Evaluation Design for the Agricultural Competitiveness and Enterprise Development Training Subactivity

Draft: October 5, 2011

Final: January 27, 2012

Jane Fortson
Kenneth Fortson
Evan Borkum

MATHEMATICA
Policy Research

Contract Number:
MCC-10-0114-CON-20/TO01

Mathematica Reference Number:
06865

Submitted to:
Millennium Challenge Corporation
875 Fifteenth Street NW
Washington, DC 20005
Project Officer: Rebecca Goldsmith

Submitted by:
Mathematica Policy Research
505 14th Street, Suite 800
Oakland, CA 94612-1475
Telephone: (510) 830-3700
Facsimile: (510) 830-3701
Project Directors: Jane Fortson,
Kenneth Fortson

Evaluation Design for the Agricultural Competitiveness and Enterprise Development Training Subactivity

Draft: October 5, 2012

Final: January 27, 2012

Jane Fortson
Kenneth Fortson
Evan Borkum

EXECUTIVE SUMMARY

A. Overview of the ACED Activity

As part of its compact with the government of Moldova, the Millennium Challenge Corporation (MCC) is sponsoring two projects in Moldova: the Transition to High-Value Agriculture (THVA) and Road Rehabilitation projects. The impact evaluations designed and carried out by Mathematica will focus on the THVA project. This project entails a wide range of activities intended to (1) increase rural incomes and (2) catalyze future investments in high-value agriculture (HVA).

The ultimate purpose of the evaluation of the THVA project is to determine whether, and the extent to which, the Compact activities are effective at increasing investment in HVA activities and whether and how much those HVA activities promote economic growth and poverty reduction. The evaluation focuses on three THVA activities: (1) the Access to Agricultural Finance (AAF) activity, (2) the Irrigation Sector Reform Activity and Centralized Irrigation Systems Rehabilitation Activity (ISRA-CISRA), and (3) the training subactivity of the Agricultural Competitiveness and Enterprise Development (ACED) activity. In this chapter, we describe the design of the evaluation of the ACED training subactivity, which is only one element of the ACED activity.[1] MCC, MCA-Moldova, and the U.S. Agency for International Development (USAID) will assess the full range of ACED activities through a more comprehensive impact assessment; although the details of this assessment are to be determined, it will likely rely more heavily on descriptive methods and qualitative data collection.

The ACED activity, which is being funded jointly by MCC and USAID, is designed to "increase incomes and generate jobs in rural Moldova by addressing the most critical impediments to the development of a competitive HVA sector" (ACED Contract). The implementation of the ACED activity will use a value chain approach, identifying and addressing binding constraints within particular value chains—such as tree fruits or table grapes. Therefore, the program might affect input suppliers, farmers, packers, consolidators, processors, transporters, exporters, and a variety of other value chain actors. Depending on the constraints identified, program activities could range from developing new markets to improving transportation procedures to meeting market standards for quality and appearance of produce to promoting the adoption of new crop varieties. The ACED activity consists of two components, which are being implemented in parallel: (1) growing HVA sales and (2) enterprise development in Transnistria. The first component is, in turn, organized into four subactivities: (1) HVA market development and expansion, (2) training to upgrade production and meet buyer requirements, (3) demand-driven technical assistance, and (4) the improvement of an enabling environment for HVA. Mathematica's evaluation focuses on the training subactivity.

[1] The ACED activity is also referred to as the Growing High-Value Agriculture Sales (GHS) activity.

The ACED training subactivity aims to help HVA farmers upgrade production and improve the efficiency of post-harvest activities such as processing, transporting, and delivering products to consumers. ACED trainings may involve classroom instruction, demonstration plots, farmer field days, and other methods. The expectation is that farmers will be made aware of the benefits of product upgrading and the available training opportunities and will choose to participate in the trainings. Direct participation in trainings and/or information received from others who attended trainings, together with the simultaneous relaxation of other value chain constraints through other ACED subactivities, is expected to lead to the adoption of innovative production and post-harvest practices. Adoption of these innovative practices will result in increases in production and in product upgrading, so that farmers will increase their sales and receive higher prices for their products. Finally, this is expected to translate into increases in farm revenue, farm profits, and household income.

B. Evaluation of the ACED Training Subactivity

The evaluation of the ACED training subactivity will focus on measuring the extent, if any, to which the training activities improved the productivity and profitability of participants. In particular, the evaluation will address the following research questions:

1. What is the impact of ACED farmer training on adoption of new practices, production, sales, and farm income within the context of a value chain project? Do these impacts vary by value chain?
2. Does distance from an ACED farmer training site affect participation in ACED farmer training?
3. To what degree are new practices adopted by value chain participants who do not themselves participate in ACED farmer-training activities? Can adoption by nonparticipants be attributed to program ripple effects, rather than broader trends?
4. Is the economic rate of return (ERR) for the ACED training subactivity large enough to justify the investment?

In addition to addressing these primary research questions, the evaluation will explore how impacts on practice adoption, production, sales, and farm income vary across farmers with different characteristics.

The ACED training subactivity is just one element of the ACED activity. The impact evaluation described in the present document is not designed to measure the overall impact of the ACED activity. Instead, the impact evaluation will be able to provide evidence on the impact of the training subactivity (alone) *in an environment in which other value chain constraints are concurrently addressed.* The evaluation will not necessarily be able to tell us about the impact of training in other settings or contexts: the impacts of training might be quite different when conducted outside the context of a value chain project. However, the concurrent implementation of other value chain activities—such as identifying buyers in new markets, organization of co-packing or other joint marketing arrangements, improving the cold storage chain and other key linkages, improving packaging materials, developing strategies for target markets, and training other value chain participants—is a strength of the evaluation. In other settings, training is sometimes offered, but numerous other constraints prevent trained farmers from being able to benefit. The design of the integrated value chain approach specifically addresses those binding constraints. Therefore, the evaluation is uniquely poised to provide evidence on the impact of training in what might be a best-case scenario.

The impact evaluation of the ACED training subactivity will use a random assignment evaluation design. Potential training sites were randomly assigned to a treatment group—at which training activities will be conducted—or to a control group—at which training activities will not be conducted. If all the farmers who live in (or near) a treatment site participate in training and those who live in control sites are less likely to participate, then impacts can then be estimated by comparing farmers who live in treatment sites with farmers who live in control sites.

However, some farmers who live in control sites could still choose to travel long distances to attend training, and such farmers will not be barred from participating. At the same time, some farmers who live in treatment sites may choose not to participate. Hence, random assignment does not necessarily separate farmers who attend training from those who do not. Instead, randomly assigning the location of training activities (among potential sites) changes the *probability* that farmers attend training, assuming that those who live closer are more likely to attend. If many control group farmers choose to attend training, we will account for this feature in estimating the impacts of attending training using an instrumental variables (IV) framework, which we describe in more detail in the Methodology section of this chapter.

C. Data Collection and Reporting

We will work with USAID and MCA-Moldova to collect data to inform the impact evaluation. The primary data source for the analysis will be several rounds of the Farm Operator Survey (FOS). Through the FOS, we will gather information from farm operators on key outcomes before, during, and after implementation. We will work closely with MCA-Moldova to ensure the success of the data collection effort by providing input and feedback to support MCA-Moldova's role in managing the FOS contractor. Implementation data from USAID (collected by its implementation contractor, DAI) will also inform the analysis. Finally, qualitative data from farmer focus groups (which may be collected by USAID and/or MCA-Moldova as part of the overall ACED evaluation, or in a separate effort) will provide a richer understanding of the impact of the trainings and complement the quantitative impact results.

Evaluation activities will occur between 2012 and 2016 (Figure 1). The FOS baseline will occur before the ACED activities have had a chance to influence agricultural outcomes (2012).[2] A midterm FOS follow-up will capture outcomes two years later (2013–2014); an end-of-Compact FOS follow-up will capture outcomes three years later (2014–2015); and a post-Compact FOS follow-up funded by MCC will capture outcomes four years later, providing sufficient time for final outcomes to materialize (2015–2016). The timing of qualitative data collection will be aligned with the midterm and post-Compact FOS surveys in order to best complement the quantitative analysis. MCC might, in consultation with MCA-Moldova, choose to conduct only two of the three currently planned follow-up FOS rounds, with the end-of-Compact survey (round 3 of 4) being the most likely to be removed. If data collection plans are modified, the analysis and reporting plans will be modified accordingly. Hereafter, we summarize the timeline assuming all FOS rounds are retained.

[2] Training could have begun in some treatment sites before the associated farmers have been interviewed. However, as most key FOS outcomes relate to the prior agricultural season, these baseline responses should be unaffected by training.

We will summarize our findings in several reports. In 2012, we will prepare a baseline report that will describe the data from the baseline FOS.[3] In 2014, we will prepare a midterm evaluation report based on the analysis of midterm data that will focus on preliminary outcomes and impacts. In 2015, we will prepare an end-of-Compact evaluation report based on the analysis of end-of-Compact data. Finally, in 2016 we will prepare a post-Compact report that will examine the impacts of ACED training activities after more time has elapsed.

Figure 1. Schedule of ACED Training Subactivity Evaluation Activities, by Study Year and Quarter

	2012				2013				2014				2015				2016			
	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
Data Collection and Processing** (Farm Operator Survey, ACED Sample)	B	B	B					M	M			E*	E*			P	P			
Analysis and Reporting				B	B					M	M			E*	E*			P	P	
Qualitative Data Collection and Reporting***								M								P				

* The end-of-Compact round of the FOS may be removed.

** Managed by MCA-Moldova.

*** Managed by USAID and/or MCA-Moldova.

B = baseline; M = midterm follow-up; E = end-of-Compact follow-up; P = post-Compact follow-up.

[3] This baseline report might require a modification of our contract with MCC, because it was not specified in our original scope of work.

EVALUATION DESIGN FOR THE AGRICULTURAL COMPETITIVENESS AND ENTERPRISE DEVELOPMENT TRAINING SUBACTIVITY

A. Introduction

As part of its compact with the government of Moldova, the Millennium Challenge Corporation (MCC) is sponsoring two projects in Moldova: the Transition to High-Value Agriculture (THVA) and Road Rehabilitation projects. The impact evaluations designed and carried out by Mathematica will focus on the THVA project. This project entails a wide range of activities intended to (1) increase rural incomes and (2) catalyze future investments in high-value agriculture (HVA).

The ultimate purpose of the evaluation of the THVA project is to determine whether, and the extent to which, the Compact's activities are effective at increasing investment in HVA activities and whether and how much those HVA activities promote economic growth and poverty reduction. The evaluation focuses on three THVA activities: (1) the Access to Agricultural Finance (AAF) activity, (2) the Irrigation Sector Reform Activity and Centralized Irrigation Systems Rehabilitation Activity (ISRA-CISRA), and (3) the training subactivity of the Agricultural Competitiveness and Enterprise Development (ACED) activity. In this chapter, we describe the design of the evaluation of the ACED training subactivity, which is just one component of the ACED activity.[1] The full range of ACED activities will be assessed by MCC, MCA-Moldova, and the U.S. Agency for International Development (USAID) through a more comprehensive impact assessment; although the details of this assessment are to be determined, it will likely rely more heavily on descriptive methods and qualitative data collection. Earlier chapters (submitted as separate documents) describe the evaluation designs for AAF and ISRA-CISRA.

The ACED activity, which is being funded jointly by MCC and USAID, is designed to "increase incomes and generate jobs in rural Moldova by addressing the most critical impediments to the development of a competitive HVA sector" (ACED Contract). The implementation of the ACED activity will use a value chain approach, identifying and addressing binding constraints within particular value chains—such as tree fruits or table grapes. Therefore, the program might affect input suppliers, farmers, packers, consolidators, processors, transporters, exporters, and a variety of other value chain actors. Depending on the constraints identified, program activities might range from developing new markets to improving transportation procedures to meeting market standards for quality and appearance of produce to promoting the adoption of new crop varieties. The ACED activity consists of two components, which are being implemented in parallel: (1) growing HVA sales and (2) enterprise development in Transnistria. The first component is, in turn, divided into four subactivities: (1) HVA market development and expansion, (2) training to upgrade production and meet buyer requirements, (3) demand-driven technical assistance, and (4) the improvement of an enabling environment for HVA. Mathematica's evaluation focuses on the training subactivity.

The ACED training subactivity aims to help HVA farmers upgrade production and improve the efficiency of post-harvest activities such as processing, transporting, and delivering products to consumers. ACED trainings may involve classroom instruction, demonstration plots, farmer field

[1] The ACED activity is also referred to as the Growing High-Value Agriculture Sales (GHS) activity.

days, and other methods; it is anticipated that approximately 4,300 farmers will be trained through the subactivity.

B. Research Questions and Hypotheses

The impact evaluation of the ACED training subactivity is designed to address the following research questions:

1. What is the impact of ACED farmer training on adoption of new practices, production, sales, and farm income within the context of a value chain project? Do these impacts vary by value chain?
2. Does distance from an ACED farmer training site affect participation in ACED farmer training?
3. To what degree are new practices adopted by value chain participants who do not themselves participate in ACED farmer-training activities? Can adoption by nonparticipants be attributed to program ripple effects, rather than broader trends?
4. Is the economic rate of return for the ACED training subactivity large enough to justify the investment?

The ACED training subactivity aims to help HVA farmers upgrade production and improve the efficiency of post-harvest activities, complementing other ACED subactivities. The expectation is that farmers will be made aware of the benefits of product upgrading and the available training opportunities, and will choose to participate in the trainings. Direct participation in trainings and/or information received from others who attended trainings, together with the simultaneous relaxation of other value chain constraints through other ACED activities, is expected to lead to the adoption of innovative production and post-harvest practices. Adoption of these innovative practices will result in increases in production and in product upgrading, so that farmers will increase their sales and receive higher prices for their products. Finally, this is expected to translate into increases in farm revenue, farm profits and household income. Therefore, if the training is effective—and other value chain constraints are simultaneously addressed—the ACED training subactivity is expected to lead to the adoption of new practices and increases in production, sales, and farm income. Consequently, the first of these research questions is the focus of the evaluation.

The impacts of the ACED training activity might vary by value chain. Farmers in different value chains might have different characteristics, leading them to respond differently to trainings. In addition, the nature of trainings (or complementary activities conducted through ACED) might differ across value chains, leading the impacts to differ. Therefore, we will also separately estimate impacts by value chain; however, because statistical power will be limited (especially for high-variance outcomes), our analysis will concentrate primarily on the pooled impact estimates. When we look at estimates separately by value chain, we will focus on intermediate outcomes like practice adoption (which has a lower variance). DAI, the implementation contractor, expects that many farmers living near control sites will learn about the training programs from other farmers, agricultural cooperatives, and other peer networks. These farmers may choose to attend training despite having to travel farther to do so. Question 2 will examine the relationship between distance from a training site and participation in training to determine the extent to which control group farmers participate in training. In addition to helping describe implementation, the answer to this question is important for determining the most appropriate statistical model for the evaluation, as we describe in Section C.

Another key component of the ACED design is the "ripple effect", or innovation diffusion, in which nonparticipants adopt promoted practices through farmer networks (stimulated by the ACED activity). We will examine the degree to which there are training program ripple effects on nonparticipants (Question 3).

Finally, we will provide estimates of the ex post economic rate of return of the ACED training subactivity (Question 4). MCC uses economic rates of return (ERRs) in deciding which potential projects merit funding; ex ante ERRs must typically be above a specified threshold to consider a project for funding. The calculation of ex post ERRs will draw on parameter estimates generated by the impact analysis and, in particular, the impact of the training subactivity on farm income. Ex post ERRs will be compared with MCC's threshold values; MCC will use ERR estimates for these activities to inform its decisions on future activities.

In addition to measuring the marginal impact of ACED-provided training (Question 1), our analysis will also address a related supplemental question:

5. How does the impact of ACED farmer training on adoption of new practices, production, sales, and farm income vary with the characteristics of farm operators and farm households?

In particular, MCC is interested in understanding whether impacts are smaller or larger for farmers with particular characteristics (such as those defined by gender and baseline farm size) or by farm household composition. Our ability to measure statistically significant differences across farm households will be somewhat limited, given anticipated sample sizes. Nevertheless, we will conduct exploratory analyses that could provide suggestive evidence on the relative impact of the ACED training subactivity for farmers in different groups; our ability to measure differences will likely be better for more proximate outcomes and outcomes with lower variance (such as adoption of new practices).

As part of Question 5, MCC is also interested in examining whether impacts differ inside and outside of the centralized irrigation system (CIS) areas that are part of ISRA-CISRA. In general, DAI is concentrating its training activities in areas that already have reliable access to irrigation water, as these are the areas that it believes are best positioned to benefit from implementing new practices early in the project. Because most of the areas covered by ISRA-CISRA will not have good irrigation access until after irrigation rehabilitation is complete, few are likely to be included in the ACED value chain activity until later in the project. Hence, there will likely not be sufficient overlap of the ACED and ISRA-CISRA activities during the timeframe of the evaluation to explore activity complementarities as part of the evaluation.

As described earlier, the evaluation will be able to provide evidence on the impact of training in a setting in which other value chain constraints are addressed. Therefore, our findings might not apply to other settings, particularly those with other, significant impediments to a competitive HVA sector. Instead, the impact evaluation is expected to generate key lessons on the marginal impact of project-provided training in a value chain project.

C. Methodology

To evaluate the ACED training subactivity, we will use a random assignment design in which the locations that can be used for training activities are randomly determined from among a set of potential training sites. DAI has identified 80 sites where it *could* provide training activities. If

selected for the treatment group, a given site would host trainings specific to the value chain(s) relevant in that region. These sites were randomly assigned to two groups: a treatment group in which trainings will be conducted during the first year or two of the Compact and a control group in which trainings will *not* be conducted. A third group of eight sites (the "A list") were excluded from the random assignment procedure. DAI identified these eight A list sites as high-priority sites that were particularly well suited to host training sessions. All eight A list sites, in addition to Chisinau, will host trainings.

By randomly selecting which locations can be used for trainings, we ensure that the evaluation will provide unbiased estimates of the impact of ACED farmer training on practice adoption, income, and other outcomes. The strength of a random assignment design rests on the fact that farmers in or near randomly assigned treatment and control sites should be no different, on average, before the initiation of training activities in treatment sites. Therefore, differences in outcomes can be attributed to the availability of training activities, rather than differences in characteristics or other experiences. In contrast, an evaluation that does not involve the random assignment of program activities is subject to the criticism that program activities were conducted in sites that were predisposed to succeed (or fail); if program sites are purposefully selected, we cannot convincingly disentangle the impacts of program activities from other pre-existing differences between selected and unselected sites.

Though random assignment will determine where ACED training activities are held, it will not necessarily determine which farmers participate in training. Farmers living in communities that are near control sites will be free to attend trainings held in other communities and may travel to do so; likewise, not all farmers living near treatment sites will attend trainings. If all farmers in treatment sites attended training while all farmers in control sites did not, the impacts of training could be estimated by comparing the outcomes of treatment group farmers to the outcomes of control group farmers at follow-up. If instead some farmers living near treatment sites choose not to attend training while some farmers living near control sites do attend training—which is our expectation—our evaluation approach will have to account for this phenomenon.

We will be able to measure the impacts of the ACED training subactivity as long as farmers living near treatment sites are *more likely* to attend ACED training activities than farmers who live near control sites. Our estimation approach will exploit the variation in the likelihood of attending ACED training activities induced by random assignment. In particular, the impact of the ACED training subactivity will be estimated using an instrumental variables (IV) framework, using distance from training as an instrument for participation in training. In this context, using an IV approach is not unlike a comparison of farmers in treatment and control sites, except that it adjusts for the fact that some control farmers will participate in ACED training activities and some treatment farmers will not participate. As we expect that this adjustment will be necessary, hereafter we focus on the IV methodology.

The IV approach is credible in this context because training sites were assigned randomly. To be valid, an IV approach must meet two conditions. First, the instrument (here, distance to training) must be strongly predictive of the endogenous variable (here, participation in training), as described earlier. Second, the distance to training site must be uncorrelated with outcomes except through its impact on participation in training. Without random assignment, this second assumption would likely be violated, as the determinants of participation (such as farmer motivation) might be correlated with outcomes. However, because training locations were assigned randomly, we can assume that farmers near treatment sites are the same, on average, as farmers living near control sites (before training activities take place). The IV approach isolates the component of participation that

is driven by the instrument (here, distance). As we discuss later, our IV estimates can be interpreted as the impact for a key group affected by the training subactivity—farmers who undertake training if it is offered nearby, but not if it is offered far away.

This evaluation design will enable us to measure the impacts of participating in ACED training activities. Importantly, all value chain participants could benefit from DAI's activities, whether or not they participate in training; furthermore, other activities in the value chain could amplify the benefits of training. Therefore, impacts measured through the evaluation will tell us the impacts of training in an environment in which other value chain barriers are addressed; they will not tell us the full impact of all of DAI's activities or what the impact of training would be in the absence of other, related activities.

1. Random Assignment

To implement this approach, DAI has provided a list of 80 potential training sites, along with information about value chains that are likely to be served in each site. We assigned 60 percent of these sites (48 communities) to the treatment group and 40 percent (32 communities) to the control group.[2]

To ensure representation of all value chains and regions in both the treatment and control groups, we conducted random assignment within value chain and region (North, Center, and South). In addition to ensuring balance across treatment and control groups, stratifying random assignment by value chain and region improves statistical power. In some sites, DAI expects to offer trainings serving more than one value chain. Sites with more than one value chain were grouped into separate strata for random assignment (Table 1). For example, Northern sites identified as candidates to host trainings for both tree fruits and vegetables were grouped together (in one stratum), whereas Northern sites serving only the vegetables value chain were a separate stratum. Random assignment was then conducted within stratum (that is, within region and value chain group), with the assignment ratio as close to 60:40 as possible (with some differences due to rounding). Some of the strata defined by region and value chain had too few sites for random assignment to be valid within that stratum alone. We combined smaller strata as necessary to ensure each stratum had enough sites for random assignment to be meaningful.

[2] These proportions were chosen to balance the implementation needs of DAI (having a sufficient number of locations to work in for all value chains) with those of the evaluation (having a sufficient number of control areas that are not too close to training sites).

Table 1. Strata for Random Assignment.

Region	Value Chain(s) Identified for Training	Total Sites	Treatment Sites	Control Sites
North	Tree Fruits only	12	7	5
North	Vegetables only	4	2	2
North	Tree Fruits and Vegetables	9	6	3
Center	Table Grapes (with or without Tree Fruits and/or Vegetables)	9	5	4
Center	Vegetables only	11	7	4
Center	Tree Fruits (with or without Vegetables)	8	5	3
South	Table Grapes (with or without Tree Fruits and/or Vegetables)	16	10	6
South	Vegetables only	5	3	2
South	Tree Fruits (with or without Vegetables)	6	4	2

In treatment sites, DAI will be able to offer any training activities (including farmer field schools, classroom trainings, and demonstration plots) for any value chains in which it would like to work (even if those value chains were not originally specified in the list). In control sites, DAI will not offer any training activities of any type for any value chains over the entire length of the project.[3,4] This is not to say, however, that trainings will not actually occur in control sites. Other providers of training may provide trainings (covering related or unrelated topics) in control sites (and in treatment sites as well); this is beyond the control of DAI. (In fact, the project is designed to spur private actors to deliver training.) If private providers want to offer trainings in control sites, that need not change for the evaluation. The provision of private training could, however, affect the estimated impacts of ACED training. The evaluation intends to measure the impacts of ACED training relative to services that farmers would have access to in the absence of the ACED training subactivity. Therefore, the evaluation will aim to distinguish between private trainings spurred by the ACED activity and private trainings unrelated to ACED. Examining participation in non-ACED training at follow-up, especially for farmers in control sites, will help to assess the availability of outside training. Comparing follow-up and baseline data will help to disentangle private trainings that ACED training might have spurred from private training that would have been available whether ACED were implemented or not.

[3] For these purposes, training activities are defined as activities that are conducted in a particular location and involve more than one organizational participant. Exceptions to the control sites' participation include national-level activities (such as a Buy Moldovan campaign) and activities that involve only one organizational participant at a time, such as technical assistance to a single commercial farm or other value chain company, cooperative, or association. This means that all region- and *raion*-level training activities will be part of the randomization.

[4] DAI can offer training activities in a fourth group of sites—those that were not involved in the randomization and were not on the A list. However, these sites must be located sufficiently far from control sites; ideally, additional sites would be at least 15 kilometers from any control sites. Exceptions to this could be made with approval from USAID.

2. Estimating Impacts

As described earlier, we will estimate impacts using an IV approach, which is typically estimated in two stages. In the first stage, we will measure the impact of distance from a training site on participation in training:

$$(1) \quad t_{ijk} = \alpha + \eta d_{jk} + \delta y_{ijk,pre} + \varphi X_{ijk} + \pi Z_{jk} + \lambda_k + \mu_{jk} + \varepsilon_{ijk}$$

where t_{ijk} is a binary indicator for participation in a ACED-provided training for individual i in community j in random assignment stratum (region and value chain) k; d_{jk} is the distance from community j in stratum k from the nearest ACED training site;[5] $y_{ijk,pre}$ is the outcome of interest (for example, farm income) for the same individual at baseline; X_{ijk} and Z_{jk} are vectors of baseline individual and community characteristics, respectively, that can be related to the outcome of interest;[6] λ_k is a stratum fixed effect, estimated using a vector of binary indicators, one for each random assignment stratum (value chain and region); μ_{jk} is a community-specific error term; and ε_{ijk} is an individual-specific error term. The parameter η is the impact of distance on the probability of participation in training, which we expect to be negative (farmers located closer to training sites are more likely to participate). We can use parameter estimates from this regression to estimate a predicted probability of participating in training, $\hat{t}_{ijk}$.[7]

In the second stage, we will estimate the following regression model:

$$(2) \quad y_{ijk,post} = \phi + \beta \hat{t}_{ijk} + \rho y_{ijk,pre} + \xi X_{ijk} + \psi Z_{jk} + \theta_k + \nu_{jk} + \gamma_{ijk}$$

where $y_{ijk,post}$ is the outcome of interest for individual i in community j in stratum k at follow-up, $\hat{t}_{ijk}$ is the predicted value from the estimation of Equation (1), and all other variables are as defined previously, with associated parameters analogous to those in Equation (1).[8]

---

[5] Distance could be measured as either estimated travel time or geographic distance. Travel time would be the preferred approach but is harder to measure; we will work with MCA-Moldova to determine if this would be measurable. We will also explore the sensitivity of our results to a discrete measure of distance, obtained by coding distance into intervals and using the set of binary interval indicators in the regression. This alternative approach is less parametric than using the continuous distance measure.

[6] The specific baseline variables included as controls will be systematically selected by first choosing a small set of priority outcomes (training participation as well as the outcomes highlighted in the power calculations below) and then iteratively determining which baseline variables have the strongest predictive power for those priority outcomes, conditional on the other control variables specified above.

[7] This framework could be expanded to include other measures of access to ACED training activities; using other measures of access to training as instruments for participation would be valid if those instruments are plausibly exogenous, that is, unrelated to outcomes except through their impacts on training participation. In our analysis, we will explore whether there are other plausible instruments for participation.

[8] The standard errors in Equation (2) must be adjusted to account for the estimation error in the predicted value of t_{ijk}. Standard statistical software packages perform this adjustment automatically when both stages of the estimation are specified. In addition, standard errors should be clustered by community in order to take into account the community-specific error terms.

The parameter of interest in Equation (2) is β, which gives the estimated impact of participating in an ACED-provided training on individual outcomes. Our estimate of β is likely to change depending on the round of data; mid-term impacts might be smaller than longer-run impacts if impacts take some time to materialize. On the other hand, if impacts fade over time, mid-term impacts could be larger than longer-run impacts.

If distance from training is not a good predictor of participation, it will be difficult to measure the impacts of the ACED training subactivity. The IV approach relies on the premise that distance to a training site will be related to participation in training. If farmers readily travel to training sites in distant locations—so much so that distance does not predict training participation—the IV approach will break down; however, as long as there are differences in the probability of participation (even if some farmers living near control sites participate in training), the evaluation is likely to be viable.

Therefore, any outreach efforts to farmers in control communities should be carefully considered to meet implementation needs, while at the same time maintaining the viability of the evaluation. Specifically, although the evaluation design does not preclude the recruitment of farmers in control communities for participation in training, it is essential that any recruitment efforts maintain a distinct difference in participation rates between treatment and control communities. For example, indirect recruitment efforts in control areas (such as those occurring through media and producer associations) or direct invitations to a small number of prominent farmers are likely to maintain a difference in participation rates. On the other hand, aggressive recruitment efforts in control areas, such as door-to-door recruitment or busing, may shrink the difference in participation rates to such an extent that the evaluation breaks down.

Finally, our estimation strategy focuses on t_{ijk} as a binary indicator for training, so that we are effectively comparing those who undertake any training to those who undertake no training. However, training participants may participate to varying degrees; for example, some farmers might participate in several training activities, whereas others might attend only one relatively short training. It would therefore be interesting to estimate the impact of different degrees of participation in training on outcomes. Unfortunately, our strategy of randomizing the location of trainings only gives us (partial) experimental control over the participation decision itself (which we exploit in the IV estimates); we have no way of controlling the degree of participation for training participants. Therefore, because farmers participating more intensively in the trainings might have different characteristics compared to those who participate less intensively, we will not be able to interpret differences in outcomes between these groups as the *impact* of greater participation. Nevertheless, we will still examine the relationship between outcomes and the degree of participation in order to provide some suggestive evidence.[9]

[9] Measures of the degree of participation could include the number of trainings or the number of hours spent in training. Since these measures may not be accurate if they are self-reported, we intend to obtain them from DAI's MIS by matching survey respondents to DAI's training lists (if feasible).

3. Interpretation

a. Local Average Treatment Effect (LATE)

If the impact of the training is identical (homogeneous) across all farmers in the population, then the parameter β in Equation (2) can be interpreted as the Average Treatment Effect (ATE). This is the effect that the average farmer (or indeed, any farmer) would experience if he or she undertook training. However, the assumption of homogeneous impacts may not be plausible in practice; different farmers may experience different (heterogeneous) impacts depending on their specific characteristics. Some of these characteristics may be observable (for example, farm size), while others may be unobservable (for example, farmer motivation). Since training is optional, farmers with certain characteristics that would lead them to expect large impacts from the training may be more likely to participate. The impact estimate β could then be picking up differences in farmer characteristics rather than the impact of the training. For example, if training is more useful for highly-motivated farmers, these farmers might select into training. Conversely, if highly-motivated farmers have already learned the skills taught as part of the training through other training programs or contacts, they may have less to gain from attending the training.

If impacts are heterogeneous, then β can no longer be interpreted as the ATE due to selection into training. However, under certain assumptions, Imbens and Angrist (1994) show that β still identifies a type of average treatment effect but only for a specific population.[10] In particular, it identifies a Local Average Treatment Effect (LATE), which is the treatment effect for the farmers who undertake training if it is offered nearby, but not if it is offered far away ("marginal takers"). This is a subgroup of all the farmers in the population (Table 2); the other subgroups are farmers who would not take up training even if it were nearby ("never takers") and farmers who would take up training even if it were far away ("always takers"). These other subgroups of farmers may experience different impacts if they underwent training, and since we cannot identify these impacts we cannot compute the average impacts across all farmers and identify the ATE. Nor can we identify the effect of "treatment on the treated" (TOT), the impact on those who are actually trained, since this is composed of the impact for marginal takers and always takers and we cannot identify the latter.[11]

Although the LATE applies to a specific subgroup of the population, it is a parameter of considerable policy interest since it gives the effect for the marginal takers, the farmers who are induced to undertake training by this type of intervention and would not undertake it otherwise. These are the farmers who are materially affected by the intervention. The impact on always takers is less relevant, since these farmers would likely identify other training opportunities or sources of information regardless of the project. Similarly, the impact on never takers is of less policy interest, since training cannot be made mandatory and these individuals would never experience it in practice. Were the training to be scaled up in Moldova, or were a similar program implemented in another country, the marginal takers are more likely to represent the farmers who would participate in the expanded program.

[10] The key assumptions are that the instrument (distance from training) only affects outcomes through training, and that being located closer to training increases the probability of training.

[11] If no farmers in the control group take up treatment, then there are no always takers and the LATE can be interpreted as the impact of treatment on the treated.

Table 2. Subgroups of Farmers in the Population

Subgroup Type	Description	Status in Treatment Communities	Status in Control Communities
Marginal Takers	Take up training if it is close (treatment), but not if it is far away (control)	Trained	Not Trained
Always Takers	Take up training regardless of distance	Trained	Trained
Never Takers	Do not take up training regardless of distance	Not Trained	Not Trained

Note: For illustrative purposes, this table presents a simplified scenario in which all treatment communities are near training and all control communities are far. In reality, some control communities will be farther from training than others.

Nevertheless, it is still useful to understand the extent to which the LATE differs from the ATE for the population as a whole and from the TOT. For example, this may be important in computing ex post ERRs. We will address this using several approaches. First, we will estimate the fraction of marginal takers by comparing take-up rates in treatment and control communities and examining how baseline characteristics relate to take-up. We cannot identify specific individuals as marginal takers, but we can identify the size and characteristics of this group. If a large majority of farmers fall into this group, the LATE is likely to be closer to the average effect for all farmers. Second, since the issues of interpretation arise only due to heterogeneous effects, we will investigate whether heterogeneity is present. This could involve a formal test, such as that suggested by Heckman, Urzua, and Vytlacil (2006); if the test indicates that heterogeneity is not present, then the parameter β can be interpreted as the ATE.[12] Alternatively, we could look at differences in impacts by subgroups defined by observable characteristics related to outcomes (such as farm size); heterogeneity in impacts along observable characteristics is would suggest possible heterogeneity along unobservable characteristics as well. Third, we will attempt to use the weighting approach of Heckman, Urzua, and Vytlacil (2006) to convert the estimated LATE into an ATE, though this approach has strict data requirements that cannot be tested until we have complete data. Finally, though we cannot directly estimate the treatment effect for always takers, we can develop reasonable bounds that can inform us about the ATE. Lee (2009) develops a method that provides informative "tight" bounds with reasonable assumptions. By assessing the extent of differences between the impact for marginal takers (given by the LATE) and the bounds for always takers (calculated using Lee's approach), we will have some sense of the extent to which the LATE differs from the TOT.

b. Intent-to-Treat (ITT)

In addition to the LATE, we can estimate another treatment parameter, namely the "intent-to-treat" (ITT). This parameter is estimated as the average difference in outcomes between all targeted farmers in treatment and control communities, including all training participants and non-participants. In a regression framework, the ITT can be computed using Equation (2), only replacing $\hat{t}_{ijk}$ with an indicator for the farmer's community's random assignment status. The

[12] This involves testing whether outcomes are linear in the propensity score, which is the predicted probability of training as a function of the instrument.

ITT is driven by both the impacts of the training and differences in take-up rates for the treatment and control groups, and can be interpreted as the average impact on the target population of living near a training site. From a policy perspective, the ITT is of interest since it gives the impact of introducing training, which will necessarily involve non-universal take-up. The ITT has the advantage of having the same interpretation regardless of any heterogeneity in impacts, since it ignores selection into treatment; however, it does not allow us to draw any conclusions about the impact of the trainings on those actually participating in training. The ITT will also be of greater interest if we find that treatment and control assignments are strongly predictive of training participation, that is, if farmers are unlikely to attend training unless it is located nearby.

The ITT and LATE estimates therefore apply to different populations and have different interpretations; together they will provide complementary evidence on the impact of the training subactivity. We will report both estimates in our reports, together with a detailed discussion of their interpretation, in order to provide a comprehensive impact evaluation.

4. Assessment of Ripple Effects

A key feature of DAI's implementation plan is the intention for DAI's training lessons to reach untrained farmers through farmer networks and other information sharing. These ripple effects could mean that nonparticipating farmers could benefit from the ACED training activities, even if they do not participate directly.[13]

It is difficult to incorporate these ripple effects directly into our IV framework, since we only have experimental control over the physical location of the trainings and not over farmer information networks. Instead, the evaluation will take a two-step approach to assess the presence and magnitude of ripple effects. Initially, we will assess if there is preliminary evidence to support the existence of these effects. Then, if such evidence exists, we will attempt to quantify the magnitude of ripple effects to complement our overall impact analysis.

We will use a triangulation approach to initially assess the plausibility of ripple effects. First, we will examine whether farmers are connected to networks through which they could have plausibly learned the practices taught as part of the training without attending training themselves. For example, farmers who rely on cooperatives or on relatives in other communities for farming information might be more likely to experience these ripple effects; the Farm Operator Survey (FOS), described in Section D, will gather data on sources of farming information to inform this analysis. Second, we will consider the simple pre-post change in practice adoption among non-participants. If practices promoted by the DAI training activities are adopted by nonparticipants, then this provides suggestive evidence that ripple effects are plausible; it would not be conclusive, however, as this could also simply reflect broader trends in adoption. The FOS will gather data on practice adoption to inform this analysis. By focusing on specific practices promoted by the DAI training activities—rather than practices that other programs might promote—we will be able to

[13] There may also be "vertical" ripple effects from other ACED subactivities not targeted directly at farmers. For example, firms affected by the ACED project may communicate information about the nature and benefits of product upgrading down the value chain until it reaches farmers, encouraging them to upgrade their products. Although our assessment of ripple effects is focused on training (the subject of the evaluation), we will consider asking in the Farm Operator Survey about information received from firms and conducting descriptive or correlational analyses to relate this to practice adoption.

better assess the plausibility of ripple effects. Finally, we will use evidence from the qualitative assessment, particularly evidence related to practice adoption, to assess whether and how ripple effects occurred.

If our preliminary analysis suggests that ripple effects are plausible, we will attempt to quantify their magnitude by analyzing information networks in greater detail. Together with information on adoption of specific practices promoted by DAI training activities, the FOS will gather data on the name and location of the person or entity from whom the farmer learned about each of these practices. We assume that these self-reports will usually be reliable, given that farmers will have adopted these practices relatively recently and should therefore remember from whom they obtained the information. If feasible, we will attempt to link the names of these informants to the lists of farmers trained by DAI and will estimate the extent to which untrained farmers adopted new practices learned from training participants, including the number and types of practices adopted in this way. We will also be able to distinguish between within-community and across-community spillover effects. If it is not feasible to link specific names to the training lists, we will conduct a similar (albeit slightly weaker) analysis focusing on the location of the informants in order to analyze the extent to which untrained farmers adopted specific practices from farmers in treatment communities (who are more likely to have attended training than other farmers). We will also be able to compare the magnitude of ripple effects in different rounds of the survey; we expect that they may emerge to a greater degree in later survey rounds.

Our estimates of the magnitude of ripple effects will not be as conclusive as our main impact estimates, since they are not based on an experimental design. Moreover, we will have limited ability to completely map information networks and estimate higher degree spillovers (for example, farmer A attends training and informs farmer B about the practice, who in turn informs farmer C). Nevertheless, these analyses will provide important information on the existence and magnitude of ripple effects, which will be informative for the program itself as well as for the interpretation of the impact estimates.

5. Gender Analysis, Beneficiary Analysis, and Complementarity with ISRA-CISRA

As noted earlier, MCC is particularly interested in disaggregating program impacts by gender. Our ability to precisely estimate these disaggregated impacts is largely limited by the available sample sizes.

For individual-level outcomes recorded for both the farm operator and his or her spouse, we will be able to calculate disaggregated impact estimates with similar statistical power to the estimates for the full sample because most households will provide one observation for each gender.[14] Key outcomes of this nature include participation in training and adoption of practices, since one of the ACED activity's main goals is to have men and women benefit equally from the project. Therefore, we expect to address questions such as:

- How does participation in training activities differ between male and female household members?

[14] The original FOS asked virtually all questions of both the farm operator and his or her spouse. However, this resulted in a very lengthy survey; we will therefore restrict our surveying of both genders to specific domains in which differences by gender are of particular interest.

There are also other outcomes available for both genders which are not explicitly targeted by the ACED activity, but for which evidence on impacts may be useful in order to provide additional context for the overall findings and to inform future program design. This set of outcomes includes outcomes reported for the farm operator and his or her spouse, such as community leadership and access to credit, as explained in Section D. It also includes outcomes that will be reported only by the farm operators, but for other household members as well. For these outcomes—such as employment, non-farm income, and ownership/control of productive assets—we will also be able to measure impacts separately for male and female household members with reasonable precision. (For non-farm income, the high variance of income will limit statistical precision.) In addition to disaggregating impacts by gender, we will also answer descriptive questions, such as the agricultural credit constraints facing women in Moldova. Therefore, the set of questions that we expect to address in order to provide general evidence on gender dynamics in the ACED activity context include the following:

- How does non-farm income generated by female household members change as a result of the ACED training subactivity?[15]
- What are the constraints for women regarding access to finance?
- Has women's leadership in the community increased as a result of the ACED training subactivity?
- Has the ACED training subactivity increased female ownership/control of productive assets in the household?

On the other hand, we will have very limited ability to disaggregate impacts for household-level outcomes by the gender of the household head, the primary farm operator, or the person owning or controlling productive assets in the household, because we expect that relatively few households will have women in these roles. It will therefore be difficult to obtain statistically significant answers to questions such as the following:

- Is the impact of the ACED training subactivity on farm income different for male- versus female-headed households? For male- versus female-operated farms?
- Is the impact of the ACED training subactivity on agricultural practice adoption different for male- versus female-operated farms?
- Is the impact of the ACED training subactivity on farm income different based on the gender of the person controlling/owning productive assets in the household?

Nevertheless, we will still present these estimates as providing suggestive evidence on differences in impacts by gender.

In a similar vein, MCC is interested in disaggregating any estimated impacts by other beneficiary groups. We will conduct subgroup analyses that explore the following research question:

[15] Non-farm income will be disaggregated by household member only for categories for which respondents can report income separately by household member. For example, it might be straightforward for farm operators to attribute non-farm wage income to a particular household member, but more difficult to do so for remittances received by the household.

- Is the impact of the ACED training subactivity on farm income different for households that had different levels of income or farms of different sizes at baseline?

MCC has also expressed interest in understanding interactions between ISRA-CISRA and the training component of the ACED activity and, in particular, measuring whether the impacts of ISRA-CISRA combined with training exceed the impacts of ISRA-CISRA alone. For this comparison to be viable, we would have to be able to identify a sufficient number of ISRA-CISRA beneficiaries who have also participated in the ACED training subactivity and a sufficient number who have not.

Initially, ISRA-CISRA and the ACED training subactivity were expected to affect many of the same farmers over the same time period—in fact, so much so that it was unclear if there would be ISRA-CISRA beneficiaries who did not also participate in ACED trainings. However, current ACED training implementation plans suggest that there will in fact be very little overlap between the sites where training will be offered as part of the ACED training subactivity and the sites that will benefit from ISRA-CISRA, at least until several years into the program (after irrigation rehabilitation is complete). Though ACED trainings might be offered in the 11 CIS areas selected for rehabilitation and some farm operators may travel from other communities to attend training, only 10 of the 80 communities selected at the outset for the ACED training impact evaluation are located in these 11 CIS areas.[16] As a result, it will be difficult to measure interactions between the ACED training subactivity and ISRA-CISRA. DAI is also expected to conduct additional trainings (focused on irrigation use) inside the 11 CIS areas selected for ISRA-CISRA. However, those trainings will not be part of this impact evaluation because DAI will likely advertise those trainings to all farmers in the 11 areas (with the exception of ACED training control communities), making it impossible to rigorously estimate the impacts of those training activities relative to a counterfactual.

6. Statistical Power

To determine the size of the effects that we will be able to detect given our anticipated sample size, we computed minimum detectable impacts (MDIs)—the smallest impacts that our design will be able to statistically distinguish from zero. The MDIs depend on the sample size (both the number of training sites and the number of survey respondents), assumptions on key parameters (such as the intracluster correlation coefficient and the regression R-squared), the power with which we would like to detect effects (typically 80 percent), and the variance of the outcome. We have also accounted for the expectation that some farmers in treatment sites will choose not to participate in training while some farmers in control sites will receive training.[17] Table 3 shows MDIs for several outcomes of interest, including farm income. We calculated these MDIs using parameter estimates obtained from existing data and previous studies.

[16] These 10 overlapping communities are located in 8 of the 11 CIS areas; another CIS area contains one of the A-list communities.

[17] We made a simplifying assumption on the nature of the instrument in order to simplify the MDI calculations. Specifically, the calculations assume that the instrument for training receipt is a binary indicator for whether the farmer lives near a treatment site, rather than the community's distance from the nearest training site. In this simplified framework, we assume that 70 percent of treatment group farmers and 10 percent of control group farmers attend training.

The MDI for farm income is about $5,354, or about 109 percent of the anticipated baseline average (Table 3). This means that the evaluation will be able to detect only very large changes in farm income (a doubling or more), whereas the projected increase in farm income is about 20 percent. This high MDI is driven by the fact that farm income is highly variable (based on data from the 2008 Moldovan FOS).[18] Because of the limited ability to detect impacts on farm income, we will consider alternative approaches to provide evidence of impacts on this outcome.

A possible solution is to focus on a closely related outcome, the gross margin per hectare (defined as the value of output minus the cost of inputs, excluding non-monetary inputs, capital depreciation, and fixed costs of operation), possibly restricted to high-value crops. The advantage of this outcome measure is that it is the component of farm income that the program is most likely to affect directly, and yet it has much less variability than total farm income. The MDI for gross margin per hectare for HVA crops is $594, or about 27 percent of the anticipated baseline average.

Another alternative to measuring the impact on farm income is to investigate the impact of the ACED training subactivity on household consumption, which can serve as an alternative measure of well-being. MDIs on consumption tend to be smaller, as consumption is less variable than farm income. The MDIs show that, for the full sample, we would be able to detect an increase of about 16 percent in household consumption (with parameter estimates based on data from a similar setting in Armenia) (Table 3). Two important caveats, however, are that consumption is less proximate to the subactivity and that it can be difficult to measure accurately.

In addition to final outcomes such as farm income, we are interested in detecting impacts on intermediate outcomes that might reflect improvements in practices before final impacts materialize. Not only are impacts on intermediate outcomes likely to be larger than impacts on final outcomes, but the intermediate outcomes on which we would focus also tend to be less variable. Table 3 presents the MDIs for a hypothetical intermediate outcome—an indicator for whether a farmer adopts a particular practice covered in the training. If we assume that 20 percent of farmers have already adopted this practice at baseline, we will be able to detect a change of about 9.7 percentage points (49 percent) given our sample size.

[18] Particularly low and high outliers drive some of this variation; excluding these outliers reduces the standard deviation of farm income, improving statistical power. In the impact analysis, we will carefully analyze data from all outliers to ascertain whether values are accurately reported. Even if these values are legitimate, we can consider whether it makes sense to restrict the analysis to centrally distributed values.

Table 3. Minimum Detectable Impacts for the ACED Training Subactivity Impact Evaluation

	Sample Size				Minimum Detectable Impacts (As Percentage of Baseline Mean)			
	Treatment Sites	Control Sites	Treatment Sample	Control Sample	Annual Farm Income [Dollars]	Gross Margin per Hectare, HVA crops [Dollars]	Annual Consumption [Dollars]	Practice Adopted [% Points]
Full Sample	48	32	1,636	1,091	5,354 (109%)	594 (27%)	827 (16%)	9.7 (49%)
Subgroup (50 percent)	48	32	818	545	6,395 (131%)	709 (32%)	988 (19%)	11.6 (58%)
Subgroup (20 percent)	48	32	327	218	8,809 (180%)	977 (44%)	1,361 (27%)	16.0 (80%)

Sources: Authors' calculations using data from the 2008 Moldovan Farm Operator Survey (FOS) and the 2008 Armenia Tertiary Canal Survey (TCS). Data on farm income (per household) and gross margin per hectare (HVA crops) are from small and medium farms in the 2008 FOS. Data on consumption (per household) are from the 2008 TCS. Practice adoption parameters are assumed.

Note: MDIs are for a two-tailed test with 80 percent power and a 95 percent level of significance using the formula:

$$MDI = \frac{2.8\sigma}{f_T - f_C}\sqrt{\rho(1-R^2)\left(\frac{1}{N_T}+\frac{1}{N_C}\right)+(1-\rho)(1-R^2)\left(\frac{1}{rn_T}+\frac{1}{rn_C}\right)}$$

where σ is the baseline standard deviation of the outcome, f is the fraction of farmers attending training from the treatment or control sites, R^2 is the regression R-squared, ρ is the intracluster correlation coefficient, N is the number of treatment or control sites, n is the number of sampled farm operators in the treatment or control sites in the evaluation (excluding farm operators in the 8 purposely selected sites), and r is the survey response rate. The calculations assume a regression R-squared of 0.4, an intracluster correlation coefficient of 0.05, and a 75 percent survey response rate. As noted in the text, this formulation makes the simplifying assumption that treatment (treatment/control) is binary rather than continuous (distance from treatment). We assume that 70 percent of eligible farmers in treatment sites and 10 percent of eligible farmers in control sites attend the training. Baseline means (and standard deviations) are assumed to be \$4,866 (\$21,983) for farm income, \$2,220 (\$2,438) for gross margin per hectare (HVA crops), \$5,093 (\$3,396) for consumption, and 20 percent (40 percent) for a hypothetical farming practice. All foreign currencies were converted into U.S. dollars using nominal exchange rates at the time the surveys were conducted.

Finally, as mentioned earlier, MCC is also interested in analyses for subgroups—for example, those defined by gender. However, we will have limited power for these analyses (Table 3). For example, for a subgroup comprising one-fifth of the full sample, the MDI for farm income increases by more than 60 percent. Because some subgroups of interest might be even smaller (for example, women exclusively operated only about 3 percent of farms in the 2008 Moldovan FOS), the associated MDIs could be even higher. Therefore, although we will report impacts by subgroup when possible, the estimates will be imprecise.

7. Qualitative Analysis

We will supplement our quantitative impact analysis with qualitative analysis of focus group discussions with farmers targeted for the trainings in treatment and control communities. Qualitative analysis has the potential to provide important insights—such as how well the trainings worked in practice and what the key facilitators and barriers to the adoption of new practices are—that could not be obtained through quantitative analysis alone. By identifying themes and patterns in the responses and categorizing and sorting them appropriately, a set of key findings will be developed.

Results from the qualitative analysis will be used for three main purposes. First, they will provide input for program monitoring, course corrections, and future program design by identifying key challenges and barriers to successful trainings in the context of a value chain project. Second, they will help inform the interpretation of the quantitative data analysis by helping to understand why impacts did or did not occur. For example, if adoption of practices covered in training is higher in some treatment communities compared with others, focus groups could highlight some of the key differences in trainings that might explain this. Finally, the qualitative results will be helpful in understanding the extent to which ripple effects occurred within and across communities, and the mechanisms through which any ripple effects occurred. This will both inform the quantitative analysis and provide important information on the operation of the program.

We plan to include two rounds of qualitative data collection in the evaluation. The first round should be conducted in late 2013 (after the trainings have been completed in the treatment areas) and the second in late 2015 (after some time has passed since the trainings were completed). These rounds of data collection are timed to meet the purposes of the qualitative analysis described earlier, including providing input for monitoring and course corrections at a critical phase of the project (shortly after completion of the trainings), informing the interpretation of the quantitative data analysis (at both midterm and post-Compact follow-ups), and providing evidence on ripple effects (at the post-Compact follow-up).

We anticipate that the qualitative information we require might be collected and analyzed by MCA-Moldova and/or USAID as part of the overall evaluation of the ACED activity. If this is the case, we would attempt to leverage this data collection effort in order to avoid duplication of effort. This may require additional discussions with MCA-Moldova and USAID in order to ensure that the information relevant for this evaluation is obtained. Specifically, we would have to ensure that the focus group protocols capture information relevant to the training component of the ACED activity (for example, practice adoption) and that the focus group sample includes a representative sample of relevant value-chain participants in the research communities (for example, by selecting at least one treatment and one control community from several strata used for randomization and using the FOS sampling frame to ensure that farm operators with a range of characteristics are included). If it is determined that separate qualitative data collection and analysis is required for this evaluation, we anticipate that MCA-Moldova will play the lead role, with Mathematica providing input and

feedback on the protocols, on the terms of reference for the data collection contractor (if required), and on the summary of findings.

The focus groups will concentrate on farmers' experiences with training and the adoption of new agricultural practices. We will finalize topics for these focus groups through discussions with MCC, MCA-Moldova, and USAID, but the topics are likely to include the following:

- Awareness of the trainings and reasons for participation or nonparticipation
- Attitudes towards the training and main sources of satisfaction and dissatisfaction
- Suggestions for improvements in trainings in the future
- Sources of information and mechanisms of adoption of new agricultural practices (especially those covered in trainings) and barriers to practice adoption
- Constraints to benefitting from new practices and how these have changed over time

8. Supplemental Analyses

In addition to the main impact analysis, we will conduct several exploratory analyses to interpret and frame our findings. These analyses will include the following:

- **Outcome analysis in A list sites.** Though we will not include the A list sites in our main impact evaluation, we will use data from communities near A list sites to describe the outcomes of farmers living near those sites and compare their outcomes with those of farmers living near treatment sites. Through this analysis, we will assess whether outcome differences might be due to differences in training implementation across A list and treatment sites and whether there are differences (at baseline) in the characteristics of farm operators served in these A list sites and those served in treatment sites.
- **Analysis of impacts by program intensity.** ACED program beneficiaries may participate in ACED activities in varying degrees (including training and other ACED activities). Therefore, we will conduct descriptive analyses summarizing participation rates and the intensity of engagement with the ACED program, including the training subactivity and other subactivities to the extent possible; this will help us to interpret our impact findings. We will also look at the relationship between outcomes and the degree of participation, although, as noted earlier, we will not be able to interpret differences between these groups as the *impact* of greater participation.
- **Impact analyses by value chain.** In addition to the subgroup analyses discussed earlier, which focused on differences by gender and other beneficiary groups, we will also explore differences in impacts by value chain, broadly defined (for example, vegetables or tree fruits). This will provide suggestive evidence on which value chain trainings were more effective than others, although these estimates may be imprecise due to sample size limitations and could be driven partly by selection of different types of farmers into

training in different value chains.[19] Because of the sample size limitations and the substantial variability of some key outcome measures (such as farm income), the impact estimates that are conducted separately by value chain will focus on intermediate outcomes, such as adoption of specific practices, which will have less variability and so will be more informative.

- **Interaction with other value chain constraints.** The impact evaluation described here measures the impact of ACED-facilitated trainings in an environment in which other value chain constraints are concurrently being addressed. However, these other constraints may not be addressed to a similar extent across all sites and could affect the estimated impacts of the trainings. In order to investigate this issue further, we will attempt to obtain information from DAI on market conditions and other activities likely to affect producers near training sites. Although we will have very limited ability to quantitatively assess the relationship between measures of other constraints and training impacts due to small sample sizes, a descriptive analysis of this information will provide useful context to help us interpret our impact estimates. The extent to which other value chain constraints affect farm operator behavior can also be assessed as part of the qualitative analysis.

D. Data Requirements and Sources

The impact evaluation will draw on three key sources of data. The first is longitudinal survey data from farm operators living near treatment and control sites that will enable us to track outcome changes over time. For the impact analysis, these survey data will be linked to a second source, implementation data about ACED training activities—such as locations, value chains and topics covered, and dates. The final source is the qualitative focus group data, which were discussed in Section C. We describe below the survey and implementation data in greater detail.

1. Survey Data from a Targeted Sample of Value Chain Farm Operators

To check the validity of our design and to implement our regression model, we require data on outcomes and characteristics of farm operators near treatment and control sites. Specifically, we require these data at the following points in time:

- **Baseline.** Baseline data on outcomes, such as farming practices and farm income, as well as farm operator characteristics, such as gender, education, and farm size, will enable us to control for baseline differences between the treatment and control areas, improve the statistical precision of the impact estimates, and test the validity of our design.
- **Midterm follow-up.** Data on the farm operators surveyed at baseline, collected two years after baseline, will enable us to measure the short-run impacts of the program. We expect the main impacts at this point in time to be on training participation and practice adoption, but would still analyze impacts on other outcomes as a benchmark for the final impact estimates.

[19] Exploring differences in impacts based on the strata used for randomization will be difficult, as the large number of strata means that the estimates will likely be too imprecise to be useful. Therefore we will focus on disaggregating impact estimates by value chain.

- **End-of-Compact follow-up.** Data collected at the end of the Compact period for the farm operators surveyed at baseline will enable us to estimate the impacts of the program after three years have elapsed and most ACED training activities have concluded. However, because some investments (such as adoption of new varieties) require several years to show returns, impacts on final outcomes such as farm income might not yet have materialized. Therefore, we will obtain information on more proximate intermediate outcomes in addition to information on final outcomes. As noted previously, this round may be dropped in order to concentrate data collection resources on the other three rounds, which we consider most important.
- **Post-Compact follow-up.** By surveying the same farm operators after the Compact period, we hope to obtain data on final outcomes after sufficient time has passed for impacts on these outcomes to materialize. However, because some investments (such as adoption of new varieties) require several years to show returns, it is still possible that impacts on final outcomes such as farm income might not yet have materialized at this time. Therefore, we will obtain information on more proximate intermediate outcomes in addition to information on final outcomes.

We intend to obtain these farm operator data through several rounds of the FOS, which we will use for both the ISRA-CISRA and ACED evaluations. Though there might be some overlap in the areas included in the evaluations, each will draw a separate survey sample.[20] The sample frame for the ACED FOS survey sample will include farm operators in targeted value chains who live near treatment and control sites.[21] In communities selected for the survey, we will not draw a random sample of farmers; instead, we will draw a random sample of farmers working in the targeted value chain (for example, tree fruits or vegetables). Therefore, the sample aims to be representative of farmers in targeted value chains living near treatment and control sites. As a result, sample members will generally be HVA farmers, even before training activities have taken place.[22]

In addition to collecting data on farmers in treatment and control communities, MCA-Moldova (through its data collection contractor) will also collect data on farm operators in A list communities. Though we will not include farmers in these communities in our primary impact analysis, as noted in the previous section, we can still use data from these communities for descriptive analyses to examine how key outcomes compare between farmers living near these sites and those near treatment sites. Although this supplemental analysis would not be as rigorous as the core impact analyses, it can still inform us about whether the program beneficiaries or possible effects of the program have substantially differed for farmers near these A list sites compared with those served by treatment sites.

---

[20] The FOS sample for this evaluation will also differ from that in the pre-Compact FOS, since the evaluation focuses specifically on individuals in the treatment and control communities. The two samples will therefore likely have little overlap.

[21] In most cases, these farmers will come from the same community as the treatment site; however, if there are few farmers in the training site community (for example, if the site is a *raion* center), we might sample farmers from neighboring communities. For the purposes of the evaluation, communities will be defined as incorporated or unincorporated localities/villages.

[22] The sampling strategy will be finalized in consultation with the data collection contractor, MCC, MCA-Moldova, USAID, and DAI.

Four rounds of the FOS are planned: baseline (2011–2012), midterm follow-up (2013–2014), end-of-Compact follow-up (2014–2015), and post-Compact follow-up (2015–2016). The baseline round is timed to occur before farmers will have implemented lessons learned in training (though it might occur after some trainings have taken place). Subsequent rounds of the survey (the midterm follow-up, the end-of-Compact follow-up, and the post-Compact follow-up) will capture outcomes over time. In each round, we expect that the ACED sample of the FOS will cover approximately 3,000 households in approximately 88 communities—48 communities near treatment sites, 32 communities near control sites, and 8 communities near the A list sites.

We will administer the FOS to the same individuals in each survey round. Relative to repeated cross-sectional data (in which different individuals are surveyed in each round), longitudinal data collection has several advantages. First, with longitudinal data, we can control for individual baseline characteristics—most importantly, the baseline levels of the outcomes—increasing statistical precision and controlling for any chance differences between the treatment and control groups. Second, longitudinal data ensure that the population studied is consistent over time; repeated cross-sections, on the other hand, have the drawback that the population studied could change, making it difficult to disentangle program impacts from compositional changes. With these advantages in mind, we will work with MCA-Moldova and the data collection contractor to develop plans to follow participants over time.

The FOS questionnaire will cover several key domains (Table 4). These include household characteristics, which we will use as control variables in the regression analysis, and a variety of outcome measures. The outcome measures include the main final program outcomes (production, sales, and farm income), but also several intermediate outcomes (such as participation in training activities, adoption of practices, and other agricultural investments) that might reflect improvements in HVA production.

Another important intermediate outcome is product upgrading—the extent to which new practices translate into products that better meet end-market requirements—since this is a key channel that is hypothesized to improve farm income. Our primary measure of product upgrading will be the market price received for crops produced. Even if it takes several seasons for changes in final outcomes to become apparent—so that they would not be captured in the post-Compact follow-up survey—the various intermediate outcomes might nevertheless reveal program impacts.

The survey will also include outcomes for which differences by gender could be of particular interest, such as asset ownership and community participation, as discussed in Section C. To that end, we will ask a subset of questions separately of the farm operator and the farm operator's spouse. Although we will have limited ability to examine these differences due to low statistical power, we will still collect the necessary information to explore them in our analysis.

The FOS will be used for the ISRA-CISRA evaluation as well as the ACED evaluation. For most domains, the FOS questionnaire will be the same for the ISRA-CISRA and ACED samples. However, there will be several key differences. We will use the irrigation management, satisfaction, and usage module for the ISRA-CISRA sample but not for the ACED sample. Instead, the ACED sample will be asked more extensive questions about participation in training, access to training, and practice adoption. In addition, some items will be asked of both the farm operator and his or her spouse for one sample but will be asked of only the farm operator for the other sample.

Table 4. Domains and Key Items in the FOS

Domain	Key Items
Household Characteristics	Demographic and socioeconomic information on all members of the household (**), such as age, gender, relationship to head of household, and education level; migration of family members (**)
Economic Activity	Employment (formal and informal) (**); occupation (**); number of hours worked (including on farm activities) and earnings of the members of the household (**); household decision making on agricultural activities (*)
Assets	Ownership, control, and use of productive assets (*)
Farm Production, Revenue, and Inputs	Major crops cultivated and associated value; livestock; land holdings (owned, rented); farm inputs; farm size; rent payments to landowners; wages paid to laborers; household farm labor (**); level of investment in the farm or home
Participation in Training	Extent and intensity of participation in agricultural training (**); provider of training (*); location of training (*); topics covered in training (*); satisfaction with any training received (*)
Access to Training	Awareness of agricultural training (*); impediments to participating in training (*); desired level of training (*)
Community	Women's leadership in community (*), including changes over time
Technological Level of Production	Agricultural and product upgrading practices, very specific (*); marketing practices (*); availability of other agricultural inputs; access to credit (**), technical assistance, transportation, and markets; participation in farming cooperatives
Information Networks	Sources of different types of information on agricultural practices and markets; awareness and adoption of specific agricultural practices covered by training; sources of information about these specific practices, including name and location
Income and Consumption	Farm income; other sources of income (for example, remittances, government transfers, sales of nonagricultural products, rental payments received, and so on); spending on categories of goods and services

(*) = Asked separately of the farm operator and his or her spouse.

(**) = Asked in a roster of all household members completed by the farm operator.

Not starred = Asked only of farm operator.

2. Implementation Data

The evaluation also requires information about program implementation. Linking data from the FOS to information about trainings conducted is critical for the impact analysis; for example, to measure distance to the nearest training site, we require information about locations of trainings. Implementation data will also enable us to explore findings from the impact analysis (for example, to analyze the existence and extent of ripple effects) and will inform the development of data collection

materials.[23] For each training conducted—whether at an A list site, a treatment site, or another site[24]—DAI will be asked to provide certain key details:

- Location
- Date
- Topics covered/key practices promoted
- Value chains served
- Participants' names and where they are from
- Intensity of training (for example, number of hours)

DAI will provide this information to coincide with the timing of the various rounds of the FOS.

E. Key Challenges and Risks

The evaluation faces several challenges and risks. These potential challenges and our plans to address them include the following:

- **This evaluation is not designed to measure the impacts of the *full* ACED activity.** As noted earlier, the ACED training subactivity is only one component of the ACED activity. Impacts measured through this evaluation will provide evidence on the marginal impact of training *when administered in conjunction with a range of other value chain interventions.* Importantly, the evaluation will not be able to assess whether the full ACED activity is effective.
- **The instrument used for the IV approach might be weak.** We will be able to measure the impacts of ACED training subactivity as long as farmers living near treatment sites are *more likely* to attend ACED training activities than farmers living near control sites. However, if farmers who live closer to treatment sites are not more likely to attend training, the evaluation design is not viable.
- **Training could have an impact on baseline survey responses.** Training activities are scheduled to begin in winter 2011–2012; the baseline survey will not be administered until spring 2012. Though impacts on practices and agricultural outcomes are not likely to materialize until well after training occurs (that is, until the next agricultural season, at the earliest), there is a risk that training will influence responses to the baseline FOS, even if behavior has not yet been influenced. To mitigate this risk, we will use implementation data provided by DAI to assess the degree to which baseline responses might be influenced by training activities by separating farmers who were exposed to training before the baseline survey from those who were not.

[23] Information about topics covered might be important for ensuring that the survey instruments cover appropriate topics and for interpreting the impact findings. Furthermore, some ACED trainings might have different anticipated impacts than others. For example, we might not expect trainings focused on pesticide safety to generate increases in income.

[24] DAI may provide training in other sites as long as those sites are sufficiently distant from control sites.

- **High attrition between baseline and follow-up surveys could reduce sample sizes.** In the FOS, we plan to survey the same farm operators in all survey rounds because controlling for baseline differences, especially in the key outcome measures, improves statistical power. However, there is a risk that we will not be able to locate all baseline farm operators in subsequent rounds—for example, some farm operators might have moved or be unwilling to participate—and we will have to drop these respondents from the analysis sample. To mitigate this risk, MCA-Moldova will select a survey firm with a record of and a strategy for obtaining high response rates and we will help ensure that the contact details (such as cell phone numbers) of surveyed farm operators are collected as part of the survey, as well as the details of a close contact (such as a neighbor) who will be able to assist us in locating the baseline respondent. The midterm follow-up survey will also be a useful way of staying in touch with respondents and updating their contact details so that they are easier to track for the final follow-up.

- **There might be insufficient time to detect impacts on final outcomes.** Though ACED training activities will take place starting in 2011, for many of the practices promoted by training (such as the adoption of new varieties), it will be several years before returns accrue to farmers. Therefore, when we report the post-Compact results, we will emphasize that our estimates may not fully reflect long-term impacts and that there might be some long-term impacts we cannot capture without subsequent data collection efforts. We will attempt to address this by estimating impacts on intermediate outcomes (for example, adoption of practices) that might be more responsive during the Compact period but are likely to be related to eventual impacts on final outcomes.

- **Statistical power is limited due to the high variance of outcomes (especially farm income).** As discussed earlier, the high variance in farm income means that the estimated impact of the ACED training subactivity on farm income will be especially imprecise. To address this, we will consider using other related outcomes, such as gross margin per hectare and consumption.

Overall, although we cannot entirely eliminate these challenges and risks to the evaluation, we believe that our plans will mitigate them to the extent possible.

F. Reporting

Our evaluation reports will follow the timing of the FOS baseline and follow-up surveys and will occur between 2011 and 2016 (Figure 1). The FOS baseline will occur shortly after ACED training activities have begun, but before they have had a chance to influence agricultural outcomes (2011–2012). The baseline report will include summary statistics for key characteristics and outcomes. To assess the validity of our design, the analysis will compare farmers living near treatment sites and farmers living near control sites along these dimensions.[25] A midterm FOS follow-up will capture outcomes two years later (2013–2014); the midterm report will present descriptive and impact analyses based on the midterm FOS.

[25] As noted earlier, producing a baseline report might require modification of our contract because our original scope of work did not specify a baseline report.

The end-of-Compact and post-Compact reports will conduct descriptive and impact analyses based on the respective follow-up FOS rounds. In 2015, we will prepare an end-of-Compact evaluation report based on the analysis of end-of-Compact data. Finally, in 2016 we will prepare a post-Compact report that will examine the impacts of ACED training activities after four years have elapsed. All of the follow-up reports (midterm, end-of-Compact, and post-Compact) will include tables and figures with key descriptive statistics, but will focus on impact estimates using the IV approach.

Figure 1. Schedule of ACED Training Subactivity Evaluation Activities, by Study Year and Quarter

	2012				2013				2014				2015				2016			
	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
Data Collection and Processing** (Farm Operator Survey, ACED Sample)	B	B	B					M	M			E*	E*			P	P			
Analysis and Reporting				B	B					M	M			E*	E*			P	P	
Qualitative Data Collection and Reporting***								M								P				

* The end-of-Compact round of the FOS may be removed.

** Managed by MCA-Moldova.

*** Managed by USAID and/or MCA-Moldova.

B = baseline; M = midterm follow-up; E = end-of-Compact follow-up; P = post-Compact follow-up.

References

Heckman J. J., Urzua S. S., and E. J. Vytlacil, "Understanding Instrumental Variables in Models with Essential Heterogeneity", NBER Working Paper 12574, October 2006.

Imbens, G.W. and J. D. Angrist, "Estimation of Local Average Treatment Effects", *Econometrica*, vol. 62, no. 2, 1994, pp. 467-475.

Lee, D. S., "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects", *Review of Economic Studies*, vol. 76, no. 3, 2009, pp. 1071-1102.

MATHEMATICA
Policy Research

www.mathematica-mpr.com

Improving public well-being by conducting high-quality, objective research and surveys

Princeton, NJ ■ Ann Arbor, MI ■ Cambridge, MA ■ Chicago, IL ■ Oakland, CA ■ Washington, DC