MATHEMATICA
Policy Research, Inc.

MEMORANDUM

P.O. Box 2393
Princeton, NJ 08543-2393
Telephone (609) 799-3535
Fax (609) 799-0005
www.mathematica-mpr.com

TO: MCC; MCA-Armenia

FROM: Ken Fortson; Anu Rangarajan **DATE:** 1/7/2010

SUBJECT: Irrigation Infrastructure Evaluation Design

The Millennium Challenge Account with Armenia (MCA-Armenia) aims to increase household income and reduce poverty in rural Armenia through improved performance of the country's agricultural sector. As part of this, Armenia is investing in rehabilitating the irrigation infrastructure serving rural communities. By improving living standards among rural residents, these investments can in turn lead to future economic growth in rural areas and throughout the country as a whole.

MCC has commissioned a rigorous impact evaluation to separately examine each of the main components of the MCA-Armenia program. This memo describes our design for evaluating the irrigation infrastructure activity. MCA's irrigation rehabilitation efforts cover several different types of irrigation infrastructure, including main canals, the Ararat Valley drainage system, pumping stations, gravity schemes, and tertiary canals. However, for most of these types of infrastructure, only a handful of projects will be implemented, too few to support a rigorous evaluation. Hence, it was decided to focus the evaluation effort on the tertiary canal rehabilitation efforts because there will be a sufficient number of tertiary canals to yield precise impact estimates. In addition, the primary result for all of the different types of infrastructure (other than drainage) is to increase water availability and reliability. Understanding the impact of more water and more reliable water on farm productivity by analyzing the tertiary canals will also inform us about the likely impact of the other types of irrigation infrastructure to the extent that their measurable effects on water availability and reliability are similar.

Tertiary canals route irrigation water from larger irrigation infrastructure such as main canals or reservoirs to the farmers' fields. Nearly all of the canals under consideration for rehabilitation have deteriorated considerably over the past 20 years, and water losses are substantial. MCA-Armenia estimates that only 25-40 percent of irrigation water actually reaches the fields in most of these villages.

As of October 2009, tertiary canals serving 110 communities have been selected for rehabilitation (not counting four pilot projects that will be excluded from the evaluation).[1] Work on a small batch of these communities is scheduled to begin in early 2010. The remaining communities have been divided by region into three packages. These packages will be bid upon separately by prospective contractors, with work scheduled to begin in spring 2010. MCA-Armenia provides most of the financing for the rehabilitated canals, but villages are responsible for paying a small portion of the construction costs; if they are unable to come up with the co-funding, the canal will not be rehabilitated. This co-funding arrangement is designed in large part so that villages feel ownership over the canals and are more likely to maintain them over the longer term.

Although a random assignment design is considered the most rigorous evaluation approach, randomly selecting which canals would be rehabilitated was deemed infeasible. Instead, we have developed a *comparison group* design. Under this approach, tertiary canals for which rehabilitation is planned will be matched to other canals sharing similar geography, pre-rehabilitation conditions, and where similar crops are grown. Examining how outcomes change for farmers in the comparison group, whose canals were not rehabilitated, will inform us about how those outcomes would have changed in the absence of the rehabilitation efforts.

The key research questions guiding our design of the evaluation for tertiary canal rehabilitation are:

- Did the program affect the quantity and reliability of irrigation water provided to Armenian farmers?
- Did the program affect agricultural productivity?
- Did the program improve household well-being for farmers served by those canals, especially income and poverty?

The remainder of the memo is structured as follows. Section I describes the matched comparison group design and how the matches will be identified. Section II presents the data sources that will be used. Section III outlines the analysis plan. We conclude in Section IV with a summary of the anticipated timeline.

---

[1] Because the pilot canals were rehabilitated well before the other canals and potentially could have been utilized in the previous agricultural season, we cannot obtain the informative pre-intervention baseline data that would be necessary for these canals to be included in the evaluation.

I. COMPARISON GROUP DESIGN

The comparison group design focuses on comparing communities served by rehabilitated tertiary canals (hereafter "tertiary canal communities") to similar communities whose infrastructure was not rehabilitated (hereafter "matched comparison communities").[2] We would estimate the impacts of the program by comparing the post-rehabilitation outcomes for these two sets of communities. Crucially, the analysis will compare how the outcomes have changed relative to the same outcomes measured before the rehabilitation. This approach, which estimates program impacts as the "difference in differences" for the treatment and comparison group, is stronger than simply comparing post-rehabilitation outcomes because it allows us to adjust for pre-existing differences in the two groups. Still, for this approach to be credible, we must be able to identify communities that are very similar on observable characteristics to serve as the comparison group.

Identifying Matched Comparison Communities

For a given tertiary canal community, we want to identify a comparison community (or communities) that, prior to the rehabilitation, is very similar on the characteristics that could be expected to affect the key outcomes: agricultural production and irrigation conditions. Matched comparison groups are often chosen using statistical methods that, for each tertiary canal, would find as close of a match as possible on the many community characteristics that could affect these outcomes. However, a statistical matching approach would require a data file containing information such as main crops grown, number of farmers, irrigation sources, etc. for all of the communities in the regions where irrigation projects are planned as well as all communities that could serve as possible comparison communities. Such a data file does not exist and would require considerable effort to create.

Instead, MCA staff who are knowledgeable about the agricultural conditions in these communities worked with Water User Association (WUA) directors to identify suitable comparison communities for each tertiary canal. The comparison communities were selected with a focus on the following three criteria:

---

[2] Some communities have more than one canal, and the rehabilitated canal serves only a subset of farmers in the village. In these cases, the survey and analysis will focus on farmers served by the rehabilitated canal. In the subsequent discussion, we focus on the illustrative example of a single canal per community for expositional simplicity.

1. Be in the same geographic area and served by the same Water Users Association (WUA);
2. Have similar pre-rehabilitation irrigation conditions as the communities that will benefit from the rehabilitation project, such as similar water losses and source of irrigation water; and
3. Grow similar crops.

A given tertiary canal community could potentially be matched to multiple comparison communities if more than one is a good match on the above criteria. We will include all such matches in the survey so as to maximize the sample size and, hence, the statistical precision. In other cases, multiple tertiary canal communities may share a set of comparison villages if they have similar characteristics. In both of these cases, we will use the baseline data to construct weights based on how closely-matched the villages are.

In addition, to get a second assessment of the comparability of these matches, the survey team will verify the suitability of each matched comparison community when they go into the field to conduct farmer surveys. They will collect independent assessments of the three main criteria listed above from the village mayors as part of a village mayor survey, and they will also consider other community characteristics that may indicate that, for a variety of reasons, the planned comparison community does not provide a compelling match. MCA-Armenia has also already identified five tertiary canal communities that did not have a suitable comparison community, and these five are excluded from the survey.

II. DATA

Tertiary Canal Survey

The primary data source will be a new household survey tailored to this impact evaluation, the Tertiary Canal Survey (TCS). The TCS is modeled closely after the survey used for the Water-to-Market impact evaluation, the Farming Practices Survey (FPS), and will be fielded by the same survey team led by AREG. As with the FPS, the key outcomes of interest from the TCS include crops cultivated, crop production, agricultural profit, household income, and poverty. The TCS also features questions about reliability and quality of irrigation water. We will conduct two rounds of the TCS. The baseline TCS will be fielded beginning in December 2009 and finish in February 2010. The final round was planned to begin in March 2011 and finish in June 2011; however, as many of the canals will likely not be completed in time for the 2010 agricultural season, a third round of the survey will quite possibly be added to the data collection.

The sample frame for the TCS comprises the farming households served by the rehabilitated tertiary canals and the matched comparison group. The survey team will work with village

mayors to identify the farmers served by each tertiary canal, and then the mayors will help them arrange interviews with a subset of them. In tertiary canal communities, 15 farmers will be interviewed. Twenty farmers will be interviewed in each comparison community. The survey team will strive to interview farmers who grow similar crops and have similar land sizes as the associated treatment group farmers, but the larger number of respondents will allow some cushion in case a few of the comparison group farmers are actually dissimilar to the associated treatment group farmers. The total sample size and number of communities will include approximately 3,500 farming households across 200 communities. Just over half of these are in the tertiary canal group, and the remaining communities are in the comparison group. Every tertiary canal community is matched to at least one comparison community, and a small handful of comparison communities are matched more than once.

Some of the tertiary canals currently planned may ultimately not be implemented if, for example, the community is unable to pay their portion of the funding required. Any that are known to drop out prior to data collection will be excluded from the TCS. Construction delays may also mean that canals are not completed in time to be included in the analysis. Any dropped after data collection will be excluded from the analysis, along with their matched comparison group. However, although the analysis can be adjusted to account for these dropped communities, the smaller sample size will reduce the statistical precision of the impact estimates. Therefore, although the study may ultimately have smaller sample sizes than those described above, we estimate that, even with conservative assumptions about community attrition, the minimum detectable impact on household poverty rates is approximately 5.5 percentage points. If the true impact of the rehabilitation on household poverty is 5.5 percentage points or more, we are likely to find a statistically significant impact.

The ultimate goal of the MCA-Armenia program is to increase household income in rural Armenia, and hence, these outcomes are an important focus of the TCS instrument. Because a full accounting of all sources of household income would require far longer to administer than the allotted time for each interview, the survey concentrates on sources of income that are most directly affected by irrigation rehabilitation, specifically, income from agricultural production and from employment by the farmer and his or her immediate family. We can also use the average sale price of specific crops for other farmers in the same geographic area to monetize crops that are consumed by the household or bartered. Additionally, the TCS asks for estimates of expenditures on key categories of consumption, and for income from other sources. Table 1 summarizes the key final outcomes that can be examined using the TCS data. We note that some of these outcomes, such as employment income or income from pensions and other sources, are not outcomes the program is intended to directly affect. However, we include them because they are important components of the household's total income and, hence, it is necessary to have estimates of them.

Table 1. Final Outcome Measures: Survey Data

Final Outcome Measures	Time Frame
Agricultural Productivity Total amount of specific crops grown; amount of crops grown per square meter; total value of all crops cultivated.	Last Agricultural Season
Livestock. Number of cows, pigs, and sheep owned.	As of Survey Date
Revenue from Agricultural Production. Value of crops sold; total value of all crops (including those sold, bartered, or consumed).	Last Agricultural Season
Agricultural Costs. Expenditures on fertilizers, pesticides, irrigation water, hired labor, rented equipment, and taxes.	Last Agricultural Season
Profit from Agricultural Production. Revenues minus costs—the income from agricultural activities.	Last Agricultural Season
Income from Employment. Whether household head, spouse, and any grown children were employed (besides work on the family farm); total earnings from employment.	Last Year
Income from Pensions, Remittances, or Social Programs.	Last Year
Total Household Income. Agricultural profits plus income from employment or other sources.	Last Year
Household Consumption. Expenditure on purchased food, health care, housing products, utilities, and transportation; cost of purchased goods plus value of crops consumed by the household.	Typical Month/Last Year

Although examining impacts on the key outcomes shown in Table 1 is valuable, the impact estimates may not capture the full effect of the irrigation rehabilitation on household well-being because of the relatively short follow-up period, most likely only covering one agricultural season after rehabilitation. The full effects might not be observed in this timeframe for two reasons. First, some farmers may not immediately adapt to the improved irrigation conditions, and those changes would not be observed in this timeframe. Second, changes may lead to improvements only after a longer time horizon. For example, if improved irrigation allows farmers to plant orchards in fields that are currently fallow, those orchards might not be mature enough to bear fruit in the first agricultural season. Thus, examining intermediate outcomes will be especially crucial for this impact evaluation so that we can gauge whether future improvements in household well-being are possible. We would expect an impact on households' income only if we observe that a substantial proportion of the targeted farmers actually had improved irrigation, and perhaps most importantly, are then utilizing the improved irrigation to improve their agricultural productivity. Indeed, MCC modeled the primary paths for increasing farm income through an economic rate of return model (ERR) that justified the project for funding. The ERR included two agricultural benefits as a result of improved irrigation infrastructure. Those were (1) increased land under high-value agriculture (for example, switching crops from wheat to cucumbers) and (2) higher yields. Examining the intermediate outcomes also establishes the counterfactual—what the irrigation conditions would have been

even in the absence of irrigation rehabilitation. Table 2 summarizes the key intermediate outcomes that can be examined using the TCS data.

Table 2. Key Outcome Measures: Survey Data

Intermediate Outcome Measures	Time Frame
Water Usage. Amount of land that could be irrigated; amount of land that actually was irrigated; amount of land watered using other sources (such as well or drinking water); frequency of irrigation; estimated amount of irrigation water used.	Last Agricultural Season
Quality of Irrigation System. Perceived overall quality of irrigation in the village; perceived changes in quality from previous year; main irrigation problems in the village; timeliness and sufficiency of irrigation water.	Last Agricultural Season
Investment in Agricultural Technology or Equipment. Ownership of personal reservoir or water pump; irrigation technologies used; WUA membership.	Last Agricultural Season
Cropping Patterns. Specific crops grown, especially high-value crops; amount of land devoted to cultivation of each crop; total hectares of land devoted to crops; amount of land irrigated for each crop.	Last Agricultural Season

WUA Administrative Records

Water User Associations (WUAs) are responsible for managing irrigation systems that serve nearly all Armenian farmers, collecting associated water payments, and providing irrigation to the farmers. We plan to supplement the TCS data with administrative data from the WUAs. These data will be collected for both tertiary canal communities and their matched comparisons. Although these data may only be available for some areas, they could nevertheless provide more detailed information on outcomes related to water usage, for which survey data is necessarily limited. These data may not be perfectly reliable due to reporting errors by the people who manage water supply in the fields (known as "ditchmasters") who record water provisions or data entry mistakes, but this is still the best source of data on water usage that is available. Many WUAs also collect data on types of crops grown on each plot. Although those data are not as detailed as the TCS and are probably not available for all WUAs, they have the advantage of being available for all plots in the WUA, not just the survey respondent farmers. Thus, this data source could conceivably yield more precise impact estimates, and could also be used to determine if the survey sample of farmers is representative of the broader population of farmers in the area. Table 3 summarizes the outcome measures that can potentially be obtained from WUA administrative records.

Table 3. Key Outcome Measures: WUA Administrative Data

Intermediate Outcome Measures	Time Frame
Water Usage. Amount of water supplied; water paid for; water debt owed to WUA.	Last Agricultural Season
Cropping Patterns. Type of crops grown, especially high-value crops; amount of land devoted to cultivation of each crop; total hectares of land devoted to crops.	Last Agricultural Season

III. ANALYSIS PLAN

Because the communities where tertiary canals are rehabilitated are not randomly selected, and it is unlikely that perfect comparison communities can be identified in all cases, it is critical that we have data on the key outcome measures prior to the intervention, so any baseline differences between the two groups can be accounted for in the analysis. Many of the rehabilitation efforts will begin in early 2010, which is also when the baseline survey would likely be fielded. Because this is between agricultural seasons and the survey will ask about the previous agricultural season, the timing is good. As noted before, because the pilot tertiary canals will have been rehabilitated earlier, they cannot be included in this analysis.

The impacts of these tertiary canals would be estimated based on the follow-up surveys in Spring 2011, after the irrigation projects are complete. As noted earlier, a third survey round will likely be fielded in Spring 2012, and the analysis approach for that round would be analogous to what is described here. The impact of the irrigation rehabilitation on a given outcome measure will essentially be calculated by subtracting the average value of that outcome measure for the matched comparison group from the average value for the treatment group. As noted above, however, it is crucial to adjust for pre-existing differences between the treatment and comparison groups to ensure any observed differences at the time of the follow-up survey can be credibly attributed to the program. Using regression models to control for these baseline characteristics also improves the statistical precision of the impact estimates. The basic regression model can be expressed as follows:

$$(1)\quad y_{iv,F} = \beta' x_{iv,B} + \gamma T_v + \eta_v + \varepsilon_{iv},$$

where $y_{iv,F}$ is the outcome of interest for household i in village v at the follow-up survey; $x_{iv,B}$ is a vector of baseline characteristics; T_v is an indicator equal to one if village v is in the treatment group and zero if it is in the comparison group; η_v is a village-specific error term; and ε_{iv} is a random error term for the household. The parameter estimate for γ is the estimated impact of the program.

The vector of baseline characteristics x_{iv} will include both household and village characteristics. At a minimum, we will control for village characteristics such as the geographic region, size of the community, and number of farmers. We will also control for household size and composition, and characteristics of the household head, namely, education level, gender, age, and number of years farming. Baseline measures of the outcome measures of interest will also be included in the regression. The regression models must also account for the fact that farmers, because farmers served by the same canal are exposed to the same effects of weather and other idiosyncratic shocks, their outcomes will be correlated and cannot be considered statistically independent. This "clustering" of farmers is reflected in the village-specific error term η_v. Lastly, we will weight the data to adjust for cases where different numbers of comparison canals are matched to a given treatment canal (and vice versa).

An alternative specification of (1) would define $y_{iv,F}$ as the growth in outcomes, rather than the post-rehabilitation measure of the outcome. This formulation is sometimes preferred in situations where the outcome is measured with error due to recall error, which is usually the case for complicated outcome measures such as household income. Hence, we plan to conduct sensitivity analyses using this alternative specification.

In addition to the quantitative impact evaluation described in this memo, MCA-Armenia is also funding a qualitative process analysis that investigates issues such as how the irrigation rehabilitation project was designed, the fidelity of program implementation, and stakeholders' perceptions of program implementation and benefits. This qualitative information will provide valuable insights that complement the quantitative findings by helping us determine why the expected program impacts did or did not occur.

Tables

We will conduct both descriptive analysis and analysis of the impacts of the program. The baseline report will focus primarily on descriptive statistics of the outcome measures using the data collected prior to rehabilitation. The final impact report will include a variety of tables and figures with descriptive statistics on the data; however, the focus will be on estimates of program impacts and their statistical significance. We will also report regression-adjusted means for the treatment and comparison groups—that is, the means for the two groups if they had identical village and household characteristics. Table 4 provides an example of the structure of these tables.

Table 4. Example of Impact Estimates Table

	Treatment Mean	*Comparison Mean*	*Program Impact*	*p-Value of Impact*
Total Value of Crops				
Total Ag. Expenditures				
Profit from Agriculture				
….				
….				

Notes: */**/*** indicate statistical significance at the .10/.05/.01 level. Treatment and comparison means are regression-adjusted to account for any pre-existing differences in village or household characteristics of the two groups.

Exploring Program Complementarities

We will also attempt to conduct analysis to see whether the impacts of the tertiary canal rehabilitation differ for communities where the WtM training programs were also offered compared with communities where WtM was not offered. Where possible, this would inform us about whether the two programs were complementary—for example, if farmers are better able to leverage the skills taught in the WtM program if they have access to improved irrigation. The analysis would essentially estimate separate impacts of irrigation for the WtM treatment communities and the WtM control communities. The impacts would be estimated by modifying equation (1) as follows:

$$(2)\quad y_{iv,F} = \beta' x_{iv,B} + \gamma_{W=1} T_v \times W_{iv} + \gamma_{W=0} T_v \times (1 - W_{iv}) + \eta_v + \varepsilon_{iv},$$

where W_{iv} equals 1 if the village is also a WtM treatment community and 0 if it is in the WtM control group. Testing whether $\gamma_{W=1}$ differs from $\gamma_{W=0}$ would tell us whether the impacts differ for the two groups. Analogously, we could estimate whether the WtM impact differs for communities that did or did not also have rehabilitated tertiary canals.

This analysis would necessarily be restricted to the subset of communities that overlap in both the irrigation and WtM samples, which constrains the statistical precision; only about one-third of the TCS communities are also in the WtM treatment or control group (approximately evenly split between the two). Moreover, the WtM participants are not a random subset of farmers, so the sampling approach for the TCS may not yield many WtM participants even in the WtM treatment villages. Because the FPS targets likely WtM participants, we anticipate that the next round of the FPS will have a higher density of those farmers and would serve as a better data source for this combined WtM/irrigation analysis, but we will examine both data sources to see which is more suitable. The WUA administrative data may also be useful for this analysis. Although we expect that this analysis will be limited by the constraints described above, this exploratory analysis can nevertheless provide some insight into whether and how irrigation and WtM impacts varied across different communities.

IV. SCHEDULE FOR REPORTING

We will conduct two sets of analyses based on the TCS. The first will be a report on the baseline data collected in 2009-2010. This baseline report will emphasize understanding the current agricultural situation in the targeted communities. We plan to submit a draft report by the end of July 2010, and a final draft that incorporates feedback from MCA-Armenia and MCC will be submitted in early Fall 2010. The final impact analysis will be conducted a little more than a year later, with a draft report due in Fall 2011 and a final report a couple of months later. However, as discussed previously, it is likely that a third round of data collection will be added, in which case the final impact analysis would probably be one year later but on the same approximate monthly schedule, culminating in a draft report in Fall 2012.