

Schooling, Income, and HIV Risk (Malawi)

Analytic Design

November 29, 2007

Researchers

Sarah Baird, IRPS Email: sjbaird@ucsd.edu

Craig McIntosh, IRPS Email: ctmcintosh@ucsd.edu

Berk Özler, World Bank (visiting scholar at IRPS): Bozler@worldbank.org

The main aim of our study entitled *Schooling, Income and HIV Risk (SIHR)* is to investigate the impact of cash transfers, both conditional and unconditional, on schooling, sexual behavior and HIV/STD risk among school aged women in Malawi. By school aged women, we mean never married women aged 13-22. The main policy question this study seeks to shed some light on is the following: *What are the health and behavioral benefits from maintaining female enrollment?* This question is potentially crucial for African governments that may be considering creating schooling-based conditional cash transfer (CCT) programs. We plan to implement a multi-arm randomized policy intervention which will give us direct experimental evidence on the following questions:

- *What is the marginal impact of schooling on sexual behavior and HIV/STD risk for young (school-aged) women?*
- *What is the income elasticity of (risky) sexual behavior?*
- *Are there any negative (or positive) spillover effects of increasing schooling and/or income of some on other young women (young men)?*
- *What is the effect of cash transfers on returning to school in Malawi?*

In addition, our study will provide what is to our knowledge the first information from the African context on how CCT programs can best be designed to pull secondary school-age students back into the educational system. The rest of this section further details the evaluation design and elaborates on the research questions, estimation strategy, study setting and survey design.

1. Methodology

1.1 CCT Design

Our research will provide several very important pieces of information for African governments that may be considering creating schooling-based CCT programs. These contributions include:

- providing an experimentally estimated ‘demand curve’ for CCTs, therefore allowing us to estimate the relationship between transfer size and schooling attendance in a very exact way for those who have already dropped out.
- implementing a survey among children who are enrolled at the time of the study to allow us to estimate regressions on the propensity to drop out over the subsequent two years, and therefore providing policymakers with a set of eligibility criteria which would achieve the highest schooling increases from a CCT program.
- estimating very precisely the size of the transfer that would have been necessary to achieve various levels of schooling improvements (under the assumption that the process which re-enrolls children in school is similar to the process that would have caused them not to drop out in the first place).
- conclusively demonstrating the presence (or lack thereof) of health and behavioral benefits from maintaining female enrollment.

With these contributions in mind, we now turn to more details of the CCT/CT program. The proposed CCT/CT program targets two groups of girls: current school girls and those who have recently dropped out of school. We choose to target these two groups separately to ensure that we have a significant number of dropouts in our sample. The reason this is important is because treating current schoolgirls involves giving transfers to some girls who would have stayed in school even without the transfer. A sample that mainly consisted of school girls would be a very expensive way of trying to identify the effect of schooling on HIV risk. Treating dropouts, on the other hand, allows us to focus on a subpopulation whose schooling rates are extremely sensitive to transfers. This is particularly crucial given that the impact of CCTs on HIV incidence will be the product of two treatment effects (impact of CCTs on schooling * impact of schooling on HIV).

Given that we are interested in the impact of the CCT on HIV incidence and sexual behavior, we further restrict our sample to focus on an age group where the majority of girls are likely to be sexually active. We therefore restrict the eligibility for the CCT program to never married girls aged 13-22 who are currently in school or have recently dropped out. Our eligibility requirement differs slightly between the school girls and dropouts—school girls must be eligible to attend grades Standard 7 through Form 4 in 2008, while dropouts must simply be eligible to go back to any class in primary or secondary school. The reason for this difference is that never married dropouts are a somewhat unique group and we want to sample as many of them as possible. Never married school girls, on the other hand, are a much larger sample which allows us to be more specific in our targeting.¹ That said, the majority of the dropouts sampled will be eligible to return to the same grades as the school girls.

Within our target population, there are important differences between those that are eligible to attend primary school and those that are eligible to attend secondary school. In most other areas of the world where CCTs have been implemented (Mexico, Brazil, etc.), they have been used in secondary schooling systems where access to education is both universal and free. Given that primary school is free in Malawi, the administration of our CCT program at the primary school level will look more like a standard CCT program,

¹ The main reason for choosing eligibility based on these grades is so we can offer a cash transfer for one or two years (depending on the grade they initially return to) that will allow the respondent to either complete the primary school leaving exam, the form 2 completion exam or the form 4 completion exam.

except that we propose to add an additional randomized transfer directly to the schoolgirl herself (see next section).

In Malawian secondary schools, on the other hand, fees impose a major financial burden on even a middle-class family, and access to secondary schools is strictly regulated by national entrance exams. In secondary schools therefore, we propose to implement the conditional transfer by directly paying the school fees of treated subjects (as well as adding an additional randomized transfer to the schoolgirl explained below). This process greatly simplifies the payment of the transfers for our team, and has the added benefits of making the conditionality as transparent as possible: only by paying fees directly can we guarantee that the transfers are not diverted into other forms of consumption by parents. This also most closely mimics a real government program under which school fees for some or all secondary school girls would be abolished. In order to avoid conflicts of interest when schools are asked to report on attendance of students for whom they receive transfers, we will have random spot checks of attendance by program staff through the course of the program, with pre-announced fines to both parties for lack of compliance with program requirements. In this way we achieve the removal of financial barriers towards school attendance with a simple and easily administered mechanism.

The design suggested here is in many ways similar to a standard CCT program, which provides transfers to all eligible individuals who attend school. Our program treats three kinds of children: (1) those who would have attended school anyway, (2) those that would have dropped out and do not as a result of the transfer, and (3) those who are enticed back into school by the transfer. Over-sampling the third group allows us to look at several interesting aspects. First, from a research perspective we are interested in finding those populations whose behavior will be strongly affected by the transfer which is most easily identified through group 3. Secondly, given that we are interested in the ability of schooling to protect girls from risky sexual behavior, we are focusing on the comparison where we expect that effect to be most pronounced: the comparison between girls who drop out and those who we can entice to go back to school. Third, with the currently enrolled students, we are unable to observe at an individual level *which individuals* were enticed to change their behavior as a result of the treatment (we would not know who would have stayed in school in the absence of the transfer). With the dropouts, we can identify at the individual level whose behavior has been altered by the program and hence make much more precise statements about targeting and heterogeneity of impacts. Given separate impact estimates on these subgroups, we could easily simulate the impact of a program offered to all girls using only the share of current schoolgirls and dropouts in the population.

Along with the relationship between transfers and schooling and then subsequent HIV infection, an entirely different casual pathway between transfers and HIV infection has been widely cited in the literature. Namely, the need for pocket money has lead many young women to seek partners who provide petty patronage in return for quasi-contractual sex. It is widely believed that this mechanism is primarily to blame for the very high HIV rates among teenage girls in Africa (indeed, often substantially higher than their male age-mates, who have no venues for sexual activity with older cohorts and their correspondingly higher rates of infection). Were we, therefore, to provide a standard conditional *cash* transfer, we would be left with a causal ambiguity: are any observed changes in HIV or sexual behavior coming from schooling, or from the greater financial independence that the transfers themselves engender?

In order to avoid this ambiguity, we divide transfers into three clearly distinguished types. The first is the schooling transfer outlined above. The second is a household level cash transfer that will be given to the guardian of the girl. The third is a straight cash transfer which will be given directly to the girls themselves once per month at a pre-determined central location. Within the CCT treatment group, this additional cash transfer will also be conditional on continued school attendance. The quantity of this cash transfer will be randomized at the individual level within the treatment communities in order to gain identification over the elasticity of sexual activity (and, ultimately, HIV incidence) with respect to conditional cash income. Because this randomization can be conducted at the individual level without major logistical hurdles, it provides substantially more statistical power for the identification of treatment effects. Of particular interest to the research team will be estimating the cross-elasticity with respect to baseline characteristics such as household income, age, and pre-existing sexual relationships.

The degree to which schoolgirls are able to keep control over money given directly to them is a variant of the household income pooling question, on which a large literature exists. The standard approach to such questions, as in McElroy (JHR Vol 25 #4, 1990) or Duncan Thomas (Same volume) is to try to find ‘assigned goods’, such as hair braids or girls’ clothes, which are preferred only by the recipient of the cash. By then studying the increase in the consumption of assigned goods relative to overall household consumption we can infer the share of the transfer that has remained in the girls’ hands. While this assigned share is endogenous, we can instrument for the share actually received by the household and the girl using the total (randomized) transfer to the household, and thereby gain clean identification on how both ‘familial’ and ‘personal’ transfer alter behavior.

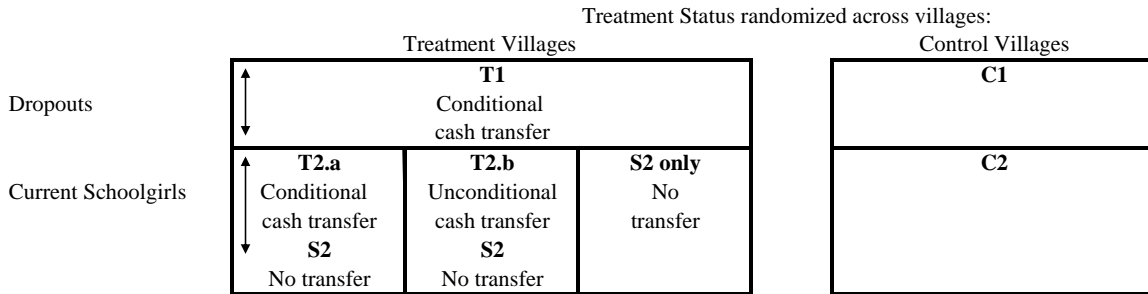
The final phase of the randomization is to extend the cash transfer program to an additional random subset of current schoolgirls that fit the same eligibility criteria. For a randomized subset of these schoolgirls, however, we will offer the transfer unconditionally with no schooling requirement. Since this cash transfer is entirely unconditional it gives a pure elasticity free of the explicit conditionality that will be attached to the transfer made to the ‘eligible’ group. As with the cash transfers to the eligible groups, these transfers will be made monthly at a pre-determined central location. The comparison between the conditional (CCT) and unconditional (CT) groups gives the pure impact of the conditionality, because all other aspects of the research design are held identical.

1.2 Evaluation Design

The specific design structure of the CCT project allows the research team to address the questions listed at the beginning of this section in a very precisely estimated manner. Exploiting both the randomized nature of the intervention and the unique structure of the design, the research team can convincingly address a range of questions that are important from both a policy and research prospective. Unlike an observational study, the randomized nature of this program allows one to infer causality, thus generating concrete policy prescriptions. This section focuses on the econometric methods used to analyze the impact of CCTs/CTs on HIV incidence and other outcomes of interest.

The following schematic illustrates the suggested research design for the Malawi project.

Malawi Research Design:



The first stage of the randomization is at the community level (where, presumably, multiple communities will be sending their children to the same secondary school and possibly even the same primary school). Within each treatment community we locate the universe of eligible girls through a listing exercise. The first group of eligible girls we identify are never married 13-22 year old recent *dropouts* who are eligible to return to primary and secondary school. We denote this core treatment group as T1. This same universe of would-be-eligible girls will be identified in the control communities, denoted by C1. Our second group of eligible girls are never married 13-22 year old *school girls* who are eligible to return to Standard 7-Form 4. Within the treatment communities, we randomly separate the communities into three categories—those where school girls receive conditional transfers (T2.a), those where school girls receive unconditional transfers (T2.b), and finally those where school girls receive no transfer (S2). In addition, within T2.a and T2.b communities we randomize the ‘saturation’ of treatment, meaning that the percentage of current schoolgirls who receive the transfer will also be randomized. These untreated school girls, as well as the untreated school girls in communities where no school girls are being treated form the counterfactual for an experimental estimation of spillover effects. This same universe of would-be-eligible school girls will be identified in the control communities, denoted by C2.

Within the treatment groups, we will randomize an additional cash transfer as described above into a ‘familial’ and ‘personal’ part, thus providing us with two different margins over which to calculate the impact of an additional dollar transfer on school attendance, sexual activity, and HIV incidence.

Given the above evaluation design we can estimate a wide range of treatment effects that are unique to this study, including the income elasticity of sexual behavior. We will now describe in detail how we will estimate the various treatment effects.

1. Impact of Conditional Transfers. Direct comparison of the experimental groups T1 and C1 gives the impact of the average transfer for dropouts. In this case, the impact of the average transfer is defined as the impact of paying school fees plus the average additional cash payment with the average ‘familial’ and ‘personal’ blend. Since we do not expect 100% compliance with the conditionality and we cannot drop non-compliers without introducing bias, we will estimate the Intention to Treat Effect (ITE) of the program, i.e. the impact of *offering* the CCT, as opposed to taking it up. This estimate will be a direct mechanical function of the compliance rate. A similar exercise can be carried out for school girls by comparing T2.a to C2.

2. Impact of Unconditional Transfers. Direct comparison of the experimental groups T2.b and C2 gives the impact of the average unconditional transfer on schoolgirls. Because there is no ‘compliance’ in this treatment, this comparison gives the Treatment Effect on the Treated (TET) of the average quantity transferred with the average ‘familial’ and ‘personal’ blend.

3. Elasticity of Sexual Behavior with respect to Transfers. .

By comparing sexual behavior, STD & HIV rates, and pregnancy rates across the size of the transfer within group T2.b, we can estimate how unconditional transfers alter risk profiles. The idea behind this study is that it is precisely the penury of girls in these treatment arms which makes them susceptible to the ‘sugar daddy’ phenomenon widely blamed for infection among young girls. Hence by increasing their economic self-sufficiency, we hypothesize that they will engage in a reduced profile of risk behaviors. The continuous and random variation in the size of the transfer gives us a very rich way of understanding this crucial behavioral question. Specifically, we can use non-parametric methods to construct an average relationship between changes in sexual behavior ΔS_{iv} and the size of the transfer τ_i in the unconditional transfer group.

A parametric regression specification for estimating the average elasticity in the sample would be:

$$(1) \log(\Delta S_{iv}) = \alpha_0 + \beta X_{iv} + \varepsilon_2 \log(\tau_{iv}) + \nu_{iv} \quad \forall v, i \in T2.b$$

ε_2 provides the elasticity for T2.b.

The school fee payments cannot be varied in an experimental fashion within the treatment group T1, but the randomized cash payments in excess of the school fees provide an individually-randomized source of variation. Because these additional payments are explicitly conditional upon continued schooling, the relationship between ΔS_{iv} and τ_{iv} in this group should be driven both by the elasticity estimated above and by the higher-powered incentives for good behavior which are provided to girls receiving large conditional transfers. We therefore hypothesize that the slope of the elasticity $\varepsilon_1 \left(\frac{dS_i}{d\tau_i} * \frac{\tau_i}{S_i} \right)$, calculated through the counterpart to 1 but estimated on T1) will be larger than ε_2 .

4. Share of Income Retained by Girls in the Study Group.

Household heads will be aware of the size of the transfers made directly to the girls, and hence either altruism on the part of the girls or dictatorial powers on the part of the household head may result in some or all of the transfer being co-opted by other household members. This question is of central theoretical interest, because it provides a direct test of the extent of household income pooling. It is also central to the policy question at hand; to the extent that ‘sugar daddies’ are able to provide consumption goods directly to the girls themselves, we may see that transfers have a negligible effect on behavior as the share actually received by the girls goes to zero.

The income pooling literature has generally approached this question by identifying ‘assigned goods’; items which are solely consumed by specific members of the household (Thomas, JHR Vol 25 #4, 1990). The claim is then usually made that certain sources of income are exogenous and unanticipated, and the share of consumption increases for general household goods versus the assigned goods allow the researcher to calculate the share of the transfer which was controlled by a specific individual. The ‘shocks’ used to identify these exercises are often of questionable exogeneity (McElroy). Because our

income shocks are randomly assigned and unanticipated, we have the ideal empirical structure for testing this central question.

5. Relationship between Transfers in Girls' Pockets and Impacts.

We can use the 'assignable good' methodology in 4 to estimate the value of the transfer which a girl directly retains herself. Through a matching exercise we can calculate individual-level impacts of the treatment. While both of these quantities are potentially endogenous to household characteristics, we are interested in understanding how these two quantities are related. We will not be able to determine, for example, that a girl experienced higher impacts *because* she retained a higher share of the transfer, but we are still interested in knowing whether this correlation exists.

We can also look at the subset of study units found to have retained virtually all of the transfer themselves. We can use other characteristics to predict a propensity score, and compare these high-retainers in the treatment to observationally equivalent control units. This comparison lets us measure the elasticity of sexual behavior to transfers in an environment where the research subjects are not transferring income to other household members.

6. Impact of Conditional Schooling Transfers Only:

We get at the sole impact of the schooling transfers by making the same average familial and personal transfers to schoolgirls in the two treatment arms T2a and T2b². Thus, the only difference between groups T2a and T2b is the conditionality itself, and so the impact of the conditionality is given by the simple difference T2a-T2b.

Along with estimation of the treatment effects detailed in (1)-(6), the research design allows for the estimation of a variety of spillover effects. Spillover effects are of particular interest in this experiment for several reasons. First, the basic idea that girls may be able to avoid HIV infection by staying in school raises the following question: what are the men whom they are not sleeping with doing instead? A natural concern is that the CCT merely 'diverts' infection rather than preventing it, thus leading to a corresponding increase in sexual activity and HIV infection in non-treated age-mates within the same village. The policy conclusions of this study would radically differ depending on whether a given ITE among the eligible is a result of HIV diversion or HIV prevention. Thus we must closely track the infection rates in the subpopulations that are identified to be the closest sexual substitutes for the treated girls. This possible negative spillover is one purpose for the intensive study of the control units within treatment communities, S2.

There is an additional spillover which, although not related to HIV, is central to the policy questions surrounding the expansion of access to education in Africa. This effect is the educational spillover effect imposed on girls who would have attended school anyway by the entrance of a wave of new classmates. References to the deterioration of the quality of primary education in Africa as a result of universal primary education are legion, and this experiment gives us a controlled setting in which to quantify these effects. Presumably, this question will be of central importance to African policymakers as they consider how best to expand access to secondary schooling across the continent.

² The average school fee payment made to the schools in the CCT group will be given to households in the CT group, so that the means are truly identical.

In what follows, we illustrate how we can use our experimental structure to estimate several forms of spillovers in experimental and quasi-experimental fashions.

7. Experimental Average Spillover Effect: Direct comparison of S2 to C2 gives us the average joint spillover effect of the three treatment arms

8. Network Spillovers within Communities: As a part of the baseline survey, we will conduct a detailed network analysis, asking standard social network questions about whom each individual in S2 spends free time with, who they go to for advice, etc. This information will allow us to identify the social network of each individual in S2 at the beginning of the experiment. The randomized nature of the experiment will generate variation in the share of each girl's social network that is treated by each of the three treatment arms, and therefore gives us a measure of the intensity of the spillover effect that we expect to observe. The intensity of the spillover comes both from the share of a girls social network that newly attends school (a'_{iv}), as well as the average transfer made to girls within her social network (τ'_{iv}). The network-based counterpart to (1) that allows us to estimate spillovers at the individual level, then, is

$$(2) \log(\Delta S_{iv}) = \alpha_0 + \beta X_{iv} + \varepsilon_a \log(a'_{iv}) + \varepsilon_s \log(\tau'_{iv}) + \nu_{iv} \quad \forall v, i \in S2$$

This regression allows us to estimate a crucial and fascinating set of relationships; how does the sexual behavior of Malawian girls respond when the schooling of their compatriots rises, when family income changes, and when their friends have more spending money in hand? Much of the interest in this regression lies in the theoretical ambiguity of the results; one could write down a model that would predict that any of these spillover effects could move in either a positive or negative direction. The only way to answer these fundamental questions is with solid experimental evidence.

9. Schooling Spillover Effects: The unit at which we expect education spillovers to be most concentrated is at the classroom level. The impact of having additional students placed into your school is likely to be muted unless they are placed into your classroom and compete directly for resources and teacher attention. In secondary school, we can calculate this classroom-level spillover effect in a very concrete way as follows: first, we will predict for each treatment and control individual the class into which they *would have gone* if they attended school. By adding this hypothetical enrollment to the real class size we get the 'potential enrollment' that would ensue if the government were to enact free secondary education without changing the eligibility criteria (and ignoring any potential incentive effect that this change might have). Taking the class as the unit of analysis, we then take advantage of the fact that our experiment creates random variation in the hypothetical enrollment, and use this as the treatment variable. Our outcome of interest will be the secondary school JCE scores (exam taken upon completion of Form 2) among untreated students, and the treatment variables will consist of the percentage of girls in each classroom who were given each of the treatments.

If non-compliance rates are found to be high, or if in the end many of our 'control units in C1 end up attending secondary school despite the fact that we predicted they would not, then we will run the educational spillovers regression using the treatment as an instrument for attendance. If we find that there is very strong heterogeneity in test scores

across the number of hypothetical enrollees, then we can include fixed effects for this attribute and estimate ‘within’ marginal effects. By explaining their test scores with the number (and/or share) of the hypothetical enrollees, we gain experimental identification on perhaps the most pertinent educational spillover question for CCTs in Africa: how much do we drive down the performance of existing school goers by adding one additional student to their class? Although primary school is already free, we can undertake a similar analysis with those finishing primary school by looking at scores on the primary school leaving exam (PSLC).

Because the ‘saturation’ of treatment among schoolgirls is randomized, we have a very simple way to estimate how the intensity of these spillover effects changes as the intensity of treatment changes. We can compare communities where only dropouts return to school to those where an increasing fraction of the current schoolgirls are treated in order to understand how giving transfer to some girls to keep them in school alters the behavior of girls in the same classes who are not getting these transfers.

10. Spillover Effects on Boys:

If the CCT works by decreasing contact with high-prevalence older men and leaving girls free to date boys their own age, it is possible that our treatment could result in an *increase* in HIV infection among younger male classmates. While we do not have the resources to track a large cohort of boys throughout the study, the experimental nature of the design makes a baseline survey of boys unnecessary. We therefore propose that at the end of the two-year study, we conduct a wide-spread testing campaign for HIV and STDs among boys in the study schools. Along with a social network survey, this single-wave study will allow us to ask the following:

- Do boys in treatment schools have different levels of infection than boys in control schools?
- Do boys in classrooms with a large share of treated girls study have higher infection rates?
- Do boys with a large share of treated girls in their social networks have higher infection rates?

Because several recent studies of female sexual behavior in Africa have ignored the obvious spillover effects that alteration of younger girls should have on younger boys (Dupas, Thornton), we feel that investigating this possibility is essential in order to understand the general equilibrium effects of our schooling intervention.

The unique design of this CCT coupled with the planned survey instruments and HIV testing will allow us to address the wide range of treatment impacts listed above, as well as research a variety of other questions not directly related to the experiment. With the aid of highly skilled researchers, and using this innovative randomized survey design along with HIV and STD testing, our study endeavors to be the first randomized CCT program that not only investigates the direct effect of the conditional transfer on schooling, but also investigates the effect of both conditional and unconditional transfers on the sexual behavior and HIV status of both the treated and non-treated individuals. To our knowledge, no CCT evaluation such as this currently exists in sub-Saharan Africa.

2. Survey Setting

Malawi was chosen to be the setting for this CCT program not only because of its high HIV prevalence rates, but also to complement a currently ongoing observational study entitled Marriage Transitions and HIV/AIDS in Malawi (MTHM). One of the principal investigators in this study (Özler) is also a principle investigator in MTHM and has field research experience in Malawi. Kathleen Beagle and Michelle Poulin, also primary investigators for MTHM, will assist us in with logistics in implementing our program. Moreover, we can draw from the survey instruments and field resources developed under MTHM, thus reducing our costs and facilitating our field work. Finally, the data collected under the MTHM observational study complements our CCT program creating an extremely rich set of data from which to draw.

Since this project entails several in-depth components which require well-managed field staff as well as the logistics of running a CCT program, the research team opted to focus the project within one district. This focus will both reduce project costs (lower fixed costs of office infrastructure and transport) and increase data quality through more careful supervision. Within Malawi, Zomba district was chosen as the site for this study for several reasons. The first reason was that it had a large enough population within a small enough geographic size rendering field work logistics easier and transport costs lower. Zomba is a highly populated district, but distances from the district capital (Zomba Town) are relatively small. The second criterion was that it had high enough dropout rates and HIV/STD rates to make our study meaningful. According to MDHS (2004), HIV/AIDS rates of women aged 15-49 in Zomba are the highest in the country at 24.6% and dropout rates are also high. The final criterion in choosing Zomba was that it was not sampled to be part of the UNICEF cash transfer program until after 2009. The UNICEF cash transfer is an unconditional cash transfer targeted at poor households. The current plan is that it will be phased into all districts within Malawi over the next nine years. To avoid both confounding our results and getting in each others way we determined that it was best to operate in a district that UNICEF will not reach until after our study is complete.

3. Survey Design

The CCT entails sampling approximately 4000 young girls from 204 EAs in Zomba district of Malawi and following these individuals for two years. Enumerations areas (EAs) were selected from the universe of EAs produced by the National Statistics Office of Malawi from the 1998 Census. The sample of EAs was stratified by distance to the nearest township or trading centre. Of the 550 EAs in Zomba 50 are in Zomba town and an additional 30 are classified as urban (township or trading center). The remaining 470 are rural (population areas, or PAs). A random sample of 29 EAs in Zomba town, 9 trading centers in Zomba rural, and 139 randomly selected population areas within 16 kilometers of Zomba town, and 27 more randomly selected EAs with a greater distance than 16 kilometers to Zomba town

After selecting sample villages, all households will be listed in the 204 sample EAs using a short two-step listing form. The first form, Form A, asks, for each household, ‘do you have any never married girls in this household who are between the ages of 13 and 22?’ This form allows us to quickly identify households that have members that fit into our sampling frame, thus significantly reducing the costs of listing. If we receive a yes on Form A, then we move to Form B which gets a list of the members of this household. For individuals in these households we ask the following additional questions:

- Name

- Age
- Marital status
- Current schooling status
- If currently in school, level attending in 2007 for current school girls
- If currently not in school, highest grade completed
- If currently not in school, the last year during which they were in school

This information collected in Form B gives us a census of all girls within the target age range, and allows us to categorize them into two groups:

1. Eligible dropouts, who have been out of school for 2 years or fewer
2. Eligible schoolgirls, those in our age group that are still in school

These two groups comprise the basis of our sample frame. In each EA, we expect to find around 80 girls who fit our sampling criteria. Of these, we expect to see 10-15 dropouts, with about two thirds of the dropouts from primary school and one third from secondary school. We expect to have 7 eligible primary dropouts and 4 eligible secondary dropouts in each EA.³ We will sample all eligible dropouts. In addition, we will sample a random subset of 15-20 school girls in the sample. We anticipate that we will treat all eligible primary and secondary school dropouts in treatment villages and that we will select a random sample of current schoolgirls to treat, some with conditional and some with unconditional transfers.

CCT treatments will be divided into three components: the amount paid to the household, which will be fixed for all recipients in a given community; a randomized transfer paid directly to the girls with notification to the head of household on the existence and size of this payment⁴; and the payment of school fees for secondary school girls paid directly to the school itself.

CT transfers are designed to exactly mirror the conditional cash transfer payments, so that the experimental difference between CT and CCT individuals is driven by nothing but the conditionality itself. The CT payments will be randomized in the same fashion, except that a varying percentage of schoolgirls in treatment communities will receive no cash payment at all. The purpose of this zero density is to allow us to get a clean measure of spillover effects by studying a random subset of ineligible dropouts and of current schoolgirls who are themselves receiving no transfers. This spillover can be measured on average by comparing these zero-transfer girls in the treatment with their counterpart control units, or by tracking the intensity of treatment within the social networks among the zero-transfer girls to see whether individual outcomes are a function of the size of transfers to their friends

A core LSMS-type household questionnaire will be administered annually to both the CCT and CT recipients. This survey will include information on household characteristics, sexual behavior, and social networks. Subsequent to this survey, schools and communities will be notified of treatment status, and then in treatment villages we will prepare a detailed informational sheet for each household detailing the quantity of transfers that each child in their household is to receive. As mentioned previously, the CCT

³ These numbers are based on data from the MTHM listing exercise and IHS-2 data.

⁴ This randomized transfer will have a density at zero, so that we have a reasonable share of observations wherein the girl herself receives nothing above and beyond what her household receives.

component will take place monthly and will be delivered at a central EA location within the treatment EAs.

The project will also include a number of additional features: a mid-year interview using a medium-length survey instrument focusing on partnering behaviors (rather than collected in a 12-month retrospective), testing for HIV and, if money permits, STDs of the respondents, a detailed survey and testing of boys at the completion of the project, and a smaller scale school based survey at baseline among children who are enrolled at the beginning of the study. The baseline survey round will be conducted during October-January 2007. Follow-up rounds will occur one year later in 2008, and then again in 2009.

The core sample for this study consists of 4000 young women that fit our sampling criteria. Approximately 20 core respondents will be selected in each village out of the two groups identified by the listing exercise. Each core respondent will be administered a core household questionnaire at baseline. The core questionnaire, to be administered each wave, resembles a traditional Living Standards Measurement Study (LSMS) questionnaire. However, the individual-specific data will be more in-depth in areas specific to the study objectives. Out of these 4000 young women, we predict approximately 2000 will receive conditional cash transfers, of which half will be in treatment villages. Thus, we will offer the CCT to approximately 1000 women.

The main components of the survey design are discussed in more detail below: core sample, annual household survey, in-depth partnership survey, CCT program, and HIV/STD testing.

4. SIHR Core Sample

We drew our core sample of young women from villages in Zomba district. A sample size of 4000 women was chosen to enable the research team to identify treatment effects on the outcome variables of interest. Across variables of interest, power calculations indicate that the proposed sample size of 4000 individuals (in 204 enumeration areas) will allow us to detect moderate treatment effects as significantly different than zero at high levels of confidence (95 percent) with considerable power (80 percent). Among these, the variable for which it is hardest to detect is changes in HIV status, due to the fact that the prevalence rates are low. Statistical power is a function of the expected effect size, the level of significance desired, and sample size. In Malawi, according to MDHS (2004), HIV prevalence among 15-22 year-old females (the age group we are studying) is 9.0% throughout Malawi, while in Zomba the rate is much higher at 17.3%. Within this age group there is a lot of variation across ages with a rate of 10.2% for those ages 15-19 as opposed to a rate of 21.1% for 20-24 year olds. Given these highly different rates across ages, we expect a rapid transition in HIV status from baseline to follow-up, particularly among the younger cohort.

We expect the take-up rate of the program to be close to 100%, and we expect to see a three percentage point change in incidence as a result of schooling, and so we must be able to detect a binary treatment effect of three percentage points in a sample of 4000 in 204 clusters with a 50% treatment rate. Power calculations, which take into account the fact that our sample is clustered at the EA level, show that our sample of 4000 individuals in 204 EAs will allow us to detect moderate treatment effects have a 80% chance of detecting a difference in HIV prevalence between treatment and control EAs (from a mean

of 0.08 in the treatment to a mean of 0.11 in the control) at a significance level of 5%,⁵ a level of power usually recognized by the research community to be sufficient (Raudenbush et al., 2004).⁶

Although our ability to detect such a large treatment effect with HIV may be somewhat optimistic, we have enough power to detect even small treatment effects with all our other outcomes of interest. Looking at HSV-2, for example, which is estimated to be around 50% among sexually active women in East Africa (and we estimate will be around 20% for our age group at baseline), power calculations show that our sample of 3000 CCT individuals in 150 EAs will have a 80% chance of detecting a difference in HSV-2 prevalence between treatment and control EAs (from a mean of 0.20 in the treatment to a mean of 0.25 in the control) at a significance level of 5%.

5. Annual SIHR Household Survey

The annual SIHR Household Survey consists of a multi-topic questionnaire to be administered to the households in which the selected sample respondents reside. Although it is described as a household questionnaire, the primary goal of the SIHR Household Survey is to collect detailed information from the individual respondent selected for the survey. Thus, the multi-topic questionnaire will consist of detailed information collected from the individual respondent as well as other household members where relevant. From individual respondents, the questionnaire will collect data on: parental background, inequalities in existing and past relationships, transfers and gifts between partners, networks, risk perception, desired fertility, socio-economic and other shocks experienced by the household during the respondent's childhood and expectations (e.g. regarding risk, but also fatalism). Household-level indicators will include traditional socio-economic indicator variables (such as household income, expenditure patterns, and labor allocation), intra-household allocation of resources (to the extent they can be measured), exogenous socio-economic shocks experienced by the household between rounds of data collection. In the expenditure section we will make sure to ask about expenditures on goods that we expect to be influenced by the cash transfer program, particularly goods we expect the girls to purchase with their share of the transfer, or 'assigned goods'. Although recognizing the limitations of such a question, we will also explicitly ask what the transfer was used for and how much of it was controlled by the household versus the girl. Community characteristics will also be collected in a separate short community questionnaire.

6. In-depth partnership interviews (IPIs)⁷

Following the structure of MTHM, we will also utilize an in-depth partnership interview (IPI) six months after baseline and again a year later. This survey focuses on partnering behaviors and related sexual activities over a shorter period, rather than relying on a 12-month retrospective questionnaire.

The IPI interviews will collect detailed information on the sexual behavior of the respondent for a short (2-4 weeks) recall period. Data collected during these interviews

⁵ HIV prevalence rates in the sample are based on calculations by the research team and the power is robust to reasonable changes in those rates.

⁷ The use of IPIs is dependent on the necessary funding.

will include: number of partners, frequency and types of sexual intercourse, presence of other STIs, information on partners, type of partnership (casual, boyfriend, spouse, with sex worker), and transactions between partners. These interviews will be conducted by a 'peer' of the respondent to elicit more reliable information on these sensitive topics.

7. Biomarker data for HIV and HSV-2

All core respondents will be tested for HIV twice during the study period. In addition, a random subset of male classmates will be tested at the conclusion of studies to look at potential spillover effects (as detailed above). The protocol for this program will follow the design and implementation successfully undertaken in the two most recent rounds of the MDCIP: MDICP 3 (2004) and MDICP 4 (2006), as well as in MTHM. The ORASURE saliva test was used, with positive results confirmed through Western Blot on the same specimen. Response rates to the HIV test were over 90% in the MDICP.

We will also pilot testing of HSV-2 and if funding permit include HSV-2 testing in our study. Including HSV-2 testing in our study is desirable for two reasons. First, with much higher rates than HIV across East-Africa (as high as 50% among women according to Oster (200?)), it will be easier to detect treatment effects with respect to HSV-2, as opposed to HIV. Second, HSV-2 is a good indicator of risky sexual behavior and future HIV risk, and as such is an important thing to prevent in its own right.

8. Additional design features of the SIHR survey

Along with the household survey, in-depth partnership interviews, and HIV/STD testing, there are some additional unique features of our survey design. Given that we are focusing our conditional transfer on drop-outs, which as mentioned above would never be a viable government policy, we will also use a school based survey to attempt to capture the determinants of dropping out to help better inform policy. To do this we will select a random sub-sample of the schools attended by individuals in our control villages. In these schools we will implement a relatively short survey at the beginning of the study among children who are enrolled. We will then observe which of these individuals drop out over the subsequent two years. This information will allow us to estimate regressions on the propensity to drop out during this time. Since the government can not directly target drop outs, these regression estimates will provide policymakers with a set of eligibility criteria that would achieve the highest schooling increases from a CCT program. This sort of information could only be gathered through a CCT program like ours.

In addition, although the focus of our study and survey effort is on females, we will also select a random sample of male counterparts in the study schools to survey and test for HIV and STDs at the conclusion of the two year study. As mentioned above, if the CCT works by decreasing contact with high-prevalence older men and leaving girls free to date boys their own age, it is possible that our treatment could result in an *increase* in HIV infection among younger male classmates. While we do not have the resources to track a large cohort of boys throughout the study, the experimental nature of the design makes a baseline survey of boys unnecessary in this case, allowing us to solely survey and test them at the conclusion of the two year study. We know of no other study that has considered the spillover effects that alteration of younger girls sexual behavior should have on younger boy.

M:\ps\HDPS Website\HIVAIDS subtopic\Ozler - Malawi\Malawi.SIHR_CCT_Analytic_Design_071129.Ozler.doc