ARMENIA'S MILLENNIUM CHALLENGE ACCOUNT: ASSESSING IMPACTS ON ECONOMIC GROWTH AND POVERTY REDUCTION IN RURAL ARMENIA

Kenneth Fortson, Mathematica
Ester Hakobyan, MCA-Armenia
Anahit Petrosyan, MCA-Armenia
Anu Rangarajan, Mathematica
Rebecca Tunstall, MCC[*]

Abstract: *The Millennium Challenge Corporation's $236 million Compact with Armenia aims to reduce rural poverty through investments in irrigation infrastructure, rural road rehabilitation, and farmer training. This paper discusses the design for a rigorous impact evaluation of the farmer training project, which focuses on water management and cultivation of high-value crops. The evaluation uses a random assignment design, whereby rural communities are randomly assigned to a treatment group, who are offered agricultural training early in the Compact, or a control group, who are not. We will then compare farmers' outcomes in treatment group villages to farmers in control group villages. The results of the evaluation will assess the farmer training program's success in accomplishing its key objectives: Increasing adoption of effective agricultural practices, increasing cultivation of higher-value crops, improving farm productivity, increasing agricultural profits and household income, and reducing poverty rates. Moreover, the lessons learned in this context can be applied to other countries with similar economic climates so that future investments can be efficiently allocated to interventions with proven economic impacts.*

JEL Classifications: O12; O13; Q12
Keywords: Poverty; Agriculture; Training; Random Assignment; Armenia

[*] This project has benefited greatly from the hard work of numerous colleagues at MCC, MCA-Armenia, and Mathematica. We are especially grateful to Emily Andrews, Sergey Balasanyan, Randall Blair, Barry Deren, Katherine Farley, Melik Gasparyan, Mary Grider, Ara Hovsepyan, Julie Ingels, Lusine Kharatyan, Sergey Meloyan, Daniel Player, Stephanie Roueche, Peter Schochet, and Sonya Vartivarian for their input throughout the project. We also thank VISTAA and ACDI/VOCA for facilitating a rigorous evaluation as they implement their programs, and AREG for implementing the baseline survey of farmer households that serves as the key data source for the evaluation.

I. INTRODUCTION

In March of 2006 Armenia signed a $235.65 million agreement (the “Compact”) with the Millennium Challenge Corporation (MCC) with the goal of reducing rural poverty through improved performance of the country’s agricultural sector. Armenia plans to achieve this goal through a five-year program of investments in rural roads, irrigation infrastructure, and technical and financial assistance to support farmers and agribusinesses.

This paper will focus on the technical training in agricultural practices provided to Armenian farmers through the Compact’s Water-to-Market Activity (WtM), and particularly, the accompanying impact evaluation.

The MCC program aims to train sixty-thousand farmers in region-specific water management techniques. These methods will help farmers to use water more efficiently, which can promote both cost savings (through water conservation) and increased quantity and quality of crops. Many farmers who will be trained in water management technologies will also receive training in higher-value agricultural (HVA) methods. The combination of these activities is expected to increase beneficiaries’ average net income by about 25 percent.

The Compact provides funding for an innovative impact evaluation to test the assumptions used in the pre-Compact economic analysis by assessing the impact of the agricultural training on farmer productivity and income. The key research questions guiding the design of the evaluation for the agricultural training are:

- Did the program affect the irrigation and agricultural practices of Armenian farmers?
- Did the program affect agricultural productivity?
- Did the program improve household well-being for the targeted farmers, including income and poverty?

The backbone of our research design is the use of random assignment to create two statistically equivalent groups of farmers, with the only difference between them being that one group can participate in training while the other cannot. This research design is considered the most rigorous methodological approach for estimating program impacts. Random assignment has been implemented to estimate the effects of programs in many contexts, and is especially widespread in developed countries (Michalopoulous, 2005; Kling, 2007). Recent research has also expanded the use of random assignment into studies of programs in developing countries (Duflo and Kremer, 2004). To our knowledge, this evaluation of the Compact’s WtM training program is the first large-scale random assignment evaluation in Armenia.

The next section provides a summary of the recent history of rural Armenia as context (Section II), followed by an overview of the Compact programs (Section III). Section IV provides a detailed description of the WtM activity. With this as background, we describe the evaluation design, beginning with the random assignment design in Section V. We then discuss the main source of data, the Farming Practices Survey (Section VI), followed by the Irrigation PIU data that could be used as a supplemental data source (Section VII). Lastly, we discuss in detail our econometric approach for estimating program impacts in Section VIII before concluding with a summary of the next steps in Section IX.

II. RECENT DEVELOPMENTS IN ARMENIA

In January 1991, Armenia began implementing a comprehensive land reform program. Rapid privatization in 1991-1992, when one-third of all agricultural land and 70 percent of arable land were transferred to family farms and the Soviet-style collective farms were abolished, created a base for a stable recovery of the sector. The decision to move rapidly reflected the fact that implementation of the land reform had the potential to boost a rapid supply of agricultural products, thus increasing the speed of moving to a market economy. In parallel with the land privatization, a set of reform measures designed to liberalize agriculture was implemented. Price controls were removed and food subsidies were abolished. The major support measures to local agricultural producers included VAT and land tax exemptions as well as subsidies for irrigation water.

After the sharp decline of agricultural output at the beginning of the 1990s, the economy of Armenia started to recover beginning in 1994. The agricultural sector has continued to register steady growth since that time. However, agriculture has lagged behind other sectors of the economy such as the industrial, service, and construction sectors; as noted in Table 2.1, agriculture represents a shrinking share of GDP in spite of increased rates of growth.

Table 2.1 GDP and Agricultural Product

	Average 1994-1999	Average 2000-2007
GDP Growth rate, percent	5.4	11.8
Agriculture Product Growth rate, percent	3.2	6.8
Share of Agriculture Product in GDP, percent	35.5	22.5

Source: Armenia in Figures, National Statistical Service of the Republic of Armenia

Another feature of Armenian agriculture had been the unfavorable changes in relative prices. Average annual growth of agricultural prices in 1995-1999 was much slower than increases in industrial and consumer prices (Table 2.2). A similar pattern, though dampened, was registered during 2000-2007. In fact, lagging agricultural prices reflect both weaknesses of demand, which steadily improved during recent years, and weak market power of farmers compared to wholesales and food processors. The disproportionate price development indicates that the growth in agriculture was mostly beneficial to consumers and the food processing industry.

Table 2.2 Changes in Price Indices

	Annual average for 1995-1999	Annual average for 2000-2007
Consumer Price Index	143.6	103.2
Industrial Price Index	166.5	105.3
Agricultural Price Index	127.3	102.4

Source: Armenia in Figures, National Statistical Service of the Republic of Armenia

The lower growth rate of the agricultural sector compared with other sectors reflects the relatively lower productivity of the agriculture that is explained by different factors:

- Small, fragmented farms (average land holding per family is 1.4 ha) limit the use of agricultural machinery. As a result production is labor intensive, and labor productivity is low relative to capital-intensive production;
- The poor condition of irrigation systems combined with climatic conditions of Armenia decreases efficiency of agricultural production;
- Because a large share of the rural population, previously working in the non-agricultural sector, became subsistence farmers in the 1990s, many farmers are inexperienced and lack basic knowledge about farming practices;
- Inadequate use and availability of improved inputs such as seed, pesticides and appropriate fertilizer applications reduces the productivity of the sector;
- Limited opportunities for storage, grading, packaging, and even processing of agricultural production reduces the value-added and limits the supply of fresh produce throughout a year, and also makes agricultural prices volatile during a year;
- Farmers have limited knowledge of the markets or access to information on demand;
- Inadequate financial services are available for the rural sector and farmers have limited access to the financial resources; and
- Collaboration among farmers in production and marketing of their products is limited.

III. MCA-Armenia Program

The Compact is designed to address many of the above issues and to increase agricultural productivity and reduce rural poverty in Armenia. MCA-Armenia Program is addressing these obstacles through two projects: Rural Road Rehabilitation and Irrigated Agriculture Projects.

The first project implemented under the MCA-Armenia Program is the Rural Roads Rehabilitation Project. The implementation of the Project will expand the access of rural communities to agricultural markets and social infrastructure, as well as increase non-farm income opportunities by improving the condition of rural roads. The Compact was designed to include MCC funding to rehabilitate up to 943 km of the rural roads from 'very poor', or 'poor', to 'good' condition. The project is expected to benefit hundreds of thousands of rural Armenians.

The second project, the Irrigated Agriculture Project, will increase the profitability of the agricultural sector by extending and improving the irrigation network, strengthening irrigation entities to better manage the network, and supporting farmers to commercialize their production. For these purposes two types of activities will be implemented under the Irrigated Agriculture Project. The first activity, Irrigation Infrastructure Rehabilitation, will improve existing infrastructure in order to increase irrigated area and to improve the efficiency and sustainability of sourcing and delivery of water. It is expected that the Project implementation will result in expansion of irrigated areas, conversion of selected irrigation areas from pump to gravity irrigation to reduce energy consumption and make water more affordable, as well as reduce water losses in the irrigation network. The primary project beneficiaries are expected to be more than 115,000 farming households (or about 40 percent of all rural households) who will be able to increase the productivity of their irrigated land.

The second activity in the Irrigated Agriculture Project, the Water-to-Market Activity (WtM), will accelerate the transition to more profitable, commercially-orientated agricultural production by introducing and encouraging best practices in irrigated agriculture, fostering the adoption of improved water management techniques, shifting or expanding to higher value crops, strengthening the post-harvest and processing enterprises linking producers to their markets, both domestic and international, and strengthening the capacity of credit providers to fund viable proposals in production and post-harvest activities.

IV. WATER-TO-MARKET TRAINING

The focus of the paper is on the Water-to-Market Activity, specifically, how this activity impacts agricultural productivity and how the impact of the activity will be evaluated. The WtM Activity is designed to ensure that MCC Funding to the Irrigation Infrastructure Rehabilitation will contribute to a sustainable increase of agricultural productivity and incomes from farming. The Water-to-Market Activity includes four types of sub-activities covering the whole chain from irrigation to the delivery of agricultural products to the consumers. MCA has contracted with ACDI/VOCA and their partners, VISTAA and Euroconsult (hereafter referred to collectively as ACDI) to implement the WtM activities that include training farmers in water management and high value agriculture, as well as credit and post-production services.

Introduction of New On-Farm Water Management Technologies. The objective of the sub-activity is to improve farmers' skills in on-farm water management techniques and their access to farm equipment to enhance the efficiency of water use in irrigation. The objective will be reached through implementation of the integrated program including in-class training and on-field demonstrations of the advanced and location-appropriate on-farm irrigation practices. Training and demonstrations will be provided to 60,000 farmers, of whom approximately 65 percent are expected to adopt water saving and productivity innovations that will increase the net benefit of their farming operations.

The demonstration farm is a critical part of the initial training and also can help reinforce the lessons even after training is completed. Farms are selected to serve as demonstration sites based on their proximity to other farmers in the village and the demonstration farmer's willingness to adopt new technologies and facilitate other farmers' understanding. ACDI provides equipment for the demonstration farms; in exchange, the demonstration farms are used as the site of the practical training, and also a resource where farmers in the village can go to see the technologies in practice, beyond the official training session.

Demonstration farms will serve anywhere from one to five villages, depending on the number of eligible farmers in those villages and their proximity. Some demonstration farms exist from previous activities, and these will be supplemented with new demonstration farms. Most demonstration farms will be new. The trainers will be agricultural experts who come from the same region so as to ensure that they are familiar with the local climatic and agricultural conditions, and so that they are available for technical assistance beyond the training sessions. Once a demonstration farm is established, ACDI will provide several rounds of training at that demonstration farm to saturate the associated villages as much as possible, because high participation rates will maximize the use of a single demonstration farm.

Coordinators will target farmers who are members of Water User Associations (WUAs). They will use posters to advertise at village centers and work with mayors to mobilize farmers to participate. In some villages, mayors may also be able to identify farmers who are particularly likely to participate, and these farmers could be targeted for additional recruitment efforts.

Transition to Higher Value Agriculture. The objective of this sub-activity is to support the farmers' transition to more profitable agricultural production through a combination of crop substitution, increased cropping intensity, and the use of higher yield generating planting materials and related inputs. Like the on-farm water management training, the objective of this sub-activity will also be reached through implementation of the integrated program including in-class training and on-field demonstrations of advanced, location-appropriate farming practices. 30,000 farmers will be trained in HVA to increase the commercial value of their farms' output. By the end of the Program, it is expected that approximately 7,800 hectares of land will be converted into higher-value agricultural cultivation as a result of increased access to water combined with an effective training program.

Post-Harvest, Processing and Marketing. The objective of this sub-activity is to introduce and expand post-harvest operations that best preserve the quality of agriculture products and add value to production. For this purpose new strategies and technologies for storage, sorting and packaging, transportation and processing of agricultural products will be introduced. Improved supply relationships between post-harvest enterprises and their farm suppliers will be another important aspect of this sub-activity. The assistance to 12 firms includes diagnostic analysis of enterprises and providing recommendations for improvements, conducting food safety and hygiene training, providing weekly market information on fresh fruits and vegetables as well as other related activities. Reliable information on market conditions and opportunities will contribute to the better positioning of fresh and processed food products in the domestic, regional and international markets and create trade relationships. Food safety and quality assurance issues such as Hazard Analysis and Critical Control Point (HACCP) and International Standards Organization (ISO) certification will also be addressed through this sub-activity. By the end of the Program, it is expected that these activities will impact 300 agribusinesses and 15,000 farmers.

Improved Access to Credit. The objectives of this sub-activity are two-fold. Improved irrigation and rural roads will present new opportunities for farmers and agricultural related businesses. However, improved irrigation and rural roads will create new opportunities for financial institutions as well. Therefore, the objective of this sub-activity is to increase the availability of longer-term, affordable credit to beneficiaries of the Project. To reach this objective, the sub-activity will provide USD 8.5 million on-lending resources to banks and credit organizations. The sub-activity will also develop the capacity of banks and credit organizations to lend efficiently and prudently in the agriculture sector. The sub-activity will provide assistance in applying for loans to the potential loan applicants as well. The assistance will improve the ability of the farmers to access financial resources and increase their awareness and understanding of credit. These actions will facilitate the intermediation process and reduce transaction costs and risk for credit providers by developing better informed and better prepared borrowers.

Table 4.1 summarizes Water-to-Market Activity's training and adoption targets for these four sub-activities. Through these activities, the direct impact of the WtM Activity will be to generate new employment opportunities and increase income among farms and rural businesses. The synergy between rural infrastructure development and agribusiness development is designed to result in a significant, sustainable reduction in rural poverty.

Table 4.1 WtM Adoption Targets

On-Farm Water Management Training	60,000 farmers
Higher Value Agriculture Training	30,000 farmers
Post-harvest Enterprises Technical Assistance	300 enterprises
Bank Loans provided to project beneficiaries and related businesses	$8.5 million in loans
Adoption of improved on-farm water management	38,350 farmers
Hectares Converted to High Value Agriculture	7,845 hectares

The implementation of the program started in February of 2007. The implementation process includes two major stages. The first stage, the pilot stage, covered February-September 2007 and it had a very specific purpose to provide lessons learned for full-scale program implementation. The first stage was mostly focused on on-farm water management training as a strategic tool to improve farmers' skills in the on-farm water management and irrigation techniques, as well as their access to new irrigation technology to enhance the efficiency of water use. Through demonstration and training, farmers are encouraged to adopt new or improved irrigation methodologies and methods which are effective and affordable for different categories of farmers. Therefore, the activities in demonstration sites have to suit all categories of farmers in terms of size and show a wide range of system improvements from the simple and inexpensive to the more costly ones. The general emphasis is on saving water and labor through proper irrigation scheduling, more uniform distribution of water within the farm and more effective delivery to the field. Development of skills and improved farmers' knowledge are viewed as the most crucial factors for achieving a satisfactory level of adoption and sustainability of the expected changes.

Full-scale implementation of the Water-to-Market program started in October 2007. By the end of March 2008 the number of participants completing the On-Farm Water Management training reached 8,186 farmers. Besides the training on on-farm water management techniques, other activities were implemented as well. The number of participants completing training in High Value Agriculture, which helps farmers identify more profitable crops for their specific regions and best practices for cultivating those crops, totaled 495 farmers.

V. USING RANDOM ASSIGNMENT TO ESTIMATE IMPACTS

The ideal method for separating program influences from other factors is to compare outcomes for the group who are provided the intervention with the outcomes for the same group if they were not offered the intervention. However, once persons or communities are offered the intervention, it is not possible to know what their outcomes would have been if they were not given the opportunity to participate. It can only be approximated by comparing their outcomes to the outcomes of some other group.

Random assignment, the most rigorous way to measure program impacts, is frequently referred to as the gold standard of evaluation designs. Essentially, when implemented carefully, random assignment leads to the creation of two virtually identical groups at baseline, with the only difference being that only one group (the treatment group) is exposed to the intervention, while the other group (the control group) is not. As a result, any changes observed between the two groups over time can be attributed to the effects of the intervention with a known degree of statistical precision. (Michalopoulous, 2005; Kling, 2007; Duflo and Kremer, 2004)

Unit of Random Assignment. Ideally, we would randomly assign individual farmers to receive training or not, and compare outcomes for the two groups. However, because these training sessions are community - level interventions, making it difficult to exclude individual farmers, such an approach is not feasible in this context. Our basic approach is to randomly assign villages to the treatment group of farmers, who are eligible for on-farm water management training and HVA training, or the control group of farmers, who are not. ACDI has grouped villages into clusters. Most clusters include only one village, but some include as many as five villages that are in close geographic proximity. All farmers who are WUA members and live in a cluster of villages selected for the treatment group will be permitted to participate in water management training. Farmers who are in the control group of village clusters would not be offered water management training until several years later. All villages will ultimately be provided training, and random assignment is used to determine *when* they are offered training.

Randomly assigning entire village clusters in this way, rather than individual farmers or villages, guards against contamination of the control group - the possibility that control group members get the same services as the treatment group. There are two types of contamination. The first type of contamination is if farmers in control group villages nonetheless participate in training. This could be problematic if control group members hear about the training activities and show up to training themselves. A different type of contamination could occur if farmers who participate in training teach farmers in the control group about the techniques they learned. Either of these types of contamination would be problematic for the evaluation because we would be unable to compare those who were offered training to those who were not offered training; with contamination, both the treatment and control group have access to or benefit from training. Generally, ACDI has chosen village clusters that are sufficiently far apart geographically to ensure that there is little chance that farmers in a control group village cluster would either participate in the training or learn about the water management techniques through other means.

However, in some areas - particularly the Ararat Valley region - many villages are located in close proximity. While we cannot completely eliminate the possibility of contamination here, it will be important for the planned implementation to strive to avoid such contamination problems by, for example, ensuring that recruiting techniques for the training attract treatment group farmers without influencing control group farmers. The WtM training program exit questionnaire will also inform us about where farmers reside, which will help us assess the extent to which control group farmers are “crossing over” and receiving training in spite of being randomly assigned to not be eligible.

Implementing Random Assignment. Random assignment was conducted for the subset of villages that have adequate water and could potentially be served early in the Compact. We randomly assigned villages to one of three groups: those who would be served in the second year of the Compact; those who would be served in either year 3 or year 4 of the Compact; and those who can be served in the final year of the Compact. The earliest group constitutes our treatment group, and the latest group our control group - impacts will be measured after the treatment group has been provided training but before the control group has. The middle group, those who are served in the third or fourth year, will not be included in the impact evaluation. Only villages that were considered ready for WtM training were included in the randomization; some villages currently have poor sources of water, and thus, would not benefit from training until their irrigation systems are rehabilitated. Such villages may receive training in the future, but they will not be included in the impact evaluation. We also excluded from the random assignment all villages that were included in the pilot phase of the WtM training or where ACDI has already developed demonstration farms.

Random assignment was conducted within strata defined by WUAs to preserve regional balance, which created balanced treatment and control groups along this dimension. The distribution of villages by treatment status for each agricultural zone is reported in Table 5.1 The probability of a village being assigned to the treatment group was approximately the same for all WUAs, with most of the deviations occurring due to rounding. An exception, however, is the Mountainous Zone, where a smaller proportion of villages were selected for the research groups (years 2 and 5), while most villages were assigned to the non-research group. This zone was undersampled largely because MCC anticipates very low impacts, so the evaluation will focus more on the other zones where MCC is more optimistic about the prospects for improvement. A total of 120 clusters were assigned to the treatment group and 80 to the control group, with these 200 clusters containing 223 villages in total. (For simplicity of exposition, we hereafter refer to village clusters as "villages.")

Table 5.1 Distribution of Village Clusters by Year of Training and Agricultural Zone

	Ararat Valley	Pre-Mountainous	Mountainous	Sub-Tropical	Yearly Total
Year 2: Treatment	44	58	12	6	120
Years 3 and 4: Nonresearch	18	19	38	2	77
Year 5: Control	28	38	10	4	80
Total	90	115	60	12	277

To ensure the process was transparent to all villages that will be served in the coming years, random assignment was conducted in public using a specially-designed computer program that randomly sorted villages into the three groups.

VI. FARMING PRACTICES SURVEY

The Farming Practices Survey (FPS) will serve as the primary data source for the impact evaluation. Approximately 25 interviews will be completed in each of the villages in the treatment and control groups, with fewer in the smallest villages and more in the largest. MCA-Armenia has contracted with AREG to field the FPS in the first, baseline year, and the FPS will subsequently be conducted each year of the follow-up period, at the end of 2008, 2009, and 2010.

Sample Frame. With the help of village mayors, the FPS targets the households of farmers who are most likely to benefit from the training programs: those who are actively engaged in farming and have been tied to the community for several years. These farmers are identified through an iterative process. First, MCA-Armenia requested that the WUAs work with village mayors to compile a list of farmers meeting our specific criteria in each village.

The number of farmers requested depended on the size of the village, but averaged about 60. The sample was then drawn from these lists.

Pretesting revealed that these lists were of mixed quality, however, often because the WUAs had not consulted with the mayors in compiling them. Thus, the sample was updated with the assistance of village mayors and marz officials, either at the marz offices or in the village itself. The mayors reviewed the lists to determine whether the farmers indeed met our criteria. If an insufficient number of farmers from the lists were eligible - that is, in cases where the WUA had failed to consult with the mayor - then the mayor helped AREG update the list in accordance with our survey eligibility criteria.

Relying on mayors to identify eligible farmers in each village is not ideal, as the lists they provide may not be representative of the village as a whole. Unfortunately, a reliable sample frame was not available, and the cost of conducting a census in project communities is prohibitive. However, as the method AREG used mirrors the manner in which ACDI identifies farmers to participate in training, the sample will still provide good coverage of the population most likely to benefit from training. Moreover, because of the random assignment design, treatment and control villages should still be comparable (on average), preserving the internal validity of the evaluation.

Intermediate Outcomes. While most of the outcomes of primary interest are longer-term outcomes, such as economic improvements, these outcomes may not be immediately observable. Consequently, we will closely examine intermediate outcomes through which the training programs are intended to improve household income. We would expect an impact on households' income only if we observe that a substantial proportion of the targeted farmers are actually participating in training, and perhaps most importantly, are then applying the techniques they learn. Examining the intermediate outcomes also establishes the counterfactual - what services the villages would have received and what practices they would have adopted even in the absence of the WtM programs. Table 6.1 summarizes the key intermediate that can be examined using the FPS data.

Table 6.1 Intermediate Outcome Measures

Intermediate Outcome Measures	Time Frame
Participation in Agricultural Training. Whether attended any irrigation or agriculture training (including training sponsored by other sources); type of training attended (e.g., classroom, video, or practical); whether received a certificate indicating the full training was attended.	Last Year
Adoption of HVA and Irrigation Practices. Which irrigation practices were used, focusing on those taught in training sessions; whether those practices had perceived time or labor savings.	Last Agricultural Season
Investment in Agricultural Technology or Equipment. Ownership of personal reservoir or water pump; ownership or rental of trucks, tractors, combines, seed planters, and sprayers.	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; whether household cultivates a kitchen plot; reason(s) for changes in cropping patterns.	Last Agricultural Season

Final Outcomes. The ultimate goal of the MCA-Armenia programs is to increase household income in rural Armenia, and hence, these outcomes are an important focus of the FPS 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 the training programs, 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 village to monetize crops that are consumed by the household or bartered. Additionally, the FPS asks for estimates of expenditures on key categories of consumption, and for income from other sources. Table 6.2 summarizes the key final outcomes that can be examined using the FPS data.

Survey Nonresponse. All interviews are conducted in person, and a limited span of time is available for interviews in the majority of villages. Therefore, survey nonresponse is a concern. Substantial survey nonresponse can damage the validity of impact estimates. Nonresponse weights can account for some of the differential nonresponse, but only to the extent that nonresponse is explained by household characteristics that are known for both respondents and nonrespondents. In the worst case, survey nonresponse might be different for treatment and control villages, contradicting the core assumptions of a random assignment design. More commonly, however, survey nonresponse affects both the treatment and control groups equally. In this scenario, the impact estimates remain *internally* valid, but may not generalize beyond the select group of survey respondents. As a salient example, if farmers' absences are due to trips to the markets to sell their produce, then the respondents may have a disproportionate share of the less engaged farmers, for whom program impacts could be minimal.

We have instituted several safeguards against survey nonresponse. Working with mayors to clean the lists in advance can help in this regard. Whenever possible, village mayors would also contact the sampled households in advance to ensure they would be available for interviews on the day AREG visited their village. In instances where a household is not available on the first attempt, interviewers would return to the household throughout their time in that village. AREG also has reserve lists of farmers which they can draw on to help them reach their targets for completed interviews within each village.

Follow-up Surveys. Ideally, each round of the FPS would interview the same set of households, yielding a longitudinal data set. Analytically, longitudinal data allow for the cleanest estimation of program impacts, and also provide the most statistical precision, because changes from the baseline to the follow-up period are not confounded with sampling variability. As a practical matter, however, it may not be as easy to track specific households from year to year. Our plan is to survey the same set of households in subsequent rounds of the FPS to the extent possible, but given the nonresponse issues described previously, we anticipate that these will need to be supplemented with additional households, yielding a mixed longitudinal-repeated cross sectional data set. The sample frame will remain consistent; so as to avoid having the samples for treatment and control villages diverge over time.

Table 6.2 Final Outcome Measures

Final Outcome Measures	Time Frame
Continuing Use of HVA and Irrigation Practices. Same as above, but focusing on changes in these practices relative to the initial follow-up years.	Last Agricultural Season
Agricultural Production. 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 (individually and in total).	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 Month
Income from Pensions, Remittances, or Social Programs. Can also be added to profits and employment income to construct a rough measure of total income.	Last Month
Household Consumption. Expenditure on purchased food, health care, housing products, utilities, and transportation; cost of purchased goods (converted from monthly to annual) plus value of crops consumed by the household.	Last Month/Last Year

VII. IRRIGATION PIU DATA

The Irrigation PIU has an impressive database for a subset of the WUAs, and by the end of the Compact, it is planned that information on members of all WUAs will be included in their databases. These data could be used to supplement the survey data in two important ways. First, they provide some outcome measures that would not be obtainable from farmers, such as energy use and water distribution. Second, they will provide data on some outcome measures for the entire population of registered WUA members in Armenia. These data items are defined in more general terms than the survey data - for example, the amount of land the WUA member plans to grow wheat on, but not actual production - but they can still be used to inform us about broad national and regional trends.

VIII. ESTIMATING PROGRAM IMPACTS

Random assignment ensures that, on average, treatment group villages and control group villages are the same, with the exception that treatment group villages are offered WtM training. Hence, the difference between the mean of the outcome of interest for the treatment group and the mean for the control group yields an unbiased estimate of the WtM program’s impact. The precision of the impact estimates can be improved, however, by controlling for other covariates in a regression model. Regression adjustment can also help alleviate any differences between the treatment and control groups in baseline characteristics that arose by chance.

Core Specification. The survey data will be cross-sectional, with a new cross-section of respondents drawn each year.[1] Given this data structure, our econometric specification is designed to compare how treatment group villages changed over time to how control group villages changed over time, controlling for idiosyncratic differences in the two groups. The basic model can be expressed as follows:

(1) $$y_{ivt} = \beta' x_{iv} + \lambda T_v \times F_t + \theta F_t + \eta_v + \varepsilon_{ivt},$$

where y_{ivt} is the outcome of interest for household i in village v at time t (where $F_t = 0$ in the baseline year and 1 in the follow-up year); x_{iv} is a vector of time-invariant characteristics of household i in village v; t accounts for any time trends between the base and follow-up years; T_v is an indicator equal to one if village v is in the treatment group and zero if it is in the control group; η_v is a village-specific error term (a village "random effect"); and ε_{ivt} is a random error term for household i in village v observed at time t. The parameter estimate for λ 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, WUA, and the baseline water conditions. 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. In the framework of a repeated cross-sectional model, however, the characteristics that are included must be restricted to those that are unaffected by the WtM programs. We must be careful with land holdings, for example, as the WtM program could conceivably induce some farmers to cultivate more land, and controlling for it would therefore understate the full program impact.

The model in equation (1) is designed to answer the general research question, "How have villages in treatment group changed from the baseline year to the follow-up year, relative to villages in the control group?" This core model can be tweaked in a variety of ways to explore alternative specifications. A simple example would be to allow the time trends to vary across regions. The specification also (implicitly) weights all respondents equally, which could be modified to either give all villages equal weight, or weights equal to the village populations.

Such explorations would not change the general interpretation of the impact estimate, but they can provide insights on two important issues. First, and of most direct interest, we can explore how robust the impact estimates are to these alternative specifications. Beyond this, however, the other regression covariates may be of independent interest, and may also provide context for interpreting the impact estimates.

Pooled Model. Instead of using data from only the base year and one follow-up year, we can also pool data from multiple waves of follow-up year surveys. The econometric

[1] As described previously, there will be substantial overlap in the household samples from the baseline year and subsequent years, but the samples will likely not be identical. If, however, the survey in subsequent years uses the same sample, we will be able to employ panel (longitudinal) data models. The intuitive interpretation of panel data models is similar to models of repeated cross-sectional data, but the estimation techniques differ somewhat from those described here.

specification would be very similar to (1), but with a separate impact estimate for each of the *n* follow-up years:

$$(2) \quad y_{ivt} = \beta' x_{iv} + \lambda_1 T_v \times F_{1t} + \lambda_2 T_v \times F_{2t} + ... + \lambda_n T_v \times F_{nt} + \theta_1 F_{1t} + \theta_2 F_{2t} + ... + \theta_n F_{nt} + \eta_v + \varepsilon_{ivt} ,$$

where $F_{nt} = 1$ if $t = n$ and 0 otherwise.

These impact estimates can then be compared to one another to see how program impacts changed over time, and could be particularly important to see whether any impacts on farming practices that are observed early on persist, and also whether impacts on longer-term outcomes, such as agricultural productivity, grow after farmers have had more time to implement new techniques and benefit from their innovations.

Clustering. The estimation techniques must take into account the correlation of outcomes for households in the same village, as they may be exposed to similar idiosyncratic influences that are not otherwise captured in the regression model, and therefore, the individual households cannot be considered statistically independent. As an example, a particular village might have abnormally good or bad weather, or could experience other economic "shocks" that are unrelated to the training program but nonetheless affect the entire village. The econometric models will account for this clustering with methods that allow flexibility in the correlation structure of the error terms. (Deaton, 1997).

Impact on Participants Only. Randomly assigning communities to be eligible for WtM training programs provides an unbiased estimate of the impact of offering this training in the villages selected for training - the "intent to treat" (ITT) effect. The ITT effect combines the effect of the intervention on both participants and non-participants in treatment villages. In many contexts, people who are offered program services but opt out of participating are unaffected by the program, while in other situations the program may nonetheless have within-village spillover effects on the outcomes of non-participants. By including questions about both participation in WtM training and adoption of WtM techniques, we will be able to determine whether there are sizable within-village spillover effects present, and how best to account for them.

When spillover effects are known to be minimal, a simple but powerful adjustment can be made to calculate the effect of the training program on participants - the effect of "treatment on the treated" (TOT). This adjustment - known colloquially as the Bloom adjustment - calculates the effect of the training program on participants by dividing the estimated impact (the ITT) by the participation rate. The intuition for this elegant result is that, if the effect of the program on non-participants is known to be zero, the estimated impact can be attributed entirely to the proportion of the treatment group that actually participated in training. Importantly, however, while the Bloom adjustment can potentially be used to account for non-participation in the impact estimate, it cannot alleviate the problem non-participation introduces for the *variance* of the impact estimate. If participation rates are low, we will not be able to detect impacts that are statistically reliable. (Bloom, 1984; Angrist, Imbens and Rubin, 1996)

Subgroup Analysis. For many of the outcome measures, it is conceivable that the effects of the interventions will vary by observable characteristics. Estimating differential impacts on female-headed households, for example, is of particular interest to MCC. We will examine whether the interventions' effects differ for key subgroups defined by the characteristics of the households such as gender, age, and level of education of the household head; size of the household; or size of farm holdings operated by the household. Similarly, we will also examine how effects vary by subgroups defined by village characteristics.

It is straightforward to embed subgroup estimates into the framework of equation (1). To do so, we include an interaction term that distinguishes treatment group members in subgroup *S* from those who are not in the subgroup:

(3) $$y_{ivt} = \beta' x_{iv} + \lambda_{S=1} T_v \times F_t \times (S_{iv} = 1) + \lambda_{S=0} T_v \times F_t \times (S_{iv} = 0) + \theta F_t + \eta_v + \varepsilon_{ivt}$$

In equation (3), the estimate of $\lambda_{S=1}$ represents the estimated impact for members of subgroup *S*, and we can test whether the impacts differ for members of that subgroup compared to everyone else by statistically testing whether $\lambda_{S=1}$ and $\lambda_{S=0}$ are equal.

Distributional Effects. The implicit focus of the analysis plan outlined above is on examining differences on the mean household. In conducting the analysis, it is also important to examine whether the interventions' effects vary at different levels of the outcome distribution. For example, the impact on agricultural real income for households with very low or very high income may differ from the impact on households at the mean. Specifically, the training programs may be such that only the higher-income households will benefit, if, for example, implementing the techniques taught in training requires investment in equipment that lower-income households cannot afford. Conversely, the poor in the community might benefit more than the wealthy if the training focuses primarily on techniques that are useful only to smaller-scale farms.

As Armenia has among the highest levels of income inequality in Europe, this distinction is not a trivial one. We will use quantile regression analysis to determine whether the intervention effects vary at different points in the distribution. Quantile regressions are analytically appealing because, similar to standard regression analysis, the quantile regression coefficients have direct and simple interpretation, thereby making it very appropriate for communicating impact estimates with policymakers.

Estimating impacts for specified quantiles starts with the same regression model as a standard model. The difference is in the methodology for estimating the parameters, which in turn, affects the interpretation of those impact estimates. While a standard regression model compares the impact for mean households, a quantile regression instead compares the impact of the interventions for a specified percentile, such as the 25th or the 75th percentile. Quantile regressions at the 50th percentile, the median, are also more robust to the influence of extreme outliers in the data, and thus can serve to validate the findings from standard regression analysis. (Koenker and Hallock, 2001; Deaton, 1997).

IX. NEXT STEPS

The analyses described above will be employed in a series of reports. The first report will cover the baseline FPS, and will be a short report focusing on the current state of the villages in the evaluation. The second report will cover the second round of the FPS, after the training programs have begun at least one round of training in most villages. This report will focus on the intermediate outcomes, to gauge participation rates and preliminary adoption rates for the new technology and practices. The final report will follow the fourth year of the Compact, the last year before the control group villages will become eligible for WtM training. This report will focus on the longer-term outcomes, but as discussed, it will also examine intermediate outcomes such that we can assess not only whether there have been tangible impacts on poverty and household income, but also whether there is evidence from the intermediate outcomes that the full economic impact of the WtM may not have been fully manifested yet.

The results from these studies will directly inform MCC and other international agencies on the efficacy of this approach for training in agricultural practices, and whether it is an effective pathway for reducing poverty in rural Armenia. Moreover, the lessons learned in this context can be applied to other countries with similar economic climates and will help shape future MCC Compacts with other countries.

REFERENCES

Angrist, J., G. Imbens, and D. Rubin (1996), *Identification of Causal Effects Using Instrumental Variables*, Journal of the American Statistical Association, vol. 91, no. 434

Bloom, H. (1984), *Accounting for No-Shows in Experimental Evaluation Designs*, Evaluation Review, vol. 8

Deaton, A. (1997), *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*, Johns Hopkins University Press: Baltimore, MD

Duflo, E. and M. Kremer (2004), *Use of Randomization in the Evaluation of Development Effectiveness*, in O. Feinstein, G. Ingram, and G. Pitman (ed.) Evaluating Development Effectiveness, Transaction Publishers: New Brunswick, NJ

Kling, J. (2007), *Methodological Frontiers of Public Finance Field Experiments*, National Tax Journal, vol. 60

Koenker, R. and K. Hallock (2001), *Quantile Regression*, Journal of Economic Perspectives, vol. 15, no. 4

Michalopoulos, C. (2005), *Precedents and Prospects for Randomized Experiments*, in Bloom, H. (ed.), Learning More from Social Experiments, Russell Sage Foundation: New York, NY

National Statistical Service of the Republic of Armenia (Multiple Years). *Armenia in Figures* series, available at www.armstat.am