non-compliers having lower consumption than compliers. All these factors are substantial in Malawi. For those in Mexico who only finished 6th grade, transfers and the share of non-compliers are large. In Morocco, the transfers are small, while in Tanzania non-compliers are not substantially poorer. In Nicaragua, there are few non-compliers.

Fourth, for each context, we decompose the CCT-relative-to-UCT welfare gains stemming from two mechanisms. The first relates to compliers and non-compliers having different baseline consumption (the "complier selection effect"). The second relates to UCT providing payments to all eligible households, which increases welfare even if non-compliers and compliers have the same baseline consumption distribution (the "unequal payments effect"). In Malawi, the complier-selection effect accounts for one-third of the difference, while in Morocco and Tanzania, we estimate that two-thirds of the welfare gain of CCTs comes from the complier selection effect. In Mexico, almost all welfare gains from a CCT-to-UCT switch come from the unequal payments effect, as expected given that consumption distributions for compliers and non-compliers are similar.

Fifth, we incorporate the value of schooling in our CCT-relative-to-UCT welfare comparisons. We use a revealed-preference argument to derive an upper bound on the value of education for marginal households (those induced to school enrollment by conditionality). Given the relatively small share of marginal households, adding this value, even under assumptions favoring larger CCT welfare gains, has a limited effect on the comparisons.[4]

We also consider a social planner who values education more than households do (due to social returns or households underestimating the returns to schooling). In contexts where the relative welfare gains of a CCT and UCT differ substantially, the social planner

[4] Moreover, evidence is mixed on whether CCTs, relative to UCTs, increase schooling. 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 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 meant to support education) and finds that conditionality "*made almost no difference*."

would need to value education significantly more (12 or 25 times) to equate the gains of the two programs. We discuss why such large multiples might be plausible.

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 (the tagging of households foregoing child labor income) and examines an internal solution involving a mix of CCTs and UCTs. Along these lines, we acknowledge that a policymaker could prefer CCTs solely based on the targeting benefits, even absent CCTs having any treatment effects on increasing schooling. For example, a policymaker could prefer to make transfers to households with school-age children, or might believe that parents that send their children to school are likely to use transfer money the most productively. Our paper addresses a different policy choice between only a CCT and only a UCT, and documents the potential "targeting *cost*" of screening households that do not meet conditions across various contexts.

While the idea that non-compliers and compliers differ, and that only the latter benefit from a CCT, is not new, it receives relatively little attention in the literature. Studies evaluating CCT programs rarely, if ever, report consumption differences by compliance status. Some papers examine how removing conditionality affects non-compliers, but usually focus on outcomes other than consumption. For example, Baird et al. (2011) finds that, compared to a CCT, a UCT reduces early marriage in Malawi, driven by those not enrolled in school. We hope that, by systematically providing consumption comparisons, offering a strategy to assess welfare consequences, and documenting substantial impacts in some contexts, our paper can inform the design of cash transfer programs.[5] [6]

Our paper also relates to the public finance literature studying the welfare and inequality consequences of taxes and transfers, particularly those intended to induce specific behaviors. For example, Kolsrud et al. (2024) studies pension design and finds that incentives to retire later have substantial inequality and welfare consequences (since re-

[5] Benhassine et al. (2015) finds that a labeled cash transfer, compared to a CCT, increases enrollment of those at higher risk of leaving school.

[6] There is also a literature studying the take-up of social programs more generally, which finds poorer households may be less likely to meet program conditions (Herd and Moynihan, 2019; Currie, 2008). To our knowledge this literature does not specifically address CCT versus UCT programs in developing countries.

tirement age and consumption are negatively correlated). Allcott et al. (2019) and Lockwood et al. (2021) explore related issues regarding the regressivity of sin taxes and lotteries.[7]

2 Contexts and Data

The five experiments analyzed in this paper were selected as follows. We began with the 18 randomized trials considered by Banerjee et al. (2017) study of cash transfers' effects on labor supply. We added Robertson et al. (2013), as it is one of the three experiments we know of that includes both CCT and UCT treatment arms.

Although our analysis does not directly use the random assignment of cash transfers, 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.

Our analysis includes experiments that meet 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 enrollment (compliance). Of 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 the two other studies. Additional literature searches did not yield other suitable experiments. See Appendix Table A.1 for details.

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). All programs are operated by governments (except Malawi). 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*." The Tanzania program is not included in Baird et al.

[7] Although not using social welfare calculations as in this paper, Björkegren et al. (2022) and Haushofer et al. (2022) investigate the welfare goals of cash transfer programs.

(2013), but Evans et al. (2014) reports that conditions were monitored and enforced.

Data harmonization. Our analysis samples consists only of children eligible for the CCT. We exclude those in non-CCT and control arms, with income above the eligibility cutoff, and with age below or above eligibility.

We measure consumption *before* CCT payments start (e.g., in a baseline survey before treatment). We adjust consumption so it measures monthly total household consumption per capita. 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).

Our measure of compliance is a dummy for whether the child is enrolled in school.[8] We always use the variable that the source papers use to estimate effects on enrollment in the period they label as the "endline." Observations with missing data on consumption and/or compliance are excluded from the analysis. To make the analysis less sensitive to outliers, we exclude households with consumption below the 5th and above the 95th percentile of their context-specific distributions. The Online Appendix provides further discussion of each context and data definitions.[9]

Measuring compliance. 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 compliance across contexts. Third, data on meeting other conditions is often unavailable. Fourth, in all our contexts, micro-data detailing which households were excluded from payments for non-compliance 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,

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

[9] In Mexico, we treat the May 1999 ENCEL as the endline since it is the last survey before the control group starts receiving treatment.

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 investments. Second, analysis in the source papers for Malawi and Morocco suggests little measurement error 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. Third, if measurement error in the enrollment variable is independent of baseline consumption, our analysis would be unaffected. We are unaware of papers documenting how misreporting of enrollment in surveys correlates with income or consumption.[10]

Program characteristics. Table A.2 provides the key program features for our analysis: share of (non-)compliers and the size of the transfers relative to household consumption, separately by age (as some programs provide different transfers based on the child's age). The Online Appendix provides a more detailed description of each program.

Malawi has the largest share of non-compliers (at 36.8%). For Mexico, transfers are a function of gender and grade of enrollment. Table A.2 shows compliance by expected grade: last completed grade plus one. Many students drop out at the end of primary school (6th grade) - the non-complier share is large at (56%) for those labeled as "7th grade" but smaller (10% or lower) for other grades. In Morocco, the non-complier share is low (below 4%) for those aged 6-10, but substantially larger (ranging from 12% to 77%) for those aged 11-15. In subsequent analyses for Mexico and Morocco, we will focus on both the entire samples and subsamples with higher non-compliance (the "7th-grade" and "ages 11-15" subsamples, respectively). The non-complier share in the Nicaragua RPS is only 4.6%. In the Tanzania program, non-compliance is high among children aged

[10] We do not analyze data from school surveys from Malawi since it is unavailable. Moreover, only 43% of participants are included in such surveys. We cannot use school records data from Morocco in our analysis since the replication files do not link it to the household survey on consumption (understandably, given privacy considerations).

7 but decreases with age until rising again for those 13 and older. Thus, differently from other contexts, non-compliance also includes a margin of delaying when younger children start school.

Regarding average payment sizes, the Mexico, Malawi, and Tanzania programs have large transfer sizes at between 17% to 25% of per capita household consumption. The Morocco and Nicaragua programs have substantially smaller transfer sizes, at 3% and 5.6% of per capita consumption, respectively.

3 Results

Consumption differences between compliers and non-compliers. Figure 1 presents the consumption distributions for compliers and non-compliers across our five contexts. Specifically, it provides histograms of the logarithm of baseline consumption per capita (before payments start). In Mexico and Morocco, we focus on the previously discussed subsamples with higher shares of non-compliers (those expected to enroll in 7th grade and those aged 11-5, respectively).[11] We treat each child as if she comes from a unique household - i.e., we do not consider potential correlations in non-compliance within households.

In Malawi, the non-compliers' distribution is clearly shifted to the left of the compliers' and mean consumption for non-compliers is 26% lower than for compliers. In Morocco, the non-complier distribution also appears shifted to the left, but by a more modest magnitude (differences in averages are 4%). In Mexico, Nicaragua, and Tanzania, there are no clearly visible differences between the distributions. On average, non-compliers consume 4% and 5% less than compliers in Mexico and Tanzania, respectively. In Nicaragua, non-compliers consume 16% *more* on average. However, we caveat that the sample of non-compliers here is small at only 33 households.[12]

[11] Appendix Figure A.1 provides similar figures for full samples.

[12] The difference in average log consumption between compliers and non-compliers in the Morocco and Malawi samples is statistically significant at the 1% level.

Welfare comparisons: empirical strategy. We compare the welfare gains from switching from the implemented CCT to a hypothetical UCT that removes conditionality and provides payments to all households that were eligible for the CCT. We take the set of *eligible* recipients as predetermined and study the welfare within this group, similarly to Bergstrom and Dodds (2021). The counterfactual UCT has the same aggregate transfer budget. We implement this by scaling down transfers: if t^{cct} is the monthly transfer paid to compliers in a CCT and s is the share of compliers in the program, we set the UCT transfer t_i^{uct} equal to $s \cdot t_i^{cct}$.[13]

Let c_i^{cct} denote household i's per capita consumption when enrolled in the CCT program. Empirically, we set $c_i^{cct} = c_i^{baseline}$ if i is a non-complier and $c_i^{cct} = c_i^{baseline} + t_i^{cct}$ if i is a complier. $c_i^{baseline}$ is the baseline monthly household per capita consumption we directly observe. Let c_i^{uct} be i's consumption in the counterfactual UCT program. We set $c_i^{uct} = c_i^{baseline} + t_i^{uct}$ for all i in our samples (which only includes eligible households). Let $u(c_i)$ denote i's utility. Its change from the transition from the CCT to the UCT, $u(c_i^{uct}) - u(c_i^{cct})$, is always positive for non-compliers, as they go from not receiving to receiving transfers, and always negative for compliers, as the size of their transfer payment decreases.

Our approach is static and abstracts from borrowing and saving (i.e., the marginal propensity to consume is one). It also abstracts from transfers' effects on labor supply, but a review (Banerjee et al., 2017) finds those to be small or zero.[14] [15]

Our key measure is the ratio between the CCT welfare gain (relative to no transfers) and the UCT utility gain (relative to no transfers):

$$\frac{\sum_{i=1}^{N}[u(c_i^{cct}) - u(c_i^{baseline})]}{\sum_{i=1}^{N}[u(c_i^{uct}) - u(c_i^{baseline})]} \tag{1}$$

[13] Transfers have an i subscript since, in Mexico and Morocco, CCT payments become larger with age. We calculate UCT counterfactual payments assuming that older children also get larger transfers: payments are all scaled down by the same factor. The Mexican program provides larger payments for girls than boys. We again maintain the same proportional difference by gender in our counterfactual UCT program.

[14] Crosta et al. (2024) studies 72 UCT programs in a Bayesian meta-analysis and finds positive effects on labor supply, but we the number of studies directly comparing CCTs vs. UCTs differential effects on labor supply are small and there does not seem to be a consensus in the literature on this.

[15] The evaluations we study provide limited information on borrowing and saving and also do not collect consumption data over enough periods to further study its dynamics.

where N is the number of households in a sample.

Equation (1) has four desirable properties. First, it has an intuitive interpretation. If it equals one, it indicates the CCT delivers the same aggregate utility gains as the UCT. A value of 0.5 would indicate that the CCT program delivers only half the welfare benefit of a UCT program that has the same aggregate transfer budget. Second, it provides a comparable metric for each context we study. Third, it is unaffected by any positive affine transformation of $u(\cdot)$. Fourth, our measure uses information on the entire distribution of consumption in complier and non-complier households (and not only their mean differences) without making assumptions about the shape of such distributions.

We abstract from the costs of monitoring and enforcing conditionality, although they can be sizable and likely to "tilt the scales" towards UCTs' effectiveness. Caldés et al. (2006) estimate that conditionality costs alone (i.e., separately from other administrative costs such as determining eligibility) account for almost 20% of Progresa's administrative costs and equal 2% of the program's transfer budget. However, one may consider that transfers themselves are not a social cost (and only the deadweight loss of raising the funds for them is a cost). Benhassine et al. (2015) indicates the administrative costs of Tayssir's CCT are 30% higher than the labeled cash transfer and provides further discussion of how administrative costs make standard CCT programs expensive.

Our method also assumes that a child's school enrollment implies the household receives the transfer. If the program is less strict and provides payments to non-enrolled children, our results would overestimate the benefits of switching to a UCT for poorer non-compliers. This discrepancy also affects the interpretation of the original experiments on CCT programs, as households may initially comply, expecting enforcement, but later adjust their behavior upon learning conditions are not strictly applied.

Welfare comparison: results. Table 1 presents the ratio 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 $\rho = 3$. To probe the sensitivity of results to the concavity of the utility function, we also provide estimates using ρ equal

to one ("ln utility"), two, and four. As before, we treat each child as if she comes from a unique household.[16]

Our baseline estimates indicate the Malawi CCT only delivers 68.4% of the utility gain of an equally budgeted UCT program. In other words, the budget-neutral switch to a UCT would increase welfare by $46.2\% = (1/0.684 - 1)$. In Mexico, when focusing on the 7th-grade subsample, a CCT only delivers 71.3% of the UCT gains. The corresponding number using the entire Mexico sample is 94.8%. The CCT-to-UCT welfare gains ratios for Morocco and Tanzania range between 90.8%-96.0%, and in Nicaragua it is close to one. We further discuss comparisons across contexts in Section 4.

As expected, ratios become closer to one as we consider utility specifications with less curvature. However, they remain substantial in the cases of Malawi and the 7th-grade subsample in Mexico. For example, with $\rho = 2$, the Malawi CCT delivers only 74.9% of the welfare gains of the UCT, and with ln utility, 85.3% of the gain.

Decomposing mechanisms: empirical strategy. Two mechanisms can affect the CCT-to-UCT welfare gains ratio. The first are differences in the consumption distribution of compliers and non-compliers. If compliers have higher baseline consumption, CCT payments will target the richer households with lower marginal utility of consumption. We label this the "complier selection effect." Second, even if the baseline consumption distribution for complier and non-complier households is exactly the same, the ratio in equation (1) would be smaller than one. This is because the CCT only targets a share of eligible households (the compliers) while the UCT provides payment 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 this as the "unequal payments effect" and define it as the ratio in equation (1) that would occur if the distribution of baseline consumption for compliers was the same as for non-compliers.[17]

We decompose our observed CCT versus UCT gains into complier selection and un-

[16] The choice of ρ in the social insurance literature usually ranges from one to four (Chetty and Finkelstein, 2013).

[17] A simple example illustrates the unequal payments effect. Assume two households have the same baseline consumption c. Providing a transfer of t to one increases aggregate utility by less than providing $t/2$ to both if there is diminishing marginal utility: $u(c+t) - u(c) < 2[u(c+t/2) - u(c)]$.

equal payments effects as follows. Let s be the share of non-compliers and N the number of eligible households in a given context. We randomly draw sN households, "pretend" these households are non-compliers, and recalculate the ratio in equation (1). We replicate this process 1,000 times. In *expectation*, the consumption distributions of those randomly labeled as non-compliers are the same. Thus, the *average* of the 1,000 simulated CCT-to-UCT welfare gain ratios is driven only by the unequal payments effect. Since our main results in Table 1 include both effects, comparing them to this average simulated ratio provides a decomposition.

Decomposing mechanisms: results. Figure **??** presents the results using our baseline utility specification (CRRA with $\rho = 3$). Each panel shows in blue the distribution of the ratio between the simulated CCT-to-UCT welfare gains ratio across the 1,000 draws. The dashed vertical line shows the mean of this distribution, and the solid line shows the "observed ratio," which is based on the actual compliers observed in the data and reported on Table 1).[18]

In Malawi, the average simulated ratio is 0.809, indicating that, even if there was no difference between the consumption distribution of compliers and non-compliers, a CCT would only deliver 80.9% of the welfare gains of a UCT given the unequal payments mechanism. Comparing it to the observed ratio from Table 1 of 0.684 indicates the unequal payments mechanism accounts for approximately two-thirds of the gap between the welfare gains that a CCT and a UCT provide, with the complier selection mechanism accounting for the other third. For Morocco and Tanzania the unequal payments mechanism approximately accounts for one-third of the CCT-to-UCT welfare gains ratios. In Mexico, almost all the CCT-to-UCT welfare gains ratio can be accounted for by the unequal payments effect. This is expected given that non-compliers' and compliers' consumption are similar in this context. For Nicaragua, there is essentially no utility difference to decompose (both simulated and observed ratios are close to one).[19]

[18] The distribution can have support above one since some draws may have (randomly assigned) non-compliers with systematically higher baseline consumption than compliers.

[19] The Mexican and Moroccan cases focus on the subsamples with higher shares of non-compliers. Appendix Figure **??** presents the same analysis for the full samples.

Incorporating the value of schooling. Given our focus on conditionality exacerbating consumption inequality and the mixed evidence on whether CCTs induce more schooling relative to UCTs (discussed in footnote 4), we have abstracted from human capital investment. However, as this is a key part of a comprehensive cost-benefit analysis, we now discuss how human capital might affect our analysis.

The key components of our approach are: i) only marginal households (induced to enroll their children by the conditionality) are relevant for social welfare comparisons, and ii) for a given cost of schooling, marginal households *not* enrolling their children absent conditionality provide an upper bound on the value they place on schooling.

We consider a binary school enrollment decision. Under a CCT, household i utility is:

$$u_i^{cct} = \begin{cases} u(c_i^* + t_i^{cct} - f_i) + D_i, & \text{if child enrolled} \\ u(c_i^*), & \text{if child not enrolled} \end{cases} \tag{2}$$

and under a UCT:

$$u_i^{uct} = \begin{cases} u(c_i^* + t_i^{uct} - f_i) + D_i, & \text{if child enrolled} \\ u(c_i^* + t_i^{uct}), & \text{if child not enrolled} \end{cases} \tag{3}$$

where $f_i \geq 0$ is the consumption opportunity cost from schooling (e.g., foregone child labor income). D_i is how much i values enrolling the child (in utils). It captures the discounted utility of future additional income and the child's preference over work, school, and leisure. c_i^* is per capita consumption absent transfers and schooling costs. Banerjee et al. (2024) uses a similar formulation to discuss the welfare consequences of conditionality.

We define a household that enrolls its child under a CCT but not a UCT as *marginal*. Other households are *inframarginal*. If i is marginal, its choice of not enrolling under a UCT and equation (3) imply:

$$D_i \leq u(c_i^* + t_i^{uct}) - u(c_i^* + t_i^{uct} - f_i) \tag{4}$$

providing an upper bound on D_i which is increasing in f_i. Intuitively, higher values of

schooling require higher costs to make marginal households prefer not to enroll (under a UCT). For example, f_i cannot be zero when $D_i > 0$.

In principle, equation (4) and knowledge of f_i and which households are marginal could be used to add upper bounds on the value of schooling to our CCT-to-UCT welfare gains ratio (equation 1). In practice, f_i and the marginals' identity are unobserved.

We thus proceed as follows. We make three "conservative" assumptions that "tilt the scales" towards CCT providing higher welfare gains than UCTs. First, UCTs, relative to no transfers, do not induce more schooling. Second, equation (4) holds with equality for all marginal households (i.e., we assume all marginal households are indifferent to enrolling under a UCT). Third, we set $f_i = 0.15 \cdot c_i^*$ (i.e., school enrollment reduces a marginal household's total consumption by 15%). This is a deliberately large cost chosen since larger f_i imply larger upper bounds on D_i and CCT welfare gains. We discuss robustness to f_i values below.

These assumptions allow us to recalculate the CCT-to-UCT welfare gains ratios (equation 1) incorporating the value of schooling. Its denominator (UCT welfare gains) is unaffected given our first assumption. The numerator only changes for *marginal* households: their welfare gains under the CCT are now the difference between the two terms in equation (2). For *inframarginal* households, the numerator remains unchanged.

The final component is "choosing" which households are marginal. We exploit that the CCT effect on enrollment (relative to no transfers) equals S^{marg}: the *share* of marginal households. For each of our contexts, we obtain S^{marg} from the effects reported in randomized evaluations (our "source papers," see Online Appendix). Since we observe the share of marginals but not their identity, we take a Monte Carlo approach. We randomly assign $S^{marg} \cdot N$ children to be marginals. We repeat this 1,000 times, calculating 1,000 separate CCT-to-UCT welfare ratios for each context. The Online Appendix provides further details on deriving the bounds and its implementation. It demonstrates that incorporating the value of education can only increase the CCT-to-UCT welfare ratio, and discusses how to reconcile this with the common intuition that imposing conditions reduces house-

hold welfare.[20]

Table 2 presents the 1st and 99th percentiles of this CCT-to-UCT welfare ratio distribution in columns (3) and (4). Column (2) replicates the baseline ratio from Table 1 for comparison. All calculations assume $\rho = 3$. In all contexts, both percentiles are closer to the baseline ratio (which excludes gains from schooling) than to one (where CCTs and UCTs provide similar welfare). The exception is Nicaragua, where the baseline ratio is above one. This suggests that incorporating the value of returns to schooling, even with assumptions favoring CCTs, has a limited quantitative impact on our welfare comparisons, leaving the conclusions from our baseline case largely unchanged. The Online Appendix discusses robustness to the different assumptions.[21]

The intuition behind the small effect of adding the value of schooling on welfare comparisons is that the share of marginal households (reported in column 1) is small. Except for the unique case of Nicaragua, this share ranges from 3.2% to 6.5%. Additionally, households cannot place an unboundedly large value on education; otherwise, they would enroll their children even without conditionality.[22]

Social returns to schooling. A common rationale for CCTs is that the social planner (government) values education more than households. This can occur given social returns to schooling or households underestimating its private returns or undervaluing the utility of future generations.[23] The Online Appendix discusses how, following Banerjee et al. (2024), we can incorporate a social planner that values the net gains from schooling θ times more than households do.

In particular, we calculate how large θ must be to make the social planner derive the same welfare gains from a CCT and a UCT. We label this θ^* and report it in column (5)

[20] It also discusses that non-complier households are not eligible to be marginal in the random assignment (they are, by definition, inframarginal) and that that choosing the households with lowest or highest consumption as marginals does not correspond to an upper bound.

[21] In particular, Appendix Table A.4 shows that doubling f_i (assuming that enrollment lowers consumption by 30%) would not substantially change this overall conclusion. The baseline case (column 2) is the same as assuming $f_i = 0$. The CCT-to-UCT welfare ratio is increasing in f_i, thus setting it to values between zero and 0.30 would also yield ratios "in between" the ones we report.

[22] The share of marginals reflects a CCT versus pure control comparison in randomized evaluations, making them conservative estimates likely larger than those from comparing a CCT to a UCT.

[23] This would include households under-valuing future non-education benefits such as reduced child marriage or teen pregnancy for CCTs targeted at adolescent girls (Baird et al., 2019).

of Table 2. For example, $\theta^* = 26.3$ for Malawi indicates the social planner would have to value education 26.3 times more than households do to make it indifferent between a UCT and a CCT. The values for Mexico vary between 12 and 17.3. Even in contexts where the baseline CCT-to-UCT welfare ratio is close to one (Morocco and Tanzania), θ^* is large (about 3.5 and 7.1). We interpret this as suggesting that, to make a meaningful change to our qualitative conclusions regarding the welfare consequences of conditionality, the difference between the value the social planner and households put on schooling has to be large.[24]

4 Discussion

We conclude by discussing why the CCT-to-UCT welfare gains ratios vary across the five contexts we study. Intuitively, three ingredients can reduce the relative gains of a CCT. The first is the transfer amount relative to baseline consumption. The second is the share of non-compliers. The third is the extent the non-complier consumption distribution is "shifted to the left" of the complier distribution.

Malawi combines all three ingredients. It has the largest transfers, the highest share of non-compliers, and the clearest difference in consumption distributions among all five contexts. Not surprisingly, it has the smallest CCT-to-UCT welfare ratio, at 68.4%.

Mexico, which also involves large transfer sizes, illustrates the role of the share of non-compliers. In the subsample where non-compliance is substantial (those expected to enroll in 7th grade), the CCT only delivers 71.3% of a UCT's welfare gains, but incorporating the remainder of the sample, where compliance is larger, moves the ratio closer to one. Moreover, the decomposition analysis highlights that, even when the consumption distributions of compliers and non-compliers are similar, the welfare losses from a CCT exacerbating inequality can be substantial due to the "unequal payments effect" if the share of non-compliers is large.

Morocco illustrates the role of transfer sizes. As they are small (on average 3% of base-

[24] In Nicaragua, θ^* is undefined as the baseline ratio is above one. Appendix Table A.4 reports θ^* assuming $f_i = .3 \cdot c_i^*$.

line per capita consumption), the CCT-to-UCT welfare ratio is not substantially far from one, regardless if one focuses on a subsample where non-compliers are more common (those aged 11-15, where the CCT-to-UCT welfare ratio is 91%) or the entire sample that has a smaller share of non-compliers (where the ratio is 96%).

In Tanzania, both the transfer size and share of non-compliers are smaller than in Malawi and Mexico, and thus the CCT-to-UCT welfare ratio is closer to one. The ratio is essentially one in Nicaragua; this can be attributed to its small transfer size and share of non-compliers, but we reiterate the caveat of the small sample of non-compliers.

Our results suggest that conditionality requirements that exacerbate consumption inequality among eligible households are quantitatively important for the overall welfare impacts of cash transfer programs. We believe the issue has not received its due attention in previous literature and hope the type of assessment we provide here can be used to inform the design of cash transfers.

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.

_ , _ , and _ , "When the money runs out: Do cash transfers have sustained effects on human capital accumulation?," *Journal of Development Economics*, 2019, *140*, 169–185.

_ , 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.

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.

Björkegren, Daniel, Joshua E Blumenstock, and Samsun Knight, "(Machine) Learning What Policies Value," *arXiv preprint arXiv:2206.00727*, 2022.

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.

Currie, Janet, "The invisible safety net: Protecting the nation's poor children and families," 2008.

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.

Heller, Sara B, Anuj K Shah, Jonathan Guryan, Jens Ludwig, Sendhil Mullainathan, and Harold A Pollack, "Thinking, fast and slow? Some field experiments to reduce crime and dropout in Chicago," *The Quarterly Journal of Economics*, 2017, *132* (1), 1–54.

Herd, Pamela and Donald P Moynihan, *Administrative burden: Policymaking by other means*, Russell Sage Foundation, 2019.

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

Lockwood, Benjamin, Hunt Allcott, Dmitry Taubinsky, and Afras Y Sial, "What drives demand for state-run lotteries? evidence and welfare implications," Technical Report, National Bureau of Economic Research 2021.

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.

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

Table 1: CCT-to-UCT Welfare Gains Ratios

		Ratio of CCT Utility Gain to UCT Utility Gain	
Sample	Share of Compliers	**CRRA ($\rho = 3$)**	Avg Simulated Ratio
Panel A: Full Samples			
Malawi	0.632	**0.684**	0.809
Mexico	0.803	**0.948**	0.959
Morocco	0.792	**0.960**	0.986
Nicaragua	0.954	**1.016**	1.000
Tanzania	0.844	**0.950**	0.984
Panel B: Subsamples			
Mexico 13-18 year olds	0.577	**0.823**	0.821
Morocco 11-15 year olds	0.585	**0.908**	0.964
Tanzania 13-15 year olds	0.782	**0.893**	0.979

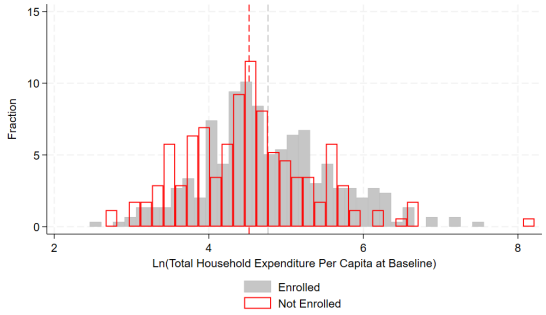
The table presents the CCT-to-UCT welfare gains ratio: the ratio between the welfare gains from the observed CCT and a counterfactual UCT program that uses the same budget to provide payments to all eligible households. In particular, it provides estimates of equation (1) discussed in Section 3. Our baseline is the case with $\rho = 3$, in bold.

Table 2: CCT-to-UCT Welfare Gains Ratios Incorporating the Value of Schooling

Sample	S^{marg} (1)	Original (2)	1st percentile (3)	99th percentile (4)	θ^* (5)
Panel A: Full Samples					
Malawi	0.058	0.684	0.687	0.696	26.33
Mexico	0.032	0.948	0.951	0.951	17.33
Morocco	0.054	0.960	0.967	0.971	3.64
Nicaragua	0.128	1.016	1.019	1.021	-
Tanzania	0.04	0.950	0.951	0.957	7.14
Panel B: Subsamples					
Mexico 13-18 year olds	0.065	0.823	0.835	0.838	11.8
Morocco 11-15 year olds	0.054	0.908	0.921	0.935	3.41
Tanzania 13-15 year olds	0.04	0.893	0.894	0.906	8.23

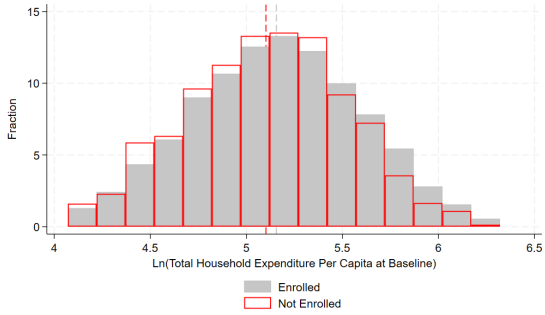
The table presents the CCT-to-UCT welfare gains ratio, incorporating the value of schooling based on our bounding approach. Column (1) presents the estimated share of marginal households. Column (2) replicates the baseline case from Table 1 (with $\rho = 3$ and no value of schooling considerations: $f_i = 0$). Columns (3) and (4) present the 1st percentile and 99th percentile of the distribution of simulated CCT-to-UCT welfare ratios based on randomly assigning which compiler households are marginal. Column (5) provides θ^*: a calculation of how many more times the social planner would have to value the net gains of schooling than households do to be indifferent between a CCT and a UCT (a welfare ratio of one). All calculations in columns (3)-(5) are based on a CRRA utility with $\rho = 3$ and setting the cost of schooling f_i equal to 15% of household i's baseline consumption. See Section 3 for further information. Although both columns (3) and (4) present the same number for Mexico, the numbers differ in their fourth digit but both round to 0.951. For Morocco and Tanzania, we use the same share of marginals (column 1) for both subsamples.

Figure 1: Baseline Consumption Distribution by Complier Status



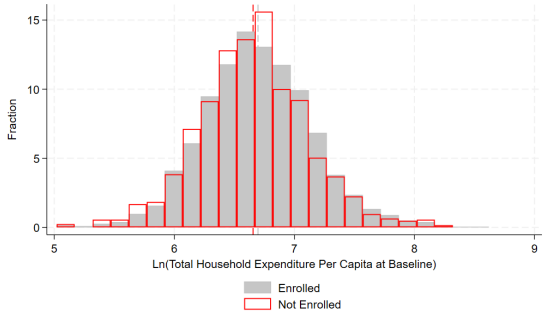
(a) Malawi

Complier mean: 4.77, Non-complier mean: 4.51



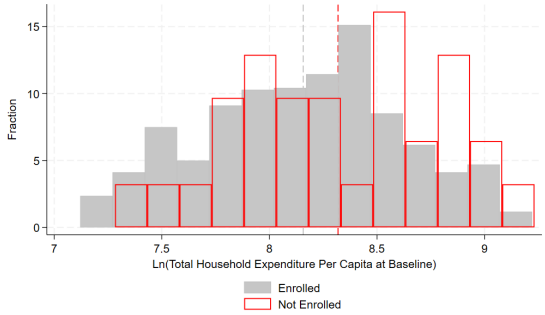
(b) Mexico

Complier mean: 5.16, Non-complier mean: 5.10



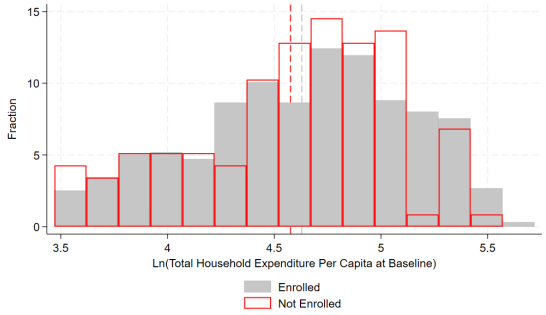
(c) Morocco

Complier mean: 6.70, Non-complier mean: 6.66



(d) Nicaragua

Complier mean: 8.16, Non-complier mean: 8.32

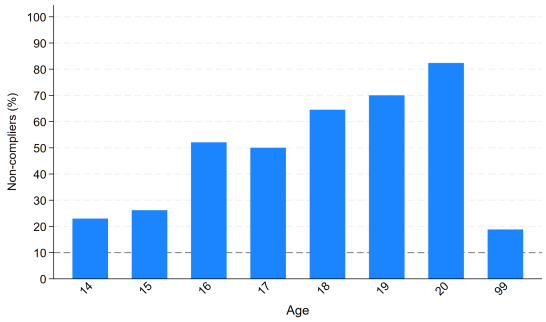


(e) Tanzania

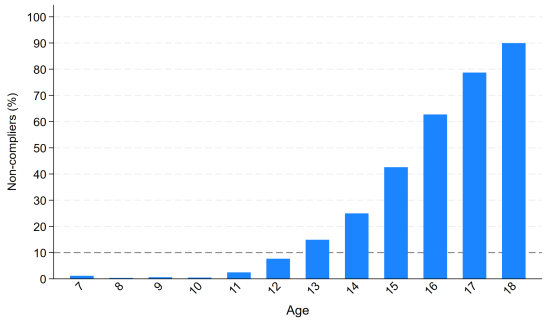
Complier mean: 4.63, Non-complier mean: 4.58

Each sub-figure shows, for a given context, histograms of the logarithm of total household consumption per capita at baseline (i.e., before CCT payments start). Solid gray-colored bars provide the distribution for compliers and the red-bordered bars do so for non-compliers. Vertical lines indicate the mean consumption for each group. The sample only includes eligible children in the CCT arm of their respective experiment. Compliance is defined as being enrolled in school.

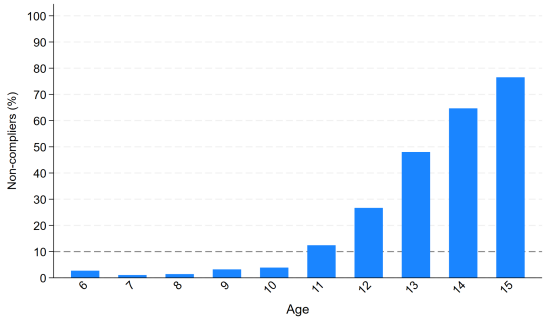
Figure 2: Share of Non-compliers by Age



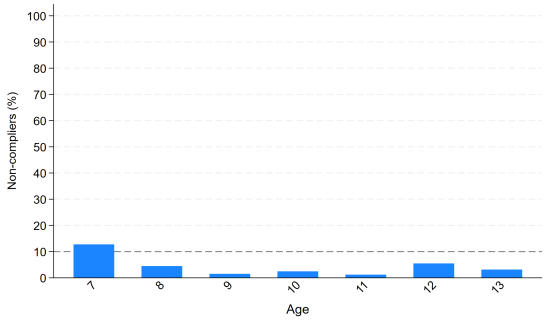
(a) Malawi
N: 470



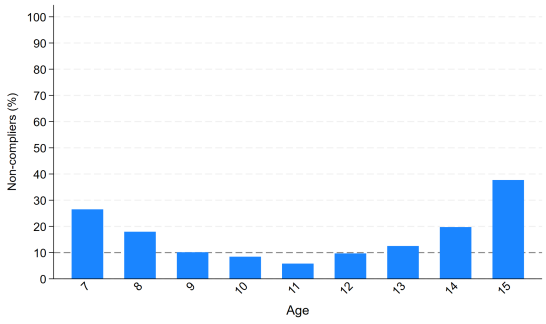
(b) Mexico
N: 11,093



(c) Morocco
N: 5,988



(d) Nicaragua
N: 712



(e) Tanzania
N: 753

Each sub-figure shows, for a given context, histograms of the percentage of non-compliers by age. Compliance is defined as being enrolled in school. A dashed line is added in each figure at the 10% level. Under each figure, we have the number of observations for each sample.

Online Appendix

A 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.

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 only on never-married women aged 13-22 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 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 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?"

Table A.2 shows the number of compliers and non-compliers by age within the CCT arm of the Malawi experiment. We define compliance as being enrolled in school in Term 1, 2010. The fraction non-complying increases substantially with age, from 23% for 14-year-olds up to 65% - 82% for 18- to 20-year-olds. The average payment size is substantial (close to 25% of average per capita consumption). Baird et al. (2011) also uses non-enrollment in Term 1, 2010 to define non-compliers in their analysis of effects on pregnancy and marriage.

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 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.[1]

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?"

Table A.2 shows the number of compliers and non-compliers by grade among eligible children in treated communities, by grade of enrollment. We define compliance as being enrolled in school in the May 1999 ENCEL survey. The survey does not explicitly collect data on current grades. We infer this by adding one to the last completed grade to determine the grade a child should be enrolling at. A substantial amount of students drop out at the end of primary school (6th grade in Mexico). Thus the share of non-compliers is large (56%) for those labeled as "7th grade" but smaller (10% or lower) for other grades. In subsequent analysis, we will focus both on the entire sample of eligible children and also on the subsample of those that only completed sixth grade (the "7th-grade" subsample, henceforth). The average payment size is substantial (approximately 23% of per capita consumption).

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

[1] 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.

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 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 6,203 children from both the CCT to fathers and CCT to mothers arms.[2]

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." Program take-up was 97%.

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

Table 1 shows the number of non-compliers and compliers by age. We define compliance 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). The share of non-compliers is low (below 4%) for those aged 6-10, but substantially larger (ranging from 12% to 77%) for those aged 11-15. In subsequent analysis, we focus both on the entire sample of eligible children and the subsample aged 11-15. A key feature of the Morocco experiment, also noted in Benhassine et al. (2015), is the relatively small payment size, close to 3% of per capita consumption (compared to 25% in Malawi and 22% in Mexico).

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 714 eligible children in the

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

treated *comarcas*.

RPS featured two core components. First, the education component provided an annual C$1,440 (1112 USD) per child aged 7-13, conditional on enrollment and maintaining at least 85% attendance, plus an additional C$275 (21 USD) for school supplies. 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.

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.

Table A.2 shows the number of compliers and non-compliers by age for eligible households in treated *comarcas*. We define compliance as being enrolled in school in the endline survey conducted in October 2002. Compared to other contexts, the share of non-compliers in the RPS program is small, at only 4.6%. Moreover, our sample only includes 33 non-compliers, making comparisons less precise. The payment size relative to per capita consumption levels is small at 5.6%.

Tanzania (TASAF). The source paper for this experiment is Evans et al. (2014). The TASAF experiment took place in three Tanzanian districts. 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 in the treated villages.

Payments were determined by household composition. Each child aged 7–15 received USD$1.5 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.

Table A.2 shows the number of compliers and non-compliers by age. We define compliance as being enrolled in school in the endline survey conducted in late 2012. Non-compliance is high among children aged 7 but decreases with age until rising again for those 13 and older. Thus, differently from other contexts, non-compliance also includes a margin of delaying when younger children start school. Payment sizes are substantial (close to 17% of per capita consumption), but not as large as in Malawi and Mexico.

B Incorporating the Value of Schooling in the CCT-to-UCT Welfare Ratio

This Appendix provides further information on our approach for incorporating the value of schooling in the CCT-to-UCT welfare ratio calculations.

Framework. As in the main text, we consider N households, each with one child and facing a binary school enrollment decision. If household i does not enroll its child in school and does not receive a transfer, it consumes c_i^*. We later discuss how c_i^* relates to observed baseline consumption. If i enrolls its child, its utility is:

$$u(c_i^* + t_i - f_i) + D_i$$

where t_i is the transfer (equal to zero, t^{uct}, or t^{cct}, depending on the scenario being considered). $f_i \geq 0$ is the consumption opportunity cost of enrolling a child in school (e.g. foregone income from child labor). D_i is how much the household values school enrollment (in utils). D_i is a combination of the net present value of the utility of future additional income from more schooling, preferences of the child relative to work, leisure, and school, and also direct utility (e.g., the pride of enrolling a child). If a household chooses non-enrollment, it does not incur f_i and does not receive D_i. Banerjee et al. (2024) use a similar a utility function to discuss the welfare effects of conditionality, but do not proceed with deriving bounds and using it in empirical analysis.

Household types. If u is concave, there are only four potential types of households when it comes to enrollment behavior under our three relevant scenarios: no transfers, a UCT, and a CCT. Type (i): those not enrolling under no transfers, UCT, and CCT. Type (ii): those not enrolling under no transfers, but doing so under UCT and CCT. Type (iii): those only enrolling under a CCT. Type (iv) those that enroll their child in all scenarios. There are only four types since the income effect on enrollment is positive.[3]

We assume there are no type (ii) households, which is equivalent to assuming the UCT does not affect schooling, relative to no transfers. This is the "first assumption" discussed in Section 3 of the main text. This is a conservative assumption in the sense that it works "against" a UCT and "in favor" of a CCT in our welfare ratio calculations. It also simplifies the analysis, as we discuss below.[4]

Assuming no type (ii) households implies there are only three types of households. First, the *marginal* households that only enroll their children under a CCT, but not under a UCT (and no transfers). Second, the *inframarginal compliers* that enroll their children under a CCT, UCT, and no transfers. Third, the *inframarginal non-compliers* that do not enroll their children under a CCT, UCT, and no transfers. Given the definitions in this paper, all marginals are compliers (they enroll their children under a CCT), but not all compliers are marginals (some are inframarginal).

Gains from UCTs and CCTs. We now define the gains from a UCT and CCT (the denominator and numerator of our CCT-to-UCT welfare ratio in equation (1)) for each of the three relevant groups.

For *inframarginal non-compliers*, the gains are the same as in our baseline analysis that

[3] For example, if i chooses non-enrollment under a UCT, it would also choose it under no transfers since $D_i \leq u(c_i^* + t_i^{uct}) - u(c_i^* + t_i^{uct} - f_i) \leq u(c_i^*) - u(c_i^* - f_i)$ always holds if u is concave and $f_i \geq 0$. A similar argument would show that, if a household chooses enrollment under a CCT, it would also do so under a UCT (since $t^{cct} \geq t^{uct}$).

[4] There is empirical evidence suggesting 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 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. Our "no type (ii)" 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 type (ii) households.

abstracts from the value of schooling (which implicitly assumes $f_i = D_i = 0$). The gain from a UCT is $u(c_i^* + t^{uct}) - u(c_i^*)$. The gain from a CCT is zero since non-compliers, by definition, do not receive a conditional transfer: $u(c_i^*) - u(c_i^*)$. In this case, it is plausible to assume $c_i^* = c_i^{baseline}$, since in the baseline period there are no transfers, the household is choosing non-enrollment, and would not be incurring f_i.

For *inframarginal compliers*, the gains are also the same as in the baseline analysis. The gain from a UCT is $u(c_i^* + t^{uct} - f_i) + D_i - u(c_i^* - f_i) - D_i$. The gain from a CCT is $u(c_i^* + t^{cct} - f_i) + D_i - u(c_i^* - f_i) - D_i$. In both cases, D_i cancels out. For this group, it is reasonable to assume that $c_i^* - f_i = c_i^{baseline}$, since in the baseline period the household is choosing enrollment and thus incurring f_i. Thus, equation (1) for inframarginal compliers remains unchanged.

For *marginals*, the gain from a UCT is $u(c_i^* + t^{uct}) - u(c_i^*)$, the same as in the baseline analysis. Again, it is plausible to assume $c_i^* = c^{baseline}$ here. In the baseline period there are no transfers, so they would choose non-enrollment. However, the gain from a CCT for marginal households is

$$u(c_i^* + t^{cct} - f_i) + D_i - u(c_i^*) \tag{5}$$

which is different from the baseline case. We do not observe f_i or D_i directly. We can, however, derive an upper bound as follows.

Deriving bounds on D_i. By definition, a marginal household, under a UCT, obtains higher utility from not enrolling their child than enrolling her. Therefore, the following holds if i is marginal:

$$u(c_i^* + t_i^{uct} - f_i) + D_i \leq u(c_i^* + t_i^{uct})$$

which can be rearranged to

$$D_i \leq u(c_i^* + t_i^{uct}) - u(c_i^* + t_i^{uct} - f_i)$$

which is equation (4) in the main text and provides the upper bound on D_i. It also implies an upper bound on the gains from a CCT, relative to no transfers, for marginals:

$$u(c_i^{baseline} + t^{cct} - f_i) + [u(c_i^{baseline} + t_i^{uct}) - u(c_i^{baseline} + t_i^{uct} - f_i)] - u(c_i^{baseline}) \quad (6)$$

which we obtain by assuming equation (4) holds with equality and substituting it in the place of D_i in equation (5), and also from the fact that it is plausible to assume $c_i^* = c_i^{baseline}$ for marginal households, as discussed above. f_i is the only argument in equation (6) not observed in the data. Thus, for a choice of f_i, it provides an upper bound on the gain from the CCT for marginal household i.

Both equation (6) and the right-hand-side of equation (4) are increasing in f_i. This implies that assuming larger f_i can only increase the potential gains of CCTs, leaving the gains from UCTs unaffected. Moreover, if $f_i = 0$, equation (6) becomes the CCT gain in the baseline case (the numerator in equation 1). Thus, incorporating the value of schooling, under our approach, can only increase the CCT-to-UCT welfare ratio.[5]

Discussion and intuition. The key components of our approach are that only the value of schooling for *marginal* households matters for a CCT versus UCT welfare comparison and, by definition, marginal households choose not to enroll under a UCT.

Intuitively, the CCT is unlikely to change the behavior of households that have high value and low costs of schooling: these children would be enrolled under the UCT as well. It is possible that these marginal households have high returns to schooling, but this statement then implies that these children must also have high costs of schooling (again, because they would also be in school under the UCT otherwise). This is also the intuition for why the gains from a CCT are increasing in f_i. We later discuss households underestimating the returns to schooling or not fully internalizing the benefits of schooling for the child.

Moreover, assuming $f_i = 0$ and $D_i > 0$ is inconsistent with equation (4). If enrolling a child does not reduce i's consumption today, then household i choosing non-enrollment

[5] The derivative of equation (6) with respect to f_i is $-u'(c_i^{baseline} + t^{cct} - f_i) + u'(c_i^{baseline} + t_i^{uct} - f_i)$, which is positive since u is concave and $t^{cct} > t^{uct}$.

implies the net discounted value of schooling (D_i) is zero or negative. This highlights that both the current costs of schooling and future benefits must be incorporated in a framework that relies on revealed preference arguments.

The intuition behind the bound in equation (6) 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. We assume, conservatively, that all marginal households are in this upper bound and are indifferent between losing f_i of consumption and receiving D_i of utility under a UCT. Under a CCT, the utility loss from incurring f_i is lower than in the UCT (because CCT payments are larger than UCT payments and a concave u thus makes marginal utility lower under a CCT than UCT). 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 (6) simplifies to the numerator of equation (1) and equation (4) becomes $D_i \leq f_i$.

At first pass, our 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 transfer to allow non-compliers to receive them. In other words, breaking budget neutrality and assuming $t^{cct} = u^{uct}$ 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 this paper: conditionality has two consequences: inducing some behavior and determining who receives the transfers (and thus their size, given a budget). The "common intuition" abstracts away from the latter consequence.

Estimating the share of marginals. While we do not observe *which* households are marginal, we can exploit that the *share* of marginal households (s^{marg}) equals the treatment effect of a CCT, compared to a pure control. The assumption that UCTs, relative to no controls, do not increase enrollment plays a role here. However, this is a conservative assumption, in the sense it works "in favor" of larger CCT gains and "against" UCT gains.

For each of our contexts, we locate in the relevant source papers their "headline" effect of a CCT versus a pure control on enrollment. Recall 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). This is the longest-run treatment effect (year 3) of their CCT program using their preferred measure of school enrollment.

For Mexico, we use the results from Table 6, row (3) from Parker and Todd (2017) (which in turn 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. All of these impacts are estimated 18 months after treatment starts. For the 7th-grade sample, we use the average of the effects across the genders for the 7th-9th grade groups: $.065 = (.045 + .085)/2$. For the entire Mexico sample we use .0323 which equals the mid-points of the within gender results, averaged acrossed the genders, and then average across the grades using the number of grades in each group: $.0323 = \frac{1}{7}[(.0155 + 0)/.02] * 4 + [(.045 + .085)/2] * 3]$.

For Morocco, we set $S^{marg} = 0.053$, 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 ($.053 = \frac{1}{2}(.074 - .019) + (.074 - .021)$). The paper does not separately report effects for the 11-15 group. We assume the same estimates hold for the Morocco-Ages 11-15 and Morocco-All samples. We prefer this approach, which "ties our hands," to estimating effects not provided in the source papers ourselves.

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). This is based on the endline results. They label this estimate as including the sample of 0-18 year olds, which may include children too young to be eligible (under 5). Evans et al. (2014) also reports an effect of 0.030 for those aged 7-14, thus $S^{marg} = 0.040$ appears as a reasonable approximation.

Empirical implementation: assigning marginals. While we know the *share* of marginal households (S^{marg}), we do not know which specific households are marginal. As previ-

ously discussed, by definition, only complier households (who enroll their child 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 complier households to be assigned as marginal. Below, we discuss that choosing the $S^{marg} \cdot N$ households with lowest or highest baseline consumption to be the marginals would not necessarily yield an upper bound in our approach.

Empirical implementation: choice of f_i. In the main text, we set f_i equal to 15% of i's baseline consumption in our main estimates: $f_i = 0.15 \cdot c_i^{baseline}$. This is a purposely large number, implying that enrolling a child will reduce *total* household consumption by 15%. This is a "conservative" assumption given that, as discussed above, the CCT gains are increasing f_i, while the UCT gains do not vary with f_i. We also report results assuming f_i equals 30% of baseline consumption in Table A.4. Since our baseline results (Table 1) are the same as assuming $f_i = 0$, we probe robustness to a range of f_i. Note, moreover, that CCT gains are increasing in f_i, so intermediate values between 0 and 0.30 would also be "in between" the ones we report.[6]

Implementation in practice. For each context (sample), we assign $S^{marg} \cdot N$ complier households to be marginal. Non-compliers are never assigned to be marginal. We then recalculate the CCT-to-UCT welfare ratio (equation 1) with only one change: if i was assigned to be marginal, its contribution to the numerator of equation (1) changes from $u(c_i^{baseline} + t_i^{cct}) - u(c_i^{baseline})$ to the expression in equation (6). The denominator of the welfare ratio for all households remains unchanged, and so does the numerator for households not assigned to be marginal.

We then repeat the above procedure 1,000 times for each context, each time taking a new random draw of the assignment of marginal status to households. This yields 1,000 calculated welfare ratios.

[6] Another possibility would be to set $f_i = x$, where x is a constant in dollars that is the same to all marginal households, 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 median household consumption), it may be larger than the poorest households' baseline consumption, implying they have negative consumption if they enroll their child, and making it impossible to compute their value of utility.

Results. Our tables report the 1st and 99th percentiles of this distribution (the 10th lowest and 990th highest ratio for each context). Figure A.2 provides the full distribution of the 1,000 calculated ratios for each context. In all cases, we use u as a CRRA utility function with $\rho = 3$.

Table 2 reports the results setting f_i equal to 15% of i's baseline consumption. Table A.4 repeats the exercise doubling f_i to 30%. The "original" column in each table reports the baseline case from Table 1 (using $\rho = 3$). It is the same welfare ratio we would obtain using $f_i = 0$. Thus Tables 2 and A.4 allow us to probe robustness to values of f_i ranging from zero to 30% of baseline consumption. Recall that the CCT gains and thus the welfare ratio is increasing in the value of f_i, so values of f_i between those reported would also yield welfare ratios between those reported (modulo sampling variation from the random assignment).

The main text discusses the interpretation of the results. Overall, the reported ratios are closer to the "original" ratios than to one, implying that incorporating the value of schooling only has a limited effect on our conclusions. This occurs although we use an upper bound on the value of schooling and make assumptions that systematically "tilt the scales" towards increasing the ratio.

Social returns to schooling. Our approach, so far, relies on a revealed preference argument and thus on the *private* costs and values that marginal households place on school enrollment (f_i and D_i). However, a common rationale for CCTs is that the social planner (government) values education more than households. This can occur given social returns to schooling that households do not internalize, or household decision-makers underestimating the (private) returns or undervaluing the utility of future generations.

This section discusses how our approach can incorporate such considerations. It builds on a formulation used by Banerjee et al. (2024) to conceptually discuss these issues. Let $R^{baseline}$ be the CCT-to-UCT welfare ratio abstracting from the value of schooling (e.g., those reported in Table 1) and let $R^{schooling}$ be the ratio that incorporates it (e.g., those reported in Table 2). $R^{schooling} - R^{baseline}$ is thus the value of enrollment for marginals

divided by the utility gains from a UCT.[7]

Assume that the social planner values the net gains from schooling θ times more than households do. The relevant CCT-to-UCT welfare ratio for such social planner is

$$R^{SocialPlanner} = R^{baseline} + \theta(R^{schooling} - R^{baseline}) \tag{7}$$

If the social planner values schooling exactly as much as households, then $\theta = 1$ and $R^{SocialPlanner} = R^{schooling}$. If it places no value on schooling, then $\theta = 0$ and $R^{SocialPlanner} = R^{baseline}$.

In equation (7), θ only multiplies the *marginal* households' value of schooling. When considering a choice between a CCT and UCT, the inframarginals' value of education is irrelevant (by definition, it occurs under both programs).

For any given value of θ, one can use the $R^{schooling}$ and $R^{baseline}$ we report to calculate $R^{SocialPlanner}$. Given there is no obvious choice of θ, Tables 2 and A.4 report instead the value of θ that would make $R^{SocialPlanner} = 1$. We label this θ^*. Intuitively, it provides an answer to "how much more does the social planner must value the net gains from schooling than households do to make it indifferent between a CCT and a UCT?"

The main text discusses the value of θ^* we calculate. We highlight here that in Mexico and Malawi, it is very large (ranging from 12 to 26). In Morocco and Tanzania, they are more modest but also quite large. We interpret this as suggesting that, to make a meaningful change to our qualitative conclusions regarding the welfare consequences of conditionality, the difference between the value the social planner and households put on schooling has to be large.

We emphasize that such large θ values may be reasonable for multiple reasons. First, while the empirical work that we and the most of the literature conducts focuses on enrollment, the incentives produced by CCTs may also increase learning, attendance beyond the CCT required minimums, and ultimately future employment and earnings. To the extent that such benefits are not internalized by households (in theory our welfare model includes the present value of all such benefits - but it assumes households internalize

[7] Both $R^{schooling}$ and $R^{baseline}$ have the same denominator: the aggregate welfare gains from a UCT relative to no transfers.

these), they can increase θ values. Further, various externalities produced by schooling could also be large (social returns to education via more informed voters, lower public costs of incarceration (Heller et al., 2017), etc.)

Lastly, we caveat that our approach involves θ multiplying the *net* gains from schooling. An alternative formulation, which is also feasible in our approach, would involve a social planner that only values the gains from education (D_i) more than households, while placing the same value on its costs (f_i).

Does assigning the households with the lowest or highest consumption as marginal provide an upper bound? Lastly, this section illustrates why our framework does not allow us to make a general statement about whether choosing the households with lowest or highest consumption households as marginals will guarantee an upper bound on CCT-to-UCT welfare gains ratio (and therefore why we choose to report results on simulated draws of marginals instead).

Consider a simple example with three households, A, B, C, enrolling their child in school under the CCT. Their baseline consumption is $c_A^{baseline} < c_B^{baseline} < c_C^{baseline}$. We explore how our welfare ratio changes if we choose A versus C as the marginal. Throughout, we can ignore the denominator in our welfare ratio (welfare gain from the UCT) because it does not depend on who the chosen marginals are. If we choose A as the marginal, the numerator of our welfare ratio is:

$$\begin{aligned} &u(kc_A^{baseline} + t^{CCT}) + D_A - u(c_A^{baseline}) + u(c_B^{baseline} + t_B^{CCT}) - u(C_B^{baseline}) \\ &\quad + u(c_C^{baseline} + t^{CCT}) - u(c_C^{baseline}) \end{aligned} \tag{8}$$

where k equals $1 - f_i$ (e.g., $k = 1 - 0.15 = 0.85$ when we set $f_i = 0.15 \cdot c^{baseline}$). The analogous expression for when C is chosen as the marginal is:

$$\begin{aligned} &u(kc_C^{baseline} + t^{CCT}) + D_C - u(c_C^{baseline}) + u(c_B^{baseline} + t_B^{CCT}) - u(C_B^{baseline}) \\ &\quad + u(c_A^{baseline} + t^{CCT}) - u(c_A^{baseline}) \end{aligned} \tag{9}$$

taking differences between both expressions and rearranging:

$$\begin{aligned} D_A - D_C + [u(c_C^{baseline} + t^{CCT}) - u(kc_C^{baseline} + t^{CCT})] \\ - [u(c_A^{baseline} + t^{CCT}) - u(kc_A^{baseline} + t^{CCT})] \end{aligned} \tag{10}$$

inserting our expression for D_i (the upper bound):

$$\begin{aligned} & [u(c_C^{baseline} + t^{UCT}) - u(kc_C^{baseline} + t^{UCT})] \\ & \quad - [u(c_A^{baseline} + t^{UCT}) - u(kc_A^{baseline} + t^{UCT})] \\ & \quad + [u(c_C^{baseline} + t^{CCT}) - u(kc_C^{baseline} + t^{CCT})] \\ & \quad - [u(c_A^{baseline} + t^{CCT}) - u(kc_A^{baseline} + t^{CCT})] \end{aligned} \tag{11}$$

Each of the four expressions in brackets is positive (since $k < 1$). However, it is unclear if the first expression is larger than the second, and whether the third is larger than the fourth, making the sign ambiguous. For example, compare the first and second terms in brackets. Both involve an increase in consumption. On one hand, the term in the first bracket involves a larger increase in consumption (since $(1-k)c_C^{baseline} > (1-k)c_A^{baseline}$). On the other hand, the term in the second bracket involves a lower increase in consumption levels, but the *utility* gain may be larger since u is concave. The same applies to a comparison between the third and fourth bracket.

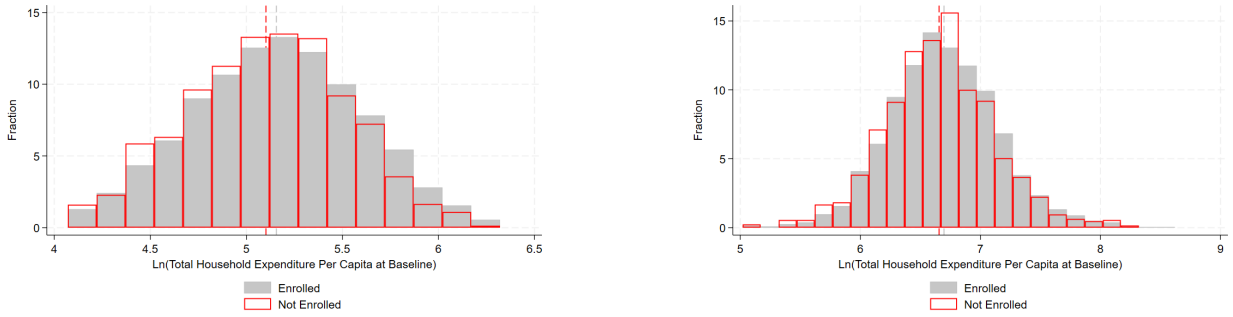
Intuitively, the net gain of schooling that our upper bound assigns to a household depends on both the level of its consumption but also on its marginal utility (the concavity of u). Thus, different combinations of the concavity of u, levels of consumption and transfers, and the differences between the households with higher or lower consumption make it so that it is ambiguous whether choosing the households with lowest or highest consumption yields an upper bound in our calculations.

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 for further discussion.

Figure A.1: Baseline Consumption Distribution by Complier Status - Additional Figures



(a) Mexico - All Grades
Complier mean: 5.16, Non-complier mean: 5.10

(b) Morocco - All Ages
Complier mean: 6.70, Non-complier mean:6.66

This figure replicates Figure 1 using the entire sample for Mexico and Morocco (as opposed to only the 7th grade and ages 11-15 subsample, respectively). Each sub-figure shows, for a given context, histograms of the logarithm of total household expenditure per capita at baseline (i.e., before CCT payments start). Solid gray-colored bars provide the distribution for compliers and the red-bordered bars do so for non-compliers. Vertical lines indicate the mean consumption for each group. The sample only includes eligible children in the CCT arm of their respective experiment. Compliance is defined as being enrolled in school.

Table A.2: Non-Complier vs. Complier Summary Statistics

	# Non-Compliers	# Compliers	Share Non-Compliers	Payment \Avg. Consumption	
				Girls	Boys
	(1)	(2)	(3)	(4)	(5)
Panel A: Malawi (SIHR)					
Age=14	20	67	0.230	0.260	-
15	28	79	0.262	0.241	-
16	38	35	0.521	0.279	-
17	30	30	0.500	0.318	-
18	20	11	0.645	0.235	-
19	7	3	0.700	0.176	-
20	14	3	0.824	0.293	-
Unknown	16	69	0.188	0.195	-
All Ages	173	297	0.368	0.252	-
Panel B: Mexico (Progresa)					
Age=7	3	263	0.011	-	-
8	3	917	0.003	-	-
9	7	1195	0.006	-	-
10	6	1342	0.004	-	-
11	29	1154	0.025	-	-
12	105	1268	0.076	-	-
13	191	1090	0.149	-	-
14	269	809	0.250	-	-
15	350	472	0.426	-	-
16	407	242	0.627	-	-
17	443	120	0.787	-	-
18	367	41	0.899	-	-
All Ages	2180	8913	0.197	-	-
Panel C: Morocco (Tayssir)					
Age=6	11	412	0.026	0.021	0.021
7	5	486	0.010	0.019	0.019
8	9	654	0.014	0.025	0.025
9	24	752	0.031	0.026	0.026
10	34	819	0.040	0.034	0.034
11	101	719	0.123	0.033	0.033
12	174	478	0.267	0.033	0.033
13	283	312	0.476	0.033	0.033
14	349	191	0.646	0.032	0.032
15	299	91	0.767	0.033	0.033
All Ages	1289	4914	0.208	0.030	0.030
Panel D: Nicaragua (RPS)					
Age=7	13	89	0.127	0.056	0.056
8	6	128	0.045	0.056	0.056
9	2	131	0.015	0.056	0.056
10	3	119	0.025	0.056	0.056
11	1	83	0.012	0.056	0.056
12	4	69	0.055	0.056	0.056
13	2	62	0.031	0.056	0.056
All Ages	33	681	0.046	0.056	0.056
Panel E: Tanzania (TASAF)					
Age=7	22	61	0.265	0.169	0.169
8	14	64	0.179	0.169	0.169
9	8	71	0.101	0.169	0.169
10	8	87	0.084	0.169	0.169
11	4	65	0.058	0.169	0.169
12	12	112	0.097	0.169	0.169
13	11	77	0.125	0.169	0.169
14	15	61	0.197	0.169	0.169
15	23	38	0.377	0.169	0.169
All Ages	117	636	0.156	0.169	0.169

The table presents, for each context: (1) the number of eligible non-complier households, (2) the number of complier households, (3) the share of households that are non-compliers, (4) the size of the CCT payment for girls as a share of average per capita household consumption, and (5) the same as (4) for boys. The sample only includes eligible children in the CCT arm of their respective experiment. Compliance is defined as being enrolled in school. Payments depend on child's age in Malawi and Morocco, and in enrollment grade and gender in Mexico. The Malawi program only targeted women. See Section 2 for further information.

Table A.3: CCT-to-UCT Welfare Gains Ratios

	Ratio of CCT Utility Gain to UCT Utility Gain		
Sample	Ln Utility	CRRA ($\rho = 2$)	CRRA ($\rho =$
Panel A: Full Samples			
Malawi	0.853	0.749	0.660
Mexico	1.150	0.954	0.913
Morocco	0.996	0.981	0.932
Nicaragua	1.003	1.011	1.020
Tanzania	0.986	0.970	0.928
Panel B: Subsamples			
Mexico 13-18 year olds	0.941	0.874	0.775
Morocco 11-15 year olds	0.990	0.956	0.841
Tanzania 13-15 year olds	0.974	0.945	0.823

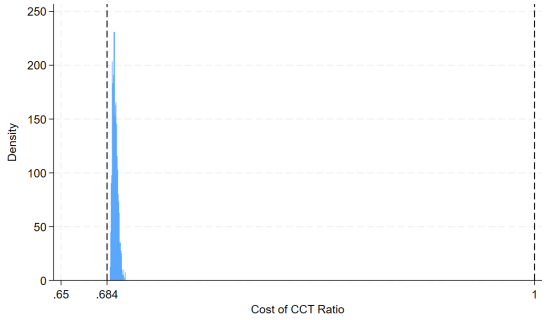
The table presents the CCT-to-UCT welfare gains ratio: the ratio between the welfare gains from the observed CCT and a counterfactual UCT program that uses the same budget to provide payments to all eligible households. This table shows the results for other functional forms for the utility function.

Table A.4: Welfare Ratios Including the Value of Schooling ($f = 0.3 \cdot c^{baseline}$)

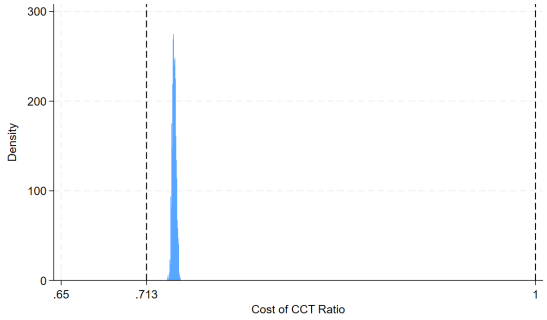
Sample	Original	1st percentile	99th percentile	θ^*
Malawi	0.684	0.691	0.714	10.53
Progresa - 7th grade	0.713	0.760	0.778	4.42
Progresa - All	0.948	0.955	0.957	5.78
Morocco (ages 11-15)	0.908	0.948	0.988	1.15
Morocco - All	0.960	0.979	0.991	1.29
Nicaragua	1.016	1.024	1.031	-
Tanzania	0.950	0.955	0.970	2.50

This table replicates Table 2, but doubling the value of f_i so it equals 30% of baseline consumption. It presents the CCT-to-UCT welfare gains ratio, incorporating the value of schooling based on our bounding approach. The *original* column replicates the baseline case from Table 1 ($\rho = 3$ and no value of schooling considerations: $f_i = 0$). The next columns provide the 1st percentile and 99th percentile of the distribution of simulated CCT-to-UCT welfare ratios based on randomly assigning which compiler households are marginal. θ^* provides a calculation of how many more times the social planner would have to value the net gains of schooling than households do to be indifferent between a CCT and a UCT (a welfare ratio of one). All calculations are based on a CRRA utility with $\rho = 3$.

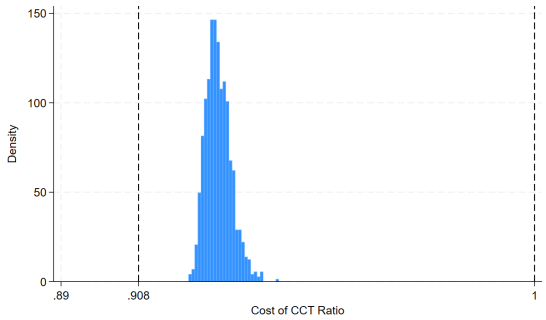
Figure A.2: CCT-to-UCT Welfare Gains Ratios Incorporating the Value of Schooling



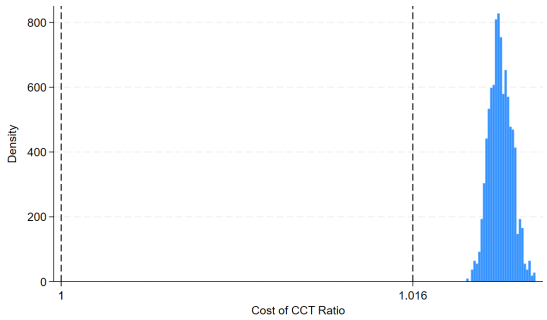
(a) Malawi



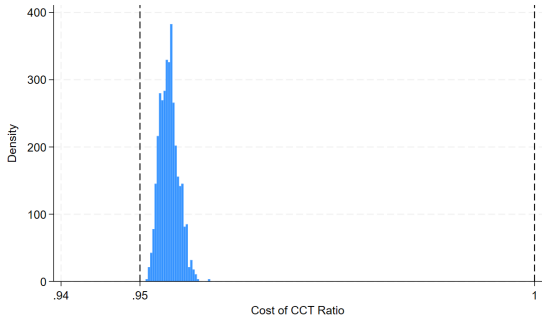
(b) Mexico - 7th Grade



(c) Morocco: 11-15 Year Olds



(d) Nicaragua



(e) Tanzania

Each panel provides a histogram (in blue) depicting the 1,000 CCT-to-UCT welfare gains ratios incorporating the value schooling, calculated by taking 1,000 random draws of the assignment of marginal households. It presents the spread of possible CCT-to-UCT welfare ratios. Each panel also provides a vertical line at the baseline welfare ratio (i.e., the $\rho = 3$ column of Table 1) and also at one (the point where UCTs and CCTs deliver equal welfare). All calculations are based on $\rho = 3$ and setting the cost of schooling f_i equal to 15% of household i's baseline consumption. In all contexts, distribution is closer to the baseline welfare ratio than to one. In Nicaragua, the observed welfare gain is to the right of one since the CCT already delivered higher gains in the baseline analysis.