Preliminary Report on the Fourth Round of Data

NORTHERN GHANA MILLENNIUM VILLAGES PROJECT

Date: **June 2017**

Authors: **Edoardo Masset with the assistance of Jorge Garcia-Hombrados**

Submitted by Itad
In association with:

Results in development

Institute of Development Studies

Participatory Development Associates Ltd.

Acknowledgements

This report has been prepared by the team for the impact evaluation of the Millennium Villages Project. The team is composed of staff from Itad, the Institute of Development Studies (IDS), the London School of Hygiene and Tropical Medicine (LSHTM), and Participatory Development Associates (PDA) Ltd in Ghana. The team is fully independent of the Earth Institute and the Millennium Promise. The principal author of this report is Dr Edoardo Masset, with analysis support by Jorge Garcia-Hombrados and additions by Dr Chris Barnett and Kelsy Nelson. The report was edited and proofread by Etty Payne. The team is nonetheless very grateful to all the researchers that have assisted with data collection, the staff at DFID, and everyone else that has provided support, information and comments – including the work of the Earth Institute during the enumeration phase. The findings of this report are the full responsibility of the authors, and any views contained in this report do not necessarily represent those of DFID or of the people consulted.

Citation

Masset, E., Garcia-Hombrados, J. 2017. *Northern Ghana Millennium Villages Project: Preliminary Report on the Fourth Round of Data*. Itad, Hove.

Disclaimer

UKaid
from the British people

Acronyms and abbreviations

CEW	Community Education Workers
CHW	Community Health Worker
CV	Control Village
DD	Difference-in-Difference
DFID	UK Department for International Development
GES	Ghana Education Service
GLSS	Ghana Living Standards Survey
GSS	Ghana Statistical Service
MDG	Millennium Development Goals
MV	Millennium Village
MVP	Millennium Villages Project
NGO	Non-Governmental Organisation
PRG	Peer Review Group
PTA	Parent-Teacher Association
SADA	Savannah Accelerated Development Authority
SMC	School Management Committees
WASH	Water, Sanitation and Hygiene

Contents

List of tables

List of figures

Executive summary

The Millennium Villages Project (MVP) has been designed to demonstrate how an integrated approach to community-led development can translate the international Millennium Development Goals (MDGs) into results. The project in northern Ghana is one of several projects instigated over the past 10 years, which together are set to reach nearly half a million people across 10 countries in Africa. Central to the MVP approach is the synergistic value of integrated community-based investments, focused on scientifically proven interventions, delivered simultaneously rather than as one-off investments. The premise is that a critical platform of basic needs must be reached before economic development can really take off. The project in northern Ghana has been running in three districts since May 2012, investing over £11 million on health, education, agriculture and infrastructure interventions in 35 communities, and reaching around 30,000 people.

In 2011, the UK Department for International Development (DFID) commissioned an impact evaluation of the MVP in northern Ghana. The original design foresaw three rounds of data collection: baseline (2012), midline (2014) and endline (2016). In response to discussions with the Peer Review Group (PRG), two smaller household survey rounds were later added to the design for the "interim" years (2013, 2015). These additional rounds focus on a subset of modules such as demographics, consumption and expenditure – and allow the evaluation team to achieve greater statistical power for these variables in order to detect relatively small effects of the intervention (Brown et al., 2013). The baseline data was collected in 2012 and the midterm data in 2014, and include an analysis of the full dataset alongside qualitative assessments (Masset et al., 2014 and Masset et al., 2016).

This report summarises the fourth year of data collection undertaken in 2015. The scope of this report is more limited than the Baseline and Midterm Reports given the data collection during the interim years, but similar to 2013. The fourth survey round collects a reduced set of information at all levels (households, individuals and villages) and there is no qualitative follow-up. The findings presented here focus only on those key variables that are covered across all four rounds of data collection.

The findings from the fourth survey round are summarised below.

- As in previous rounds, attrition rates resulting from households leaving their homes, household dissolution or refusal to conduct interviews were minimal. In addition, the rates are similar in the Millennium Villages (MVs) and control villages (CVs), thus excluding the presence of attrition bias.
- As in previous rounds, household composition underwent considerable changes. The changes, however, do not bias the comparisons between MVs and CVs and are not different from changes normally observed in similar household surveys.
- Participation rates in project activities remain high, particularly in the health sector. Progress was also made in forming groups through social mobilisation. Farmers are involved in project activities, while education interventions do not appear to reach a large fraction of the population.
- Participation rates have remained stable or slightly decreased in other sectors.
- Participation rates are uncorrelated with baseline poverty levels; which confirms that the project does not target or reach particular sectors of the community and provides services to all households.
- Poverty (measured by consumption expenditure) decreased in both the MVs and CVs in comparison to the baseline. However, poverty reduction in MV areas was not larger than poverty reduction in CV areas, suggesting the project did not have an impact on poverty or extreme poverty.

- The agricultural interventions had a large impact on incomes. They increased farmers' incomes by at least 50% on a per capita basis.
- The change in income did not translate in an increase in consumption expenditure. The analysis suggests that the additional income was invested in savings, assets and durable goods.
- The project had a very modest impact on primary school attendance.
- The data provide some support to the hypothesis that the project encourages migration into the MV area from neighbouring communities and discourages migration to outside communities in comparison to migration flows in CV areas.

We also conducted exploratory analysis in order to explain: (a) the discrepancy about the MVP's impact on income and consumption; and (b) the sources of positive impact on agricultural incomes.

To address the first issue, we investigated whether households are saving or investing gains in income. To do so, we reviewed the structure of household savings in the MVs and found that in terms of expenditure on durable goods (which is not normally included in consumption for poverty calculations), wealth is predominantly and increasingly held in the form of animal holdings and, to a lesser extent, in financial savings and investment in durable goods. This behaviour fits well with the permanent income hypothesis model of consumption behaviour. Households are perceiving the income gains resulting from the project as temporary, rather than permanent. Instead of spending their income and increasing consumption, they are saving additional income.

The second issue is addressed by decomposing the MVP's impact on agricultural outputs resulting from increased use of agricultural inputs, increased productivity of the inputs and an increase resulting from other interventions such as improved health due to fewer missed work days. The vast majority of the project impact is the result of increased use of seeds and fertiliser that was donated by the MVP or provided at highly subsidised rate. This result raises concerns about whether it will be possible to maintain achievements in agricultural outputs when the MVP ends and subsidised inputs are discontinued.

1. Introduction

This report analyses the impact of the Millennium Villages Project (MVP) using the fourth round of data collection from 2014–15. The analysis follows the Analysis Plan[1] approved by the Peer Review Group (PRG). We start by discussing the data quality and its suitability for a difference-in-difference (DD) analysis. We then illustrate the MVP's impact on selected indicators for which data are available. Finally, we conduct an exploratory analysis on the causes of the differential impact of the intervention on income and consumption and of the determinants of impact on agricultural production. This report does not disaggregate the analysis by distance from the Millennium Village (MV) cluster and therefore does not investigate spill-over effects; these will be comprehensively analysed in the endline evaluation.

In the last round of statistical analysis, the Midterm Report[2] found that the project had moderately positive impacts so far: the project did not have a sizeable impact on the MDGs, although several impacts were visible across other well-being indicators. For instance, the project showed a large increase in per capita income mostly driven by an increase in agricultural incomes. The increase in incomes appears to be supported by a large improvement in self-reported food security. Food security is measured by reporting the number of days in which the household did not have enough food to meet family needs. Since household expenditure includes home-produced food, an improvement in food security should be matched by an increase in food expenditure. Yet, household expenditure increased to a lower extent, which could be due to income gains being saved or invested – although the qualitative (participatory rural appraisal and reality check approach) studies did not reach the same conclusion. In this report, we explore further our understanding of both savings and household wealth (Chapter 4) and farm productivity (Chapter 5). In terms of health, the Midterm Report also found that there was a considerable improvement in the nutritional status of children under five. Plus, it produced a modest reduction in mortality rates – although there are difficulties calculating these effects after only two years of data – and, the project did not affect prevalence of anaemia. Health variables are not covered in this report, as this data are collected next in the final year of the evaluation. Other findings from the Midterm Report included a modest increase in school attendance rates and no improvements on the quality of education as measured by mathematics and English test scores, and no difference in impacts by sex. Similarly, no differences in impact by sex emerged on nutritional status and anaemia. This report considers further exploratory analysis of educational impacts (in Chapter 3), as well as education gender parity and the impact in female-headed households (Chapter 6).

1.1 Data collection and changes in household composition

The expenditure and income data quality from previous rounds was analysed using Benford tests and was found to be comparable to similar budget surveys conducted in Ghana by the Ghana Statistical Service (GSS). Thus, the same analysis will not be repeated in this report. The validity of a DD approach, however, requires more than just data quality as it rests on the assumption that the project and comparison groups are similar. This report repeats the analysis of changes in the composition of project and comparison groups produced by attrition, migration and measurement error, which was conducted in the previous report using midterm data. If attrition, migration and measurement error affect the two groups in different ways, the DD analysis can be biased. Note that while migration is considered here as a threat to the internal validity of the study, Section 3.5 below looks at migration as an *outcome* of the intervention. Opinions may differ with regard to whether the intervention is more likely to increase or to decrease migration and to whether an increase in migration should be

[1] The Analysis Plan is available at: www.ids.ac.uk/publication/northern-ghana-millennium-villages-impact-evaluation-analysis-plan

[2] The Mid-term Report is available at: www.ids.ac.uk/publication/millennium-villages-evaluation-midterm-summary-report

perceived as a positive or negative outcome. However, as discussed in the baseline report, it is clear that migration is an outcome[3] of the intervention that needs to be considered because of its welfare implications.

1.2 Panel structure of the data

The baseline survey targeted a sample of 755 households in the MVs and 1,496 households in the control villages (CVs). However, the baseline sample comprises of only 711 MV households and 1,461 CV households because not all the target households were found at the time of the interviews. All the following survey rounds targeted the same 755 MV households and 1,496 CV households originally selected at the baseline. The fourth round interviewed 97% of the households originally targeted at the baseline, confirming high completion rates observed in previous years. There is no discernible difference in completion rates between the MV and CV areas (Table 1).

Table 1. Completed household interviews during the first four rounds

Sample	Target	2012	2013	2014	2015
MV interviews	755	711	743	735	731
%		94	98	97	97
CV interviews	1,496	1,461	1,487	1,456	1,446
%		97	99	97	97
All interviews	2,251	2,172	2,230	2,191	2,177
%		96	99	97	97

The fourth round did not achieve the remarkable 99% completion rate obtained during the second round (the highest achieved); nevertheless the 2015 rate remains very high at 97%. Table 2 reports the reasons for not completing the interviews across rounds. Because the number of non-completed interviews is small, it is difficult to summarise this information in a meaningful way. Reasons vary from difficulties locating dwellings, since they are sometimes vacated, or a competent household member is unavailable for the interview. However, these events occur in different ways over the survey years and no clear pattern can be discerned across survey rounds.

Table 2. Reasons for not completing the interviews (numbers) in MV/CV[4]

Reason	2012 (No.)	2013 (No.)	2014 (No.)	2015 (No.)
No competent household member at home	21	1	8	13
Entire household absent	22		11	20
Interview postponed	10			
Interview refused	1			
Partly completed				
Dwelling vacant or destroyed		4	2	20
Dwelling not found	19	9	13	5
Household has relocated		6	15	8
Household dissolved or deceased		1	6	
Other	6		4	7
All	79	21	59	73

[3] The extent to which this is an intended outcome of the intervention (and thus part of the MVPs design) is debateable.
[4] Note that information on the reasons for not completing the interview is missing for one household for the year 2014 (midterm) and year 2015 (interim).

Household migration appears to occur in small numbers. The surveys record the new location of migrated households by interviewing neighbours of the household not found. The results are reported in Table 3. In 2014 and 2015, all migration occurred towards locations outside the district of origin and much migration was directed to the Ashanti region, particularly to Kumasi and gold mining areas.

Table 3. Households reported as "relocated" (numbers) in MV/CV

Location	2013	2014	2015
Kumasi, Ashanti	1	2	4
Jagsi, Kumasi, Ashanti			1
Delaasa, Kumasi, Ashanti			1
Ejisu, Ejisu-Juaben, Ashanti			1
Sariba, Northern, West Mamprusi,	1		
Obuasi, Ashanti	2		
Luisa, Builsa, Upper East	1		
Accra, Greater Accra		7	
Kentasi, Ashanti		1	
Presetia	1		
Eastern Region			1
Missing	0	5	0
All	6	15	8

Due to not all target households being interviewed every year and because the households that are not interviewed differ from year to year, the number of full panel households that are interviewed each year decreases over time. The decrement is, however, very small (Table 4). The absolute numbers of households not interviewed are so small (24 households in MV areas and 50 households in the CV areas, compared to the baseline) that a comparative analysis of the characteristics of engaged and unengaged households in the MV and CV areas is hardly feasible. In addition, these are not proper "unengaged" households, but simply households that could not be interviewed in all three survey rounds following the baseline. The attrition rate of 2015 versus the baseline survey, that is the percentage of households interviewed at baseline that were not interviewed at the fourth round, is just 4.2% (this is 1-2080/2172 from last row of Table 4).

Table 4. Completed household interviews during the first four rounds

Sample	Target	2012	2013	2014	2015
MV panel interviews	755	711	707	697	689
%		94.2	93.6	92.3	91.3
CV panel interviews	1,496	1,461	1,454	1,424	1,391
%		97.7	97.2	95.2	93.0
ALL panel interviews	2,251	2,172	2,161	2,121	2,080
%		96.5	96.0	94.2	92.4

Although the rate of household attrition is very small, attrition among individuals is quite large (Table 5). Only 87.5% of the individuals originally enumerated at the baseline were again enumerated after three years. These attrition rates are again very similar in MV areas (87.0%) and CV areas (87.8%), suggesting that they are not the result of a household's disaffection with the study, nor are they a product of migration. If households and individuals were more reluctant to be interviewed in the CVs (say, because they are not benefiting from the programme), then the attrition rates should be higher in CVs, which does not appear to be the case. High rates of individual attrition in MVs and CVs are the result of natural processes (i.e. deaths and births) and migration.

Table 5. Panel of individuals across surveys

Sample	2012	2013	2014	2015
MV individuals	5,231	5,576	5,854	6,021
MV panel		4,930	4,654	4,550
%		94.2	89.0	87.0
CV individuals	10,337	10,649	11,023	11,255
CV panel		9,869	9,378	9,072
%		95.5	90.7	87.8
Al individuals	15,568	16,225	16,877	17,276
All panels		14,799	14,032	13,622
%		95.1	90.1	87.5

Changes in household composition were analysed at great length in the Midterm Report (2016) in light of the anthropological literature on household structure in West Africa. Anthropologists (Hill, 1986) have questioned the validity of the "household" concept traditionally used by economists and pointed to the high mobility and fluidity of populations across what we define as "household units". The main conclusions of the midterm analysis were the following:

- There is large variation in reported household size across surveys, though the numbers are similar to those observed by other surveys in other contexts.
- There are considerable errors by respondents and enumerators in listing household members. The errors are observed and corrected during each survey round by verifying whether individuals previously listed as members are actually members and whether there are household members that were missed out in previous survey rounds. We call these errors Type I and Type II errors, respectively.
- Finally, at the midterm, fewer than 5% of listed household members are reported as living in the household for less than 12 months during the preceding year, which contrasts with the high fluidity of household membership depicted by anthropologists.

In this report, the analysis of changes in household composition conducted at the midterm is repeated using the balancing population equation (Preston et al., 2001). There are only two ways for people to enter a population: being born or migrating into it. Similarly, there are only two ways for leaving a population: death and out-migration. The population at any time is therefore the result of changes in the natural population growth (births minus deaths) between t and t-1, and changes in net migration (in-migrants minus out-migrants) between t and t-1, to which we add reporting errors: a reduction in population resulting from Type I errors in the previous survey and an increase resulting from Type II errors.

$$N_t = N_{t-1} + \left(B_{t-1,t} - D_{t-1,t}\right) + \left(I_{t-1,t} - O_{t-1,t}\right) + \left(e2_{t-1,t} - e1_{t-1,t}\right)$$

The breakdown of population change in natural growth, migration and errors is shown in Table 6 for MV and CV areas separately, using only those sample households (2,177) that were interviewed for three consecutive rounds. Rates of change are calculated over the number of household members in the previous round. The overall natural population change is positive and appears to be slightly larger in the MV areas than in the CV areas. The percentage population changes reported in Table 6 suggest this is the result of a larger natural change (birth minus deaths) and a larger net migration (in-migrants minus out-migrants). For both the MV and CV areas, the number of household members moving out is larger than the number moving in. This is expected as individuals move out of households as they form new households or simply migrate. Movements in and out of households are slightly higher in the MV areas compared to CV areas, but the differences are very small. This provides further evidence

that the project is not substantially changing migration patterns in either direction. Type I errors (wrongly enumerated household members) appear more numerous than Type II (household members missed out). Both errors appear to lessen over time as household membership is probed by repeat household visits. We also included an additional category to account for a particular type of error regarding individuals who have lived in the household during two consecutive surveys, but in the first survey were not considered a household member because they had not lived there for more than six months. The residual unexplained change is small but points to the fact that successive enumerations and our data analysis were unable to account for all reporting and enumeration errors. A final analysis of household changes, their determinants and measurement error will be conducted at the endline. What is important to note at this stage is that errors and unexplained differences are not substantially different in MV and CV areas.

Table 6. Population changes across surveys (per cent)

Population change	**All (%)**			**MV (%)**			**CV (%)**		
	2013	*2014*	*2015*	*2013*	*2014*	*2015*	*2013*	*2014*	*2015*
Overall change	-3.6	0.9	-0.0	-3.6	0.6	-0.3	-3.5	1.1	0.1
Natural change (births, deaths)	0.7	1.6	1.1	0.9	1.6	1.3	0.6	1.5	0.9
Migration: people moving in	0.6	1.3	1.0	0.8	1.4	0.9	0.6	1.2	1.0
Migration: people moving out	-4.9	-7.5	-7.6	-5.6	-7.8	-8.3	-4.5	-7.4	-7.3
Type I error	-3.4	-1.9	-0.6	-4.6	-2	-0.4	-2.7	-1.8	-0.7
Type II error	0.2	1.8	1.7	0.3	2	1.6	0.1	1.6	1.7
New members that previously lived with household	1.1	2.8	3.2	1.5	2.9	3.3	0.9	2.7	3.1
Residual unexplained difference	2.1	2.8	1.2	3.1	2.5	1.3	1.5	3.3	1.4

2. Participation in project activities

As in previous evaluation reports, we present levels of involvement in project activities tracked by the evaluation team. The Midterm Report found that level of involvement in MV areas was greater than in the CVs for some activities. In particular, the midterm data showed greater involvement in agricultural and health-related activities, while involvement in social mobilisation activities was modest.

This new round of data shows involvement in social mobilisation activities is high (Table 7). In particular, the project helped form groups that were for the most part absent in the areas before the MVP commenced. When compared to the CVs, participation was higher in MVs for all groups apart from the MDG School Clubs. The difference between MV and CV areas was particularly large and statistically significant for cooperatives, farmer-based organisations, women's groups, parent-teacher associations (PTAs) and village savings and loan associations. On the other hand, the promotion of water, sanitation and hygiene (WASH) groups, Millennium Development Goals (MDGs) school clubs and "Daddy's clubs", as well as farmer field schools (FFS), and school management committees does not appear to be very successful. Though most differences in participation between MV and CV areas are statistically significant for these groups (with the exception of MDGs school clubs and daddy's clubs), participation is very small and never above 3% of the population.

Table 7. Participation in groups supported by the project

Social mobilisation groups	MV	CV	P-value	Observations
Cooperative	22.6***	3.5	0.000	2,178
Farmer-based organisation	17.0***	6.0	0.000	2,178
Farmer field school	3.1*	1.4	0.063	2,178
Women's group	27.2***	14.4	0.001	2,178
Parent-teacher association	59.5***	39.3	0.000	2,178
WASH	1.1**	0.1	0.027	2,178
MDG school club	0.0	0.0	-	2,178
Water and sanitation development board	1.0**	0.1	0.016	2,178
Mother-to-mother support group	3.4***	0.1	0.002	2,178
Daddy's club	0.3	0.1	0.322	2,178
Village savings and loan association	14.8***	3.1	0.000	2,178
School management committee	1.4**	0.2	0.015	2,178

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance

The level of involvement in agricultural-related activities remains high (Table 7), particularly the membership of cooperatives and farmer-based organisations. More farmers use fertiliser in the MV areas than the CV areas, and farmers are not reporting a greater number of loans (Table 8). Although the involvement in agricultural activities is high, it appears there was a reduction when compared to previous years.

Table 8. Participation in project-related agricultural activities

	MV	CV	P-value	Observations
Any household member received a loan	1.4	0.6	0.190	2,178
Used any fertiliser	40.8*	32.6	0.072	2,134

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance

Participation in health-related activities is very high in MVs compared to CVs (Table 9). Apart from visiting a health clinic, the difference between MV and CV areas is always very large and statistically significant, including those that did not show any difference during previous survey rounds (e.g. deworming and vitamin A). Greater participation in health services related to home visits by community health workers and limited difference in access to health facilities was already observed in previous rounds and suggests that while the project is successful in increasing the supply of health services it is less successful at creating demand for them. The MVP's construction and rehabilitation of health clinics and health sensitisation campaigns should lead to higher use of health facilities resulting in a higher number of reported visits to health facilities by households. It should be noted that the fraction of households visiting a health facility is already high in the MVs (above 70%) and not knowing the health needs of the population it is impossible to make an assessment about the gap between visits to health clinics and residents' health needs. Unfortunately, we do not have baseline data on attendance rates of health facilities to estimate changes. At least three explanations for the lack of difference in the number of visits to health facilities are possible. First, the MV has largely increased usage of clinics but the same has occurred in CV areas because of the operation of other non-MV interventions. Second the MV increased the number of visits (but not having baseline data is impossible to tell) to achieve the same (higher) level of clinic usage prevailing in the CV areas. Third, individuals residing in CV locality in the proximity of the MV intervention have access to facilities in MV areas.

Table 9. Participation in health-related activities

Participation	MV	CV	P-value	Observations
Membership of NHIS	74.1***	52.5	0.000	15,605
Someone distributed bed nets	45.4***	23.2	0.000	2,178
Visit by a community health worker (CHW)	89.3***	23.4	0.000	2,178
The CHW provided condoms	4.8***	0.8	0.001	2,178
The CHW measured children's arms	69.6***	11.1	0.000	2,178
CHW advised on breastfeeding	71.1***	16.5	0.000	2,178
CHW advised on child feeding	78.1***	18.5	0.000	2,178
CHW advised on use of bed nets	78.4***	18.5	0.000	2,178
Visited a health facility	74.1	69.7	0.104	2,178
Children given deworming	43.1***	18.9	0.000	2,178
Children given vitamin A	45.4***	27.3	0.000	2,178
Children given food supplements	2.6**	1.2	0.015	2,178

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance

Involvement in education-related activities is slightly larger in MV areas compared to CV areas (Table 10). More children report having a school meal in MV areas, but the difference is not statistically significant. A large and statistically significant difference is found in the number of children receiving

incentives to attend school through donations of uniforms, stationery and other items. Similarly, a much larger number of children are reported to be receiving sanitary pads distributed by the MVP to increase girls' school attendance.

Table 10. Participation in education-related activities

Participation	MV	CV	P-value	Observations
Child had a school meal on previous day	39.2	28.1	0.176	2,178
Children received a bursary	0.7	0.6	0.869	2,178
Children received stationery, uniform, etc.	26.1***	13.4	0.000	2,178
Child given a sanitary pad	13.3***	2.3	0.000	

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance

The overall narrative offered by these data sets is that the project has successfully increased efforts to promote social mobilisation in MV areas, leading to the formation of new groups, such as women's groups and savings groups, and to an increase in overall participation in other groups. The project has also been very successful in the provision of health services by reaching a large majority of the attended population, including for services that were neglected in previous years, such as deworming and vitamin A. On the other hand, involvement in agricultural activities has decreased, particularly the provision of loans, and to some extent the use of fertiliser. Finally, involvement in education-related activities is similar to previous years and not very high.

In order to further investigate changes in levels of involvement over time we plot involvement rates over time for all the activities for which data are available (Figures 1 and 2). Since the impact of the project on involvement is the difference between participation rates in project and control areas, the charts also report the trend in this difference in order to detect patterns over time. Note also that in some cases the rates are reported since baseline (2012) while in others they are reported from the second round (2013). This is because the survey did not track all involvement activities from the baseline since many of them were unknown when it was conducted.

We are interested in detecting patterns and understanding whether levels of involvement increase or decrease over time. Participation may increase as activities expand and successful results are observed through demonstration or by word of mouth. Participation, however, may also decrease. Several factors may lead to decreased project participation after initial enthusiasm such as project fatigue in delivering outputs, multiplicity of project goals and lack of success. The difficulty of integrated rural development projects such as the MVP in managing multiple and sometimes conflicting goals has long been recognised in the development literature (see for example ODI, 1979).

The charts in Figures 1 and 2 show a few patterns and suggest that involvement in health-related activities has increased over time; involvement in education activities has remained stable; while involvement in agriculture-related activities has slightly decreased. MVP project managers have suggested that the decline in fertiliser use after 2014 is likely to be the result of the project being unable to guarantee loans for farmers affected by negative weather shocks during a bad agricultural season, which resulted in several farmers not accessing the loans. We do not observe a general decrease in involvement over time except for the small decrease in agricultural activities. In particular, we observe the following:

- The charts show a difference of about 20 percentage points between MV and CV areas in membership of farmer-based organisations, cooperatives and using fertiliser. Involvement in these activities appears to increase up until the midterm and then starts to decrease. The decrease is particularly pronounced in regard to accessing loans.

- There is a difference of less than 10 percentage points between MV and CV areas with respect to access to school feeding and provision of stationery and uniforms. There is no obvious pattern over time.
- Involvement in health-related activities has clearly increased throughout the period. Not only have visits by community health workers (CHWs) increased, but so has the provision of health services such as deworming, vitamin A and bed nets. For some of these activities (breastfeeding advice, deworming, vitamin A distribution and advice on the use of mosquito bed nets) the difference between MV and CV areas has increased over time.

Figure 1. Participation in project activities over time

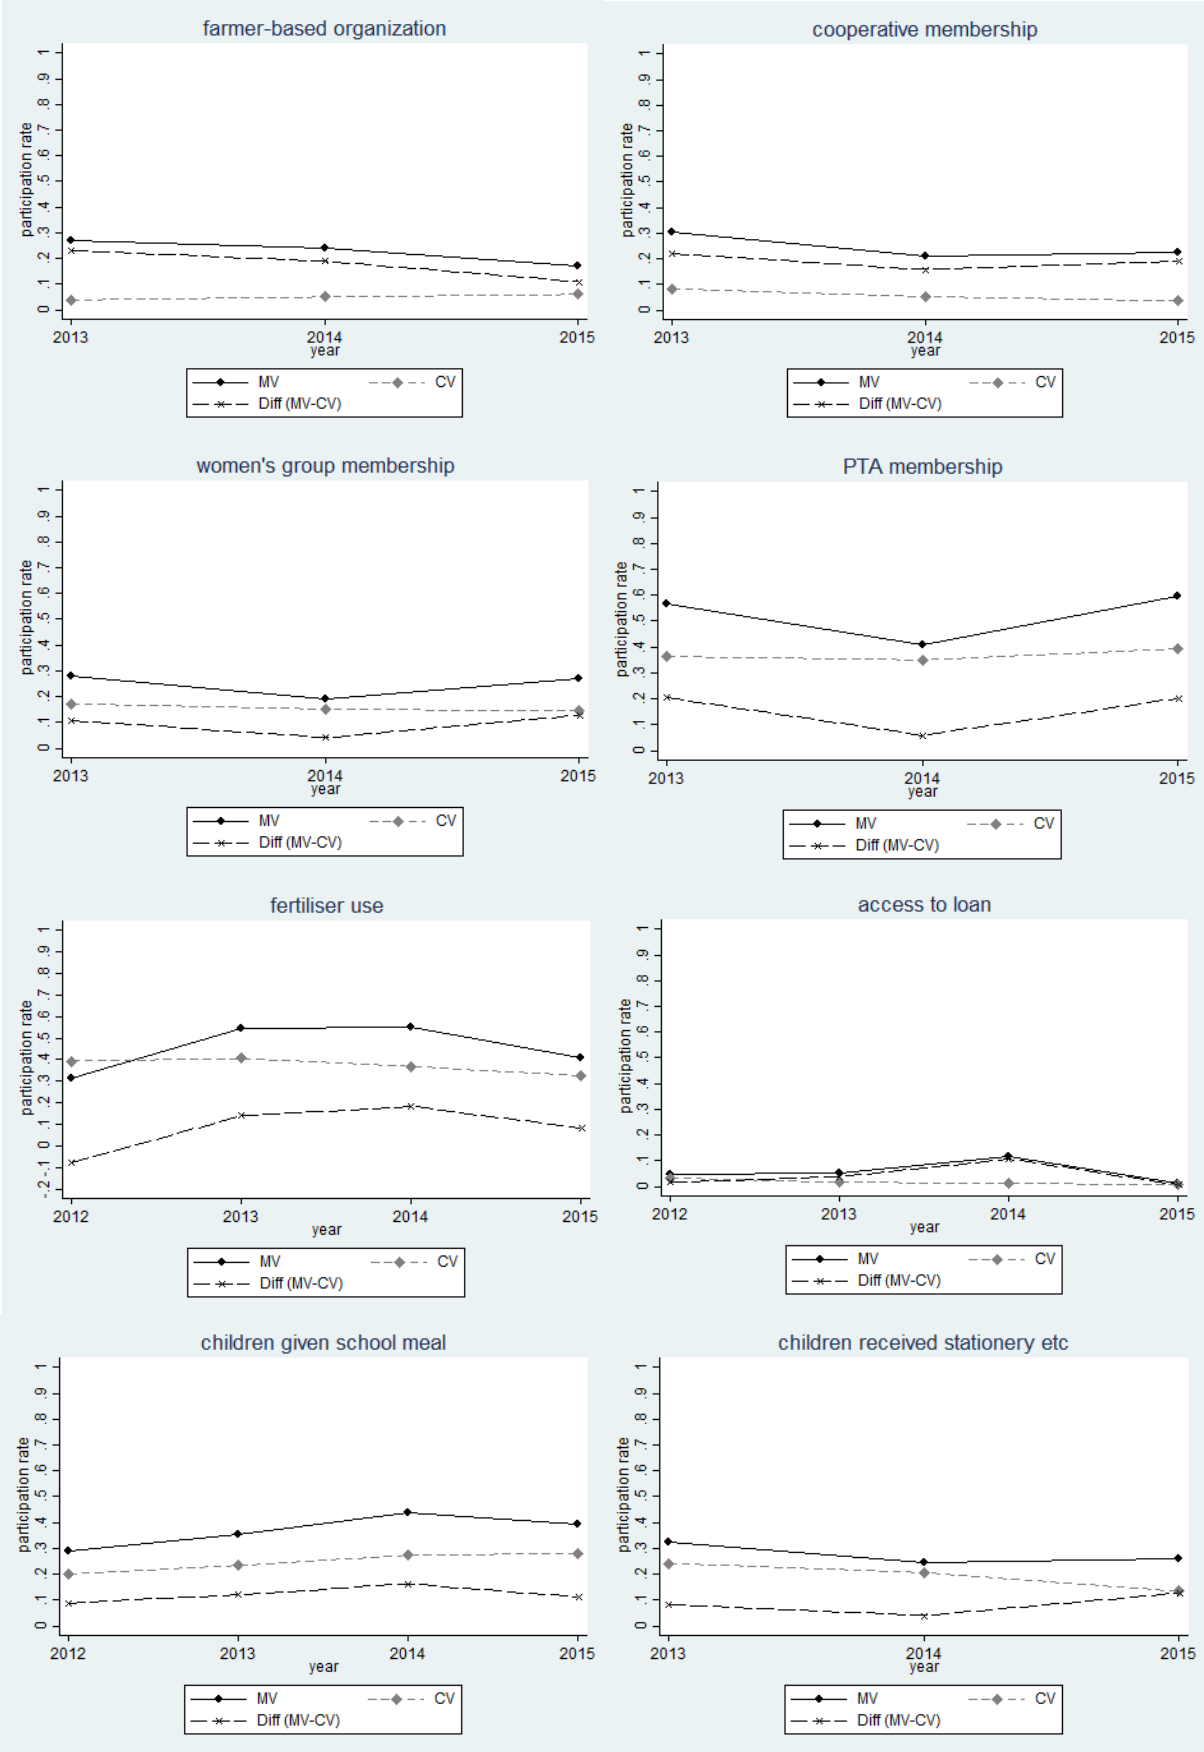

Figure 2. Participation in project activities over time: health

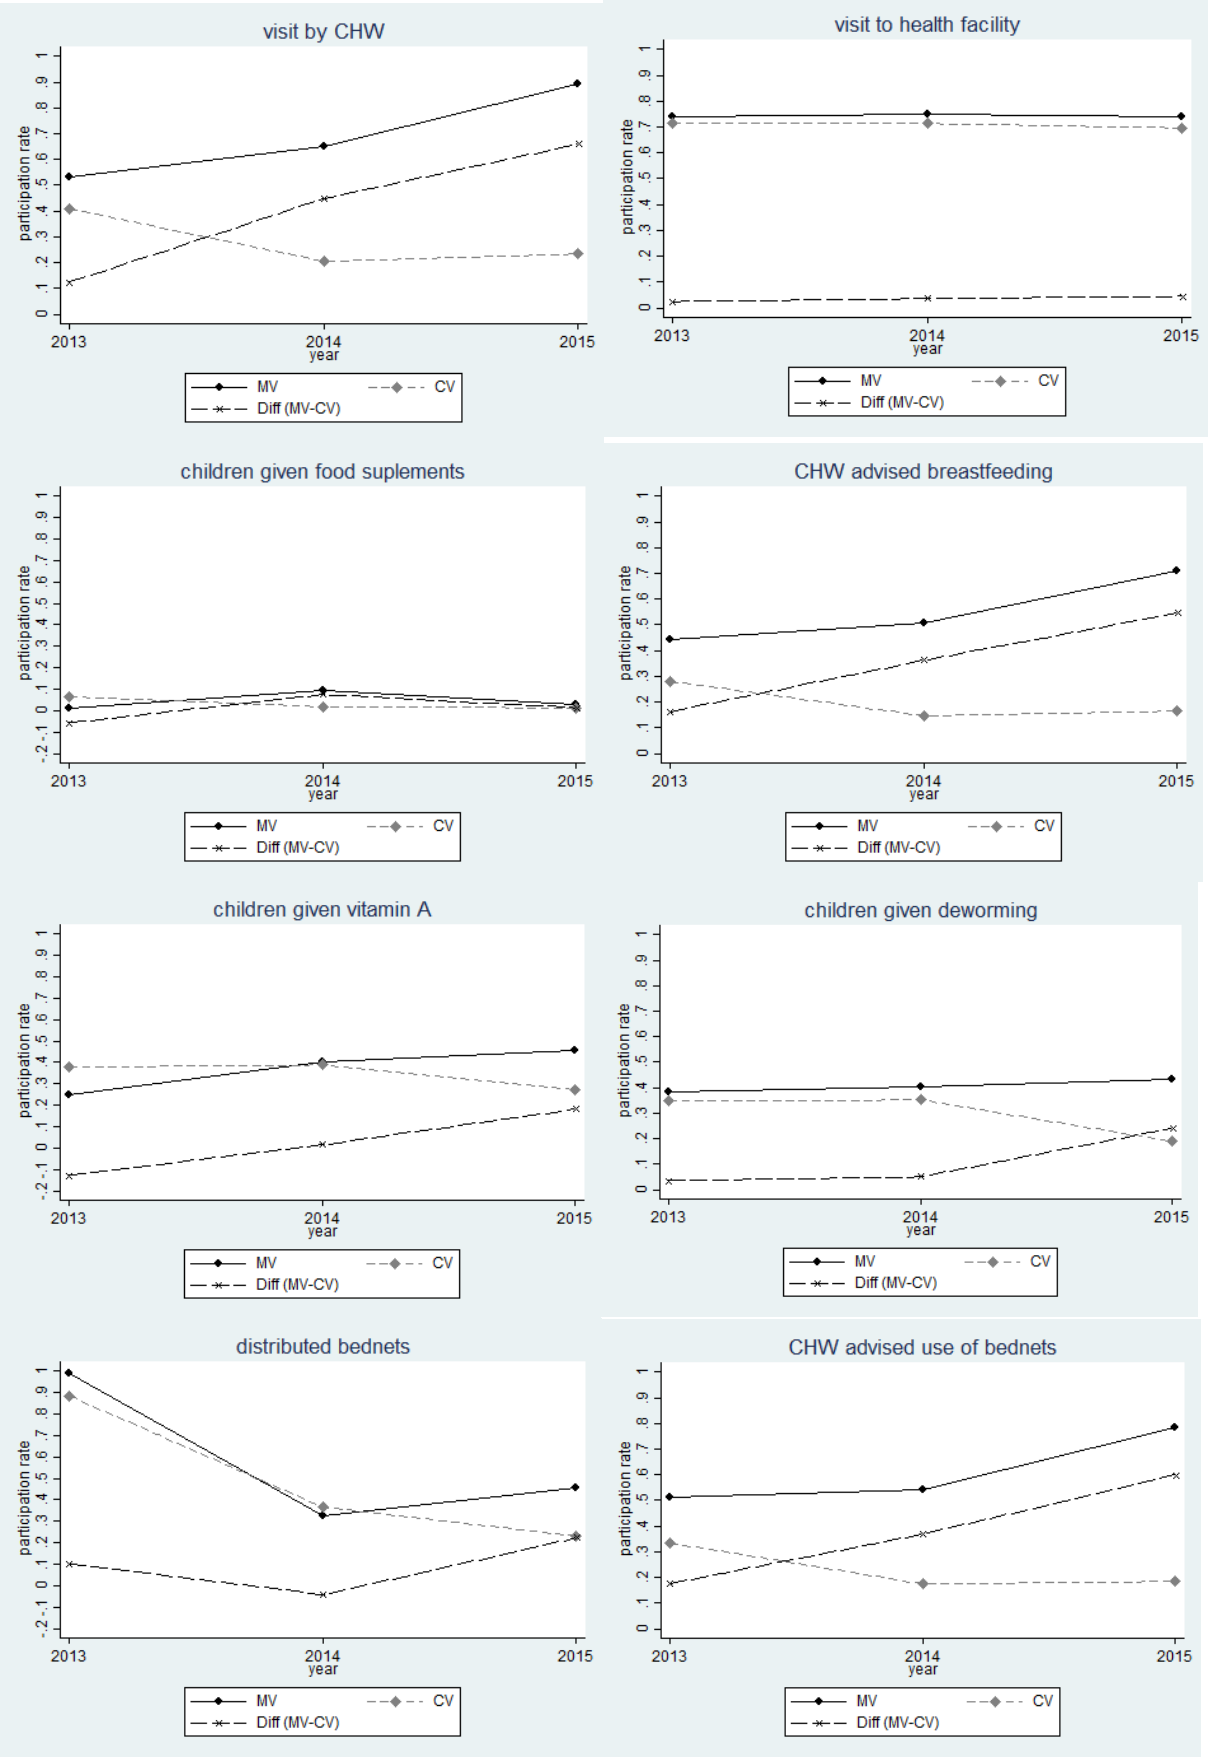

The participation patterns in activities promoted by the project are very similar in some cases in the MV and CV areas (see for example the charts on cooperative membership and school meals in Figure 1 and the charts on visits to health facilities and use of bed nets in Figure 2). We attributed this to the fact that many of the activities promoted by MVP are also promoted by other non-governmental organisations (NGOs) or by the government in neighbouring areas. An alternative explanation is that individuals living in villages near the project areas are able to share some of the benefits of the intervention, for example, by attending a clinic restructured by the project or by registering as farmers in MV cooperatives to obtain loans and fertiliser. To explore this hypothesis, we disaggregate the comparison group into two sub-groups – "near" and "far" villages – as of baseline design, under the assumption that individuals living in far-away villages are very unlikely to be able to access project benefits (Figure 3).

Overall, there is limited difference in participation in the "near" and "far" CV areas, with the exception of fertiliser use, school meals, stationery provision and deworming (larger in "far" communities), which may suggest a higher presence of government and NGO interventions in these areas. Participation in social groups is very similar in near and far comparison areas, with the exception of PTA participation in 2014, suggesting that individuals living in neighbouring villages are not joining the groups promoted by the intervention. Access to loans is nearly identical in near and far communities, while fertiliser use is much higher in far-away communities. This suggests that farmers of nearby communities are unable to access loans and fertiliser offered by the intervention.

Similarly, many children are receiving school meals and stationery in far-away villages, which makes spill-over effects of this type unlikely. The charts illustrating participation in health activities offer an interesting picture. Households near to MV areas are more likely to visit health facilities, at least in 2014, and to be visited by a CHW, particularly for the distribution of bed nets, and for advice on breastfeeding and the use of bed nets. Of course, these charts are also consistent with more NGO activities in nearby areas than in far-away areas independent of the MV project. They do, however, hint to some possible spill-over effects in the provision of health services. In the case of education activities (school meals, PTA membership and stationery provision), participation in MV areas is more similar to participation in far-away communities than to participation in near CV communities. This is likely to be related to different patterns of operation of NGO and government programmes in the districts. Unfortunately we have not been able to map the operation of projects and NGOs to this level of detail in the field.

Figure 3. Participation in project activities over time: MV and near and far control groups

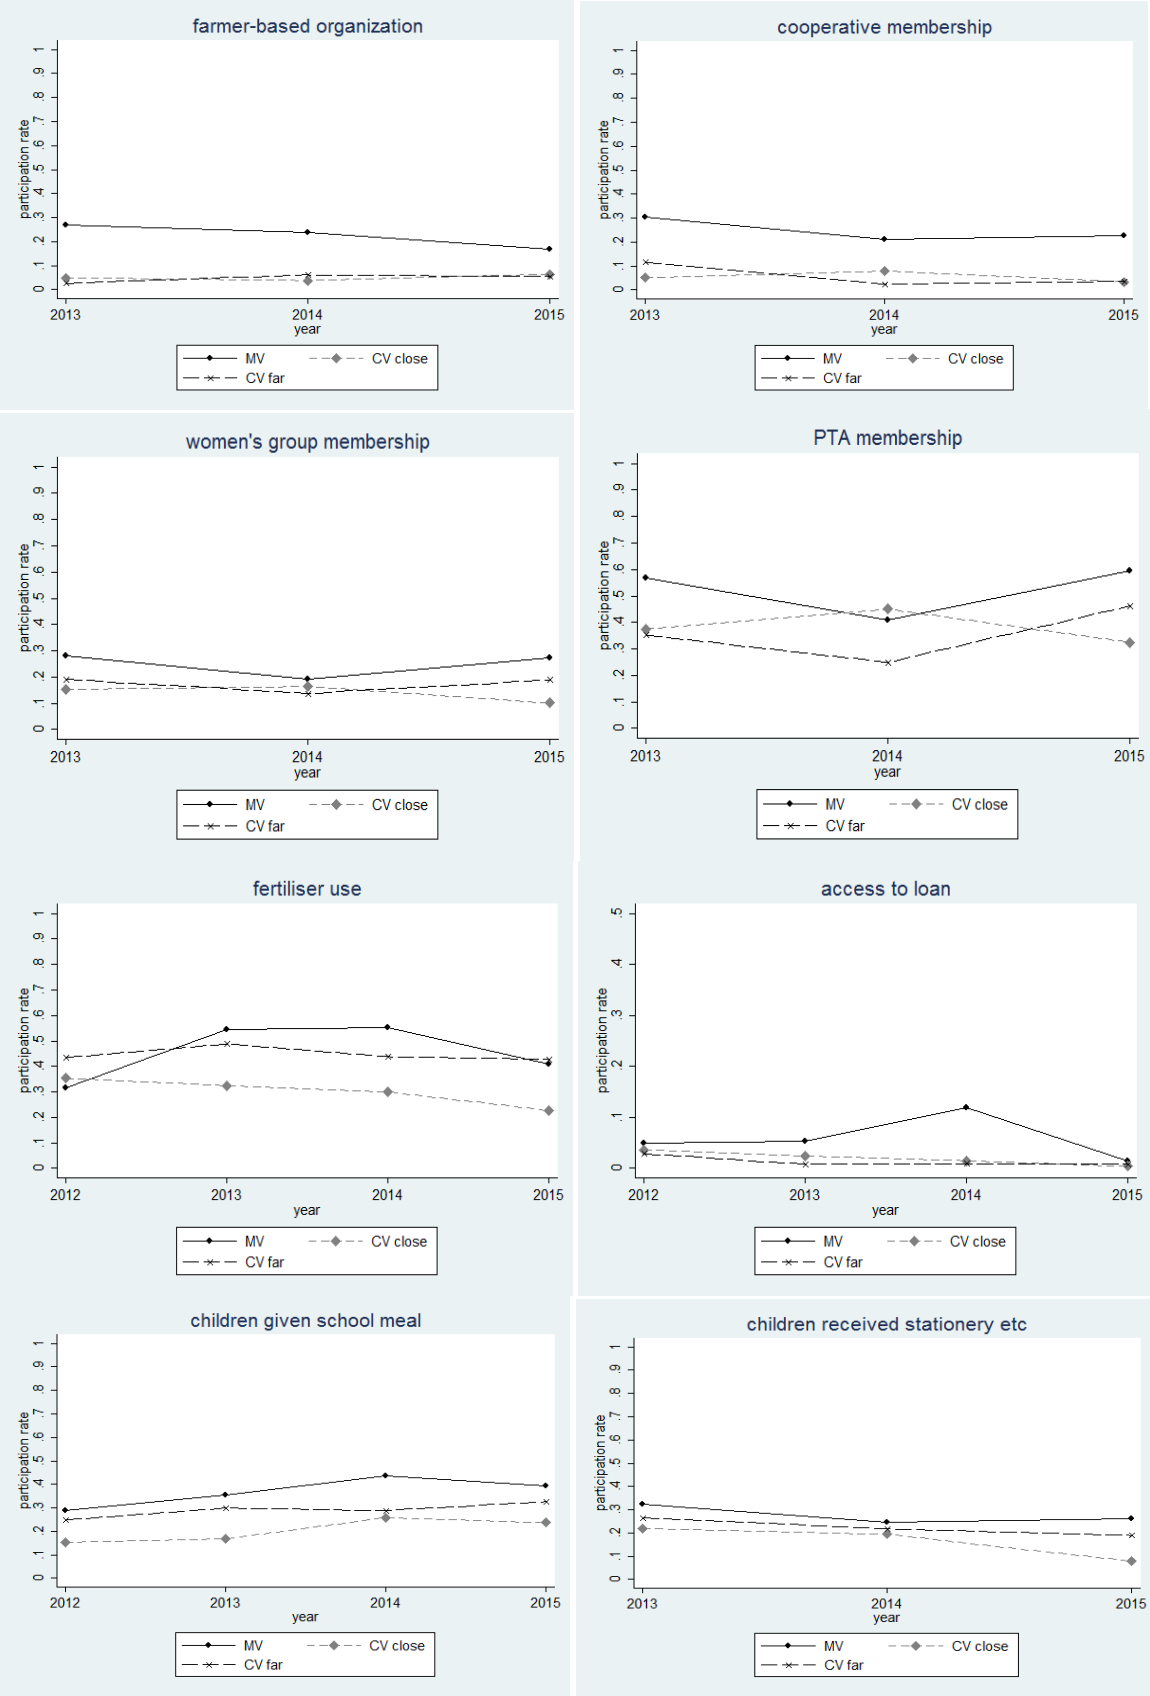

Figure 4. Participation in project activities over time: health in MV and near and far control groups

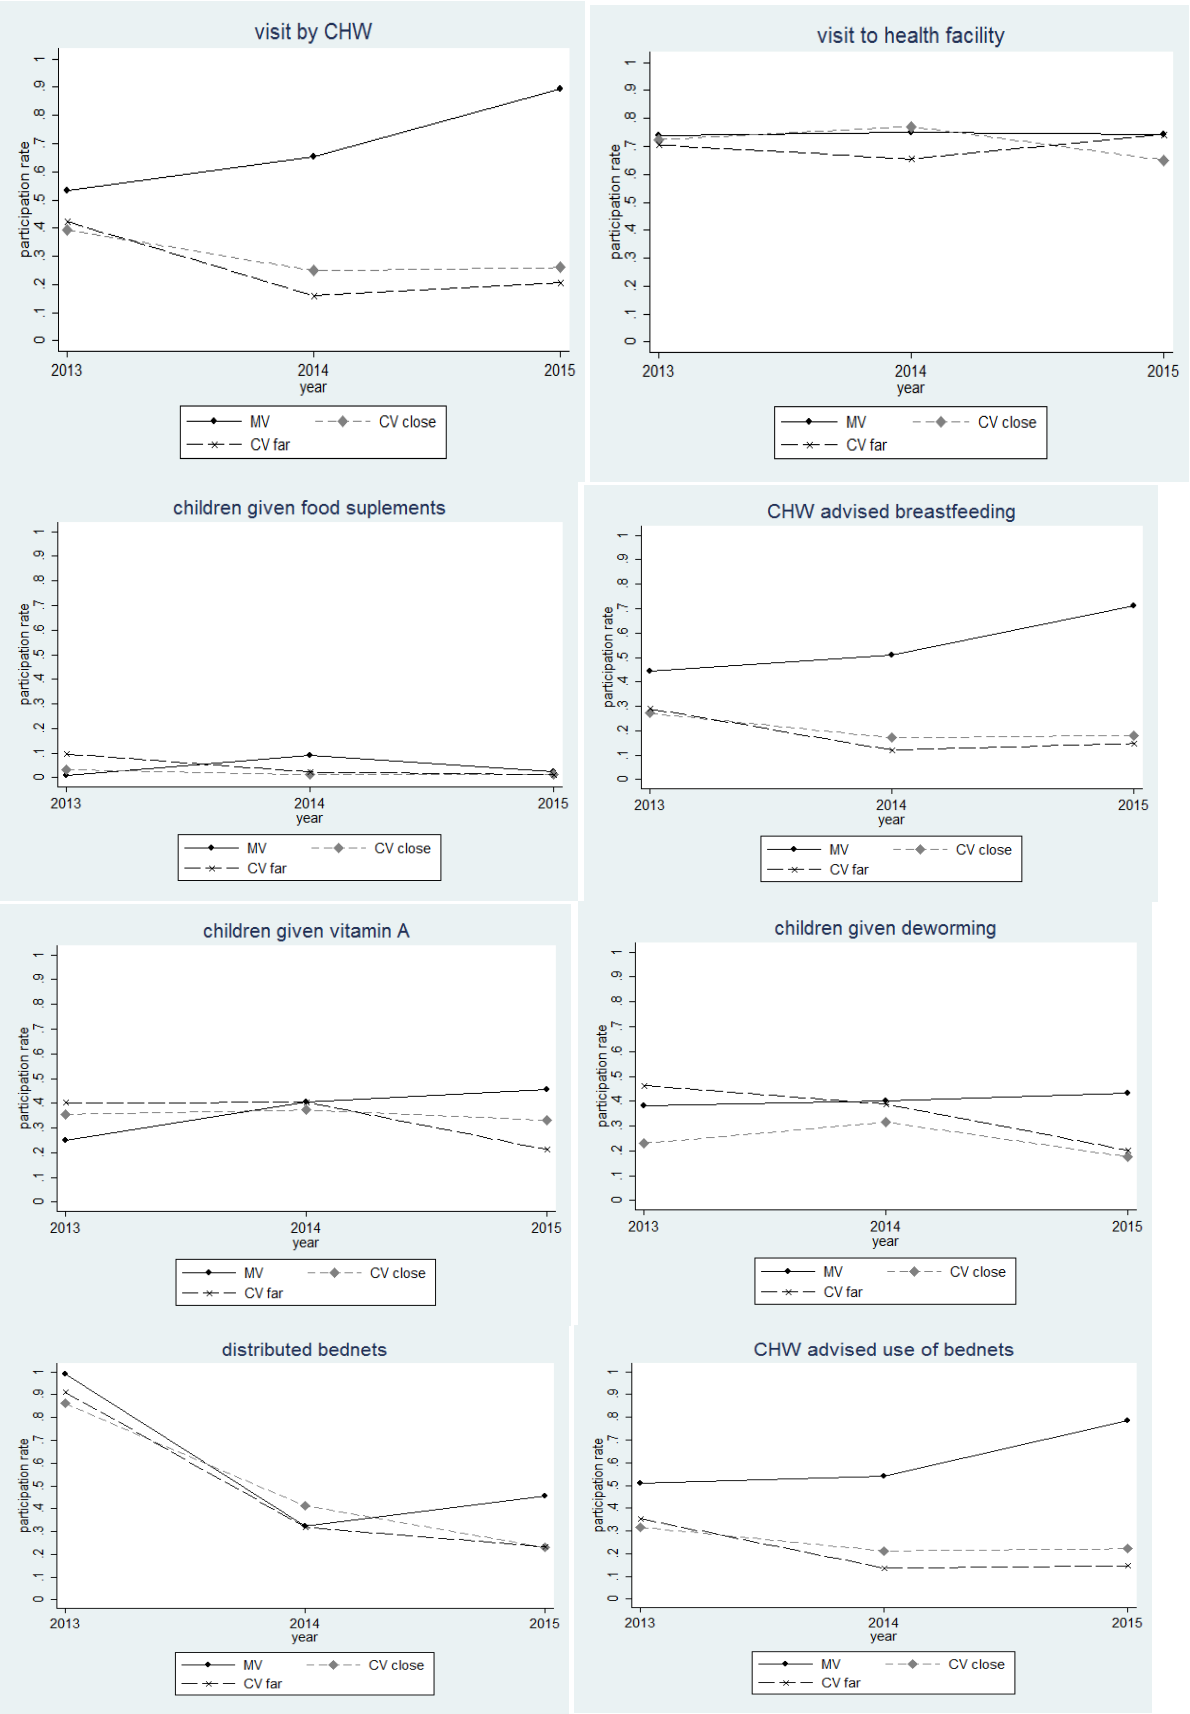

As in previous reports, we also look at participation levels across the expenditure distribution in order to detect any pattern in either participation or project targeting. The charts in Figure 5 plot involvement rates against per capita expenditure and Table 11 reports a number of statistical tests on the observed patterns.

Similar to the results from previous survey rounds, we do not find any evidence of selective targeting. The project attends equally to all population sectors regardless of welfare levels (Test 1 in Table 11). Few patterns are found in the CVs as well. In the CVs fewer people attend women's groups and PTAs as their income increases. Though MVs and CVs differ in levels of participation, they do not differ in terms of participation patterns across the expenditure distribution.

Figure 5. Participation in project activities by household expenditure

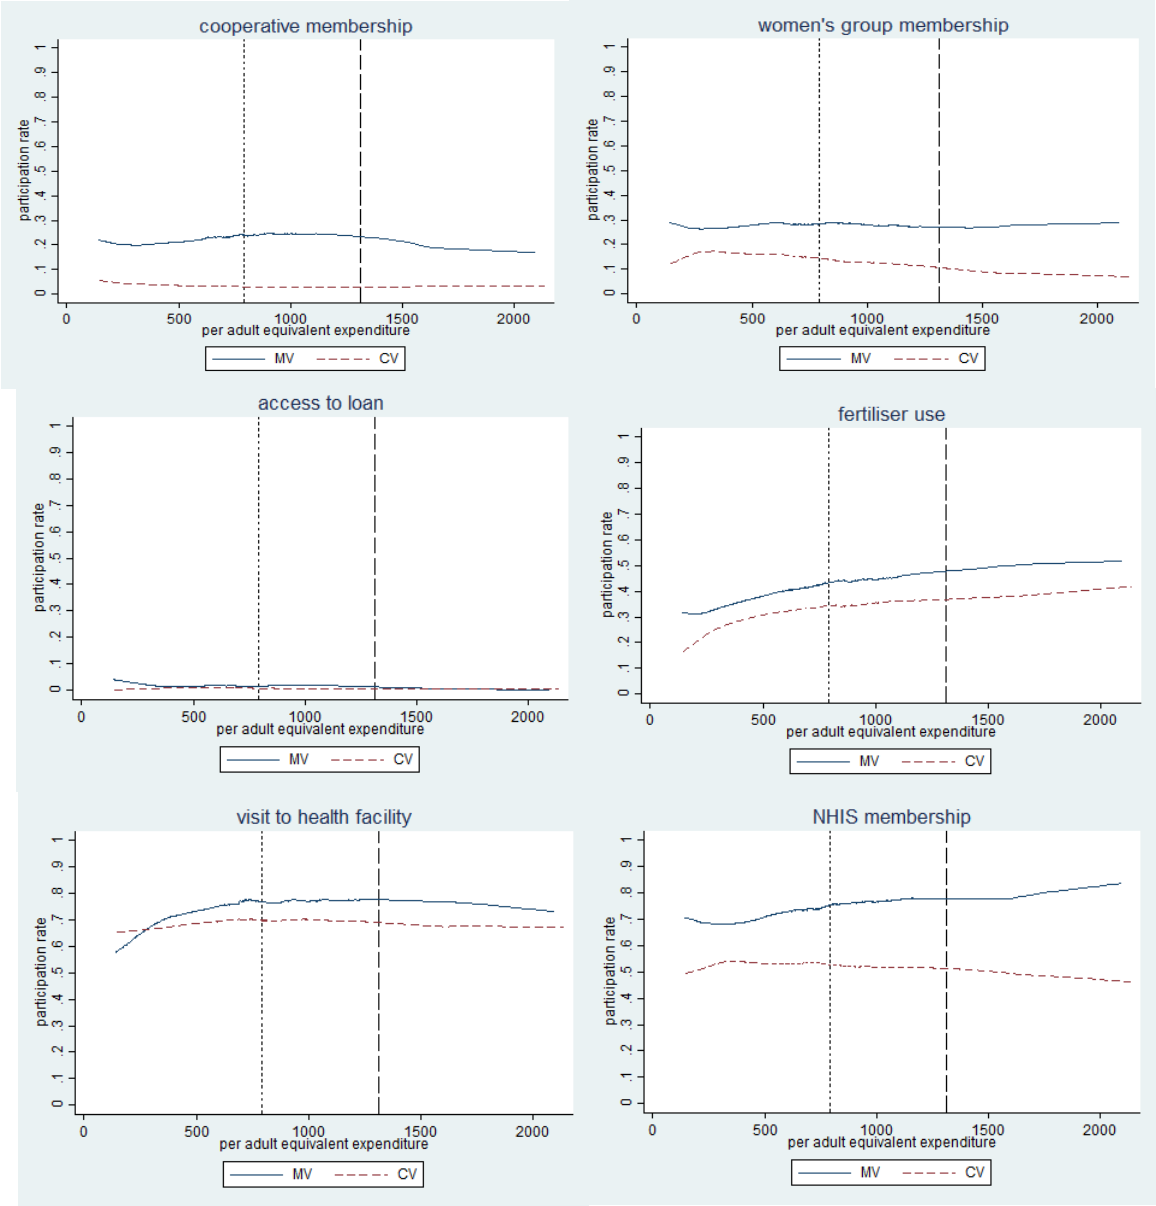

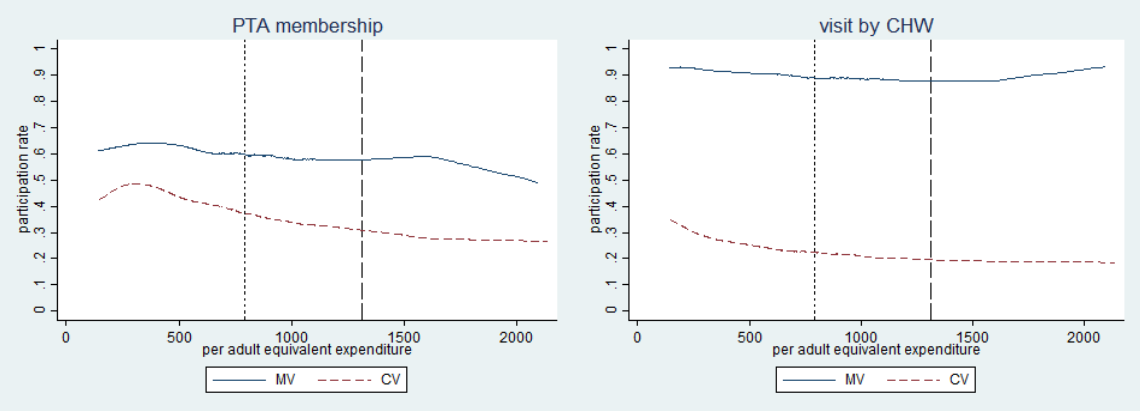

Table 11. F-test on equality of participation patterns

Equity in participation	Test 1	Test 2	Test 3	Test 4
Membership of cooperatives	1.13	1.08	9.32***	1.48
Membership of women's groups	0.43	3.77**	2.94**	1.13
Membership of PTA	1.16	4.54***	4.81***	2.51**
Loans	1.21	0.41	0.79	0.37
Fertiliser use	0.62	1.43	1.23	0.72
Membership of NHIS	2.37*	1.09	5.01***	1.84
Visits by community health workers	1.20	0.93	56.60***	1.22
Visits to clinics	0.40	1.01	1.70	0.33

Tests are conducted by running regressions of participation dummies on quintiles of baseline per adult equivalent expenditure of MV and CV households. Test 1 evaluates the presence of a pattern in the MVs (equality of all MV slope coefficients). Test 2 evaluates the presence of a pattern in CV areas (equality of all CV slope coefficients). Test 3 evaluates the difference between MV and CV areas at each quantile (all MV slope coefficients are zero). Test 4 evaluates the difference in patterns between MV and CV areas (joint difference between MV and CV slope coefficient).

3. Impact of the intervention on poverty, income, education and migration

This section covers the impact of the intervention on four welfare indicators: poverty, income, school attendance and migration. Not all MDG indicators are assessed because the second and fourth rounds of data collection only cover a limited range of indicators.

3.1 Methodology

Project impact is estimated by combining DD analysis and propensity score matching methods. However, unlike the previous reports that relied heavily on the inverse probability method to estimate project effects, in the following analysis a sub-classification method of Imbens and Rubin (2015) is employed that was only partially used in the Midterm Report. One limitation of inverse probability weights is that some observations normally have propensity scores close to one and zero, which results in extremely large weights. Large weights in turn can bias the regression coefficients used to estimate project impact (Stuart, 2010). The sub-classification method produces a less biased estimation of the project impact. The Analysis Plan identified inverse probability weights as the method of choice for matching observations in the MVs and CVs. However, after comments from reviewers on the Analysis Plan it was also recognised that inverse probability weighting can produce biased results so it was necessary to explore alternatives. In 2015, Imbens and Rubin (2015) published a book on the most exhaustive treatment of the analysis of treatment effects by matching methods to date. We decided to follow their recommended sub-classification approach in our analysis of the year four data.[5]

Following Imbens and Rubin (2015), before employing the sub-classification method, we estimate the propensity score and establish the size of the estimation sample. To do so we run a logit participation model and calculate the predicted probability of being in the MV area for each household. We then trim the sample to remove observations outside the region of overlap between the distribution of propensity scores of the project and control observations. After trimming, we re-estimate the propensity score on the trimmed sample.

Project impact is then estimated by the sub-classification method. To do so, the sample is split into blocks based on the values of the estimated propensity score until the statistical difference in the propensity score within each block vanishes. Project effects are calculated as the weighted average of the block level treatment effects, where the weights are calculated using the proportions of control and project observations in each block. The treatment effects within each block are estimated by regression analysis using three different models: simple cross-section; fixed effect model; and a lagged model (whereby the value of the dependent variable at baseline(s) is used as an explanatory variable). All tables reported below employ the sub-classification method and the results of the three models mentioned above. Further details of the sub-classification method and several diagnostics tests that led to the selection of the propensity score and of the trimmed sample can be found in the Midterm Report. We also conducted a sensitivity analysis to assess the plausibility of our matching strategy by testing the impact of MV on pseudo-outcomes and by testing the impact of pseudo-interventions on the main outcomes. The results are reported in Appendix 2. The results generally support our matching strategy, though some surprising results are found regarding income.

[5] In all cases we have estimated effects using both methods. Estimates by either method do not differ substantively.

3.2 Impact on poverty

Poverty rates are calculated using consumption expenditure data. Consumption expenditure includes all household expenditures on food and non-food items, including imputation of home-grown food and housing rents. Calculations follow the same procedure employed by the GSS and poverty rates are calculated using the official national poverty line.[6] The procedure to calculate household expenditure was described in detail in Appendix J of the baseline evaluation report and follows the practice of the GSS and best practices in the calculation of household consumption developed by the World Bank. We remind the reader that expenditure includes food purchases as well as home-grown food items; frequent expenditures of different types such as personal items, transport, health and education; and a valuation of the annual value of expenditure on housing and consumer durables. Household expenditure does not include all monetary transactions, and monetary expenditures that do not contribute to immediate household well-being are also not included. Hence, large purchases of assets and durables such as, for example, animals and cars, which provide a service over a long period of time, are not included. Similarly, purchases of production inputs, like fertiliser, and capital production goods, like tractors, are not included. Following GSS practice we adopt two definitions of poverty that have been commonly used in Ghanaian national poverty statistics: overall poverty and extreme poverty. The latter refers to the inability to purchase a minimum basket of food items that, given current diets, represent a minimum level of food consumption. Individuals from households whose expenditure is below the value of this minimum basket are classified as extremely poor or food-poor. Overall poverty is based on an enlarged consumption basket including non-food essentials such as expenditures on transportation, basic household goods, fuel and the like. Households whose expenditure is lower than the cost of this enlarged basket of food and non-food items are classified as poor.

Poverty decreased in both MV and CV areas (Tables 12 and 13) over the last three years and the reduction was larger for food poverty than overall poverty. The poverty gap and squared poverty gap also decreased substantially over the last three years in both areas and at similar rates. The distribution of expenditure became less unequal over the period and growth in expenditure that occurred benefited the poor more.

Table 12. Poverty indices (overall poverty)

	Poverty headcount		Poverty gap		Squared poverty gap	
	MV	CV	MV	CV	MV	CV
Baseline	0.870	0.878	0.486	0.483	0.314	0.307
2nd round	0.916	0.912	0.463	0.458	0.266	0.263
Midterm	0.883	0.874	0.398	0.437	0.215	0.252
4th round	0.808	0.802	0.357	0.325	0.184	0.161

[6] As in the Midterm Report, we use the 2013 national poverty line, established by the GSS based on the Ghana Living Standards Survey (GLSS) 6 survey data. The expenditure data are adjusted for regional and monthly inflation in reference to the poverty line fixed at January 2013 prices. Previous versions of the report employed the official poverty line established by the GSS in 1999 based on GLSS4 survey data, which was also employed by GLSS5 in 2004/05. The 1999 poverty line used to be updated by GSS for inflation every year but had become largely outdated in 2013.

Table 13. Poverty indices (food poverty)

	Poverty headcount		Poverty gap		Squared poverty gap	
	MV	CV	MV	CV	MV	CV
Baseline	0.658	0.655	0.287	0.280	0.159	0.149
2nd round	0.684	0.671	0.223	0.221	0.095	0.095
Midterm	0.550	0.631	0.166	0.212	0.070	0.094
4th round	0.493	0.414	0.132	0.110	0.052	0.042

The project does not appear to have an impact on consumption poverty. The chart on the left hand side of Figure 6 shows that poverty rates in MV and CV areas have been nearly identical over the three years of the intervention. However, poverty rates are very high (nearly 90% of the population in the area is poor by official standards) and differences in expenditure levels between MV and CV areas may occur below the poverty line. The right hand chart of Figure 6 shows the food poverty rate, which calculates the fraction of the population unable to afford a basic basket of food items. Some differences emerge when we look at food poverty. A smaller fraction of households were food-poor at midterm in the MVs compared to CVs. However, the exact opposite was true in the fourth round where fewer households were food-poor in the CVs compared to the MVs. Hence, there are differences in the levels of expenditure between MV and CV areas below the overall poverty line, but the direction of change is sometimes in favour of the MVP and other times against it.

Figure 6. Poverty rates in MV and CV areas over time

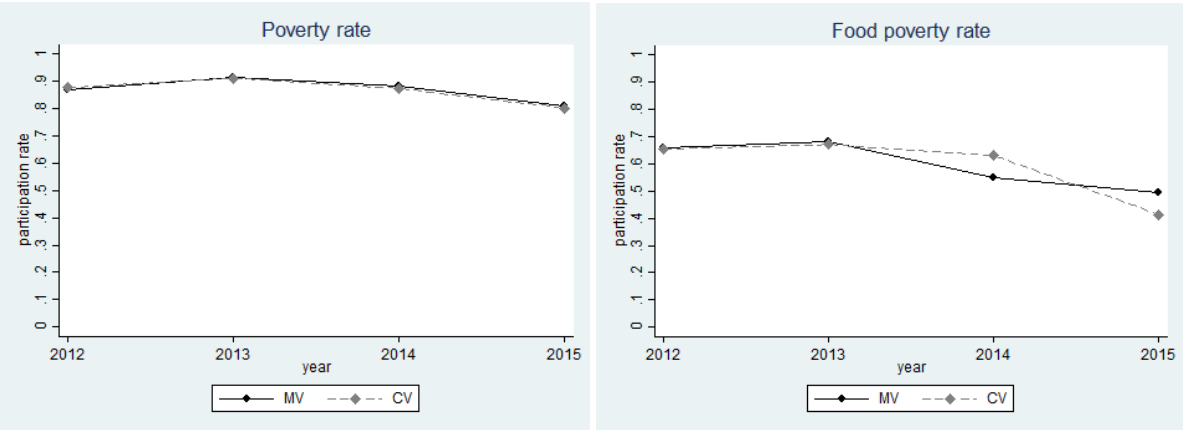

We then look at the MVP's distributional effect further by inspecting the distribution of per adult equivalent expenditure among MV and CV households (Figure 7). The charts in Figure 7 show no differences in the distribution at the baseline and at the first follow-up. There is a small shift to the right of the MV distribution at midterm and a small shift to the left during the most recent survey round (2015). While the effect at the midterm could be interpreted as a modest project impact, the effect at the fourth round is more difficult to interpret. Excluding the possibility that the MVP had a negative impact on expenditure at the fourth round, the effect could be the result of a negative shock in the MV areas. However, the poverty rate also decreased in MV areas in such a way that a negative shock can be excluded. A covariate positive shock in the CVs is also an unlikely explanation because the villages comprising the control group are spatially located around the MV cluster and are more heterogeneous than the MVs in terms of weather shocks because they cover a larger and more diverse geographic area. One possibility is that a particular cluster of villages in the control group, for example CVs in the Builsa district, was affected by an extraordinary positive shock driving the overall positive shift in the expenditure distribution. We investigated this possibility with regression analysis by

looking at reported weather-related shocks by district in MV and CV areas, but could not find evidence to support this hypothesis.

Figure 7. Densities of per adult equivalent expenditure

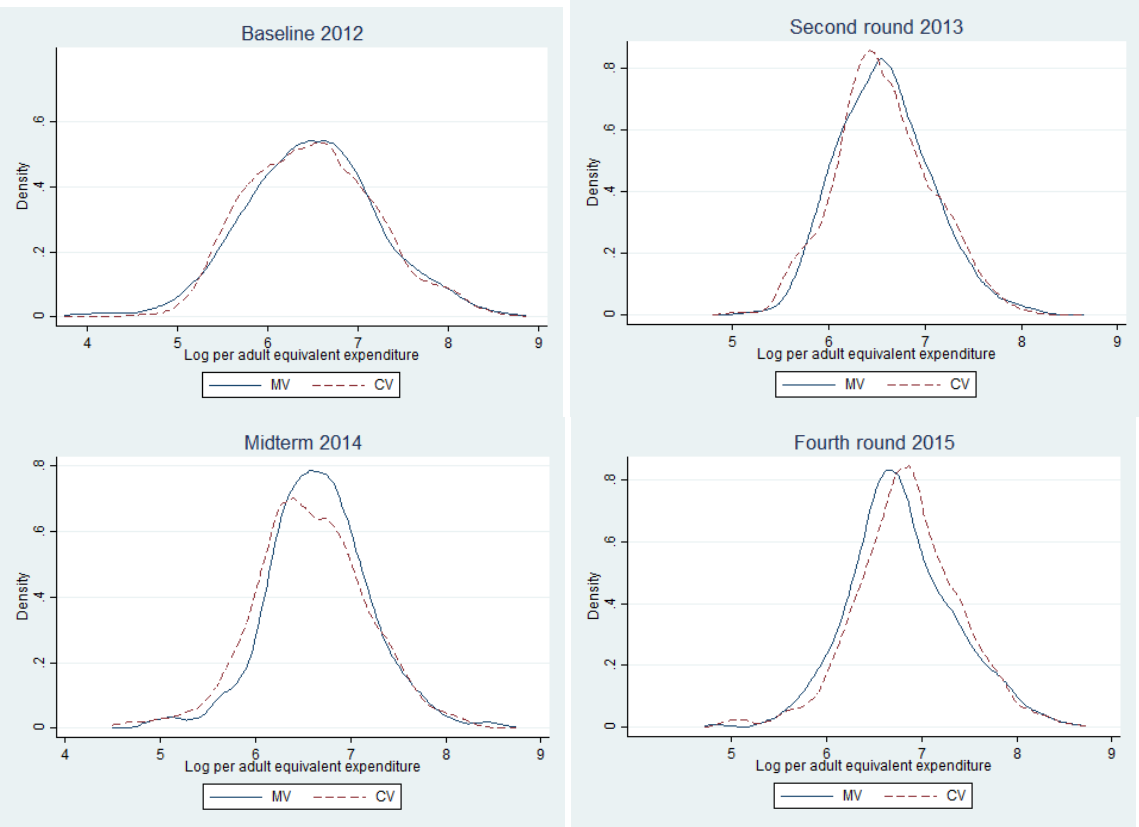

The project appears to have had a positive impact on consumption expenditure at the midterm but not at the second and fourth year of the intervention (Table 14). In fact, the observed impact at the fourth round is negative. This latter effect cancels out any positive effect observed at the midterm and the overall average effect for the three periods (first row of Table 14) is nearly zero.

Table 14. Impact on per adult equivalent expenditure

	Cross-section	Fixed effects	Lagged model
Average DD effect	-0.002	-0.006	-0.011
	(0.973)	(0.921)	(0.742)
DD effect second year	-0.003	-0.004	-0.009
	(0.970)	(0.958)	(0.842)
DD effect third year	0.076	0.072	0.066
	(0.280)	(0.305)	(0.237)
DD effect fourth year	-0.081	-0.091	-0.090**
	(0.192)	(0.151)	(0.029)
Sample size	7,847	7,844	5,857

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

3.3 Impact on income

We calculated income as revenues minus costs, including imputations for the value of non-market goods. For many households much, if not most, of food production is for their own consumption rather than to sell on the market. Nearly all households in the area farm some land, but also hold livestock, run micro-enterprises and occasionally engage in wage labour as well as receive remittances or transfers from outside the area.

The impact of the intervention on income is undoubtedly positive (Table 15). Impact is measured in terms of standard deviations. It would be ideal to use logarithmic transformations so that the coefficient for the MVP could be interpreted as a percentage effect. However, many incomes are negative because of negative income shocks affecting crop and livestock production, and logarithmic transformations cannot be used. Instead standardised income is used, whereby income is divided by the standard deviation of income at the baseline. The standard deviation of income is nearly twice the size of the average so that the effects reported in Table 15 roughly represent half the percentage change. For example, a standardised effect of 0.3 standard deviations is roughly equivalent to a 60% increase in income. The effects reported in Table 15 are very large by any standard of comparison. The effects are also consistent as they are similar in size every year and across different estimation models.

Table 15. Impact on per capita income

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.390***	0.398***	0.305***
	(0.000)	(0.000)	(0.000)
DD effect second year	0.359***	0.371***	0.269**
	(0.000)	(0.000)	(0.007)
DD effect third year	0.408***	0.425***	0.326**
	(0.000)	(0.000)	(0.003)
DD effect fourth year	0.407**	0.404**	0.321**
	(0.003)	(0.005)	(0.009)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

These positive income effects cannot be interpreted as the result of positive income shocks in the MV cluster unless we are willing to believe that positive shocks occurred every single year in the MVs over the last three years. In addition, the data does not support this hypothesis. Households in MV and CV areas report similar occurrences of drought, floods and crop losses (Figure 8) over the four rounds of data collection, including the two years before the baseline.

Figure 8. Households affected by weather-related shocks in MV and CV areas

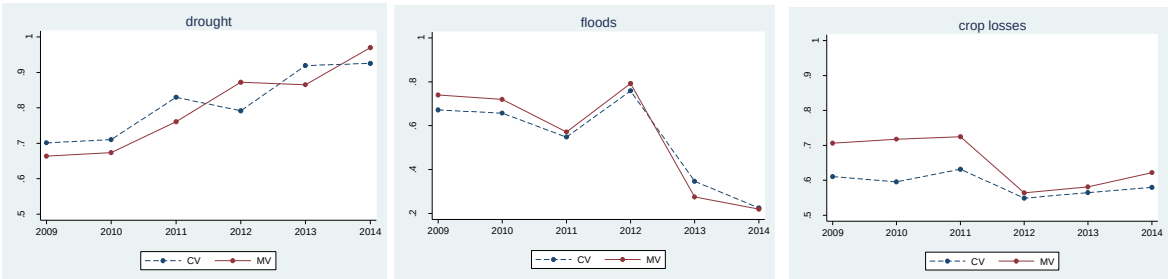

Changes in the distribution of per capita income in MV and CV areas (charts in Figure 9) suggest that the project did not significantly affect the percentages of very poor and relatively rich households, but shifted the middle distribution to the right.

Figure 9. Changes in per capita income distributions

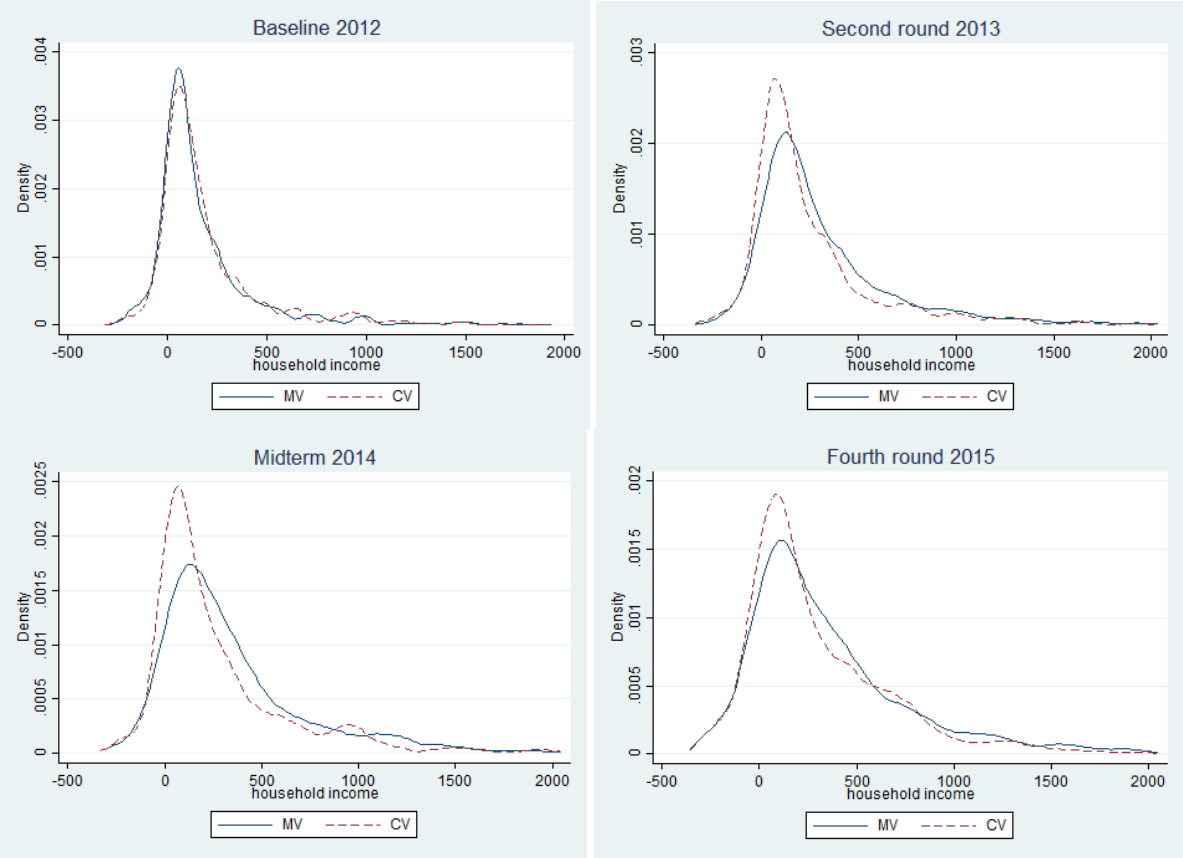

Households in the sample obtain most of their income from agriculture (Table 16). Farming accounts (approximating to the nearest integer) for 42 and 63% of total household income depending on the survey round, while livestock activities constitute between 22 and 33% of total income. Micro-enterprises account for some 7 to 17% of income and the remaining 5 to 7% is obtained from wage employment and remittances. Note that incomes from farming, livestock activities and micro-enterprises are often negative due to negative shocks. Shares cannot therefore be reliably calculated, so the column totals in Table 16 do not always add up to 100% as they should. The numbers in Table 16 nevertheless give a rough idea about the main contributors to household income. Farming is clearly at the top of income-generating activities, followed by livestock activities and, to a lesser extent, micro-enterprises. For the average household transfers and wage employment are almost irrelevant.

Table 16. Income percentage shares by source (MV and CV areas combined)

Income source	Baseline (2012)	Second round (2013)	Midterm (2014)	Fourth round (2015)
Agricultural	63.1	44.9	41.6	51.3
Livestock	22.3	30.1	32.6	29.9
Business	7.0	17.1	16.8	11.7
Employment	4.5	4.3	5.5	5.5
Transfers	0.8	0.8	0.8	1.1

Project impact is estimated separately on different income sources listed in Table 16 to understand which contributes more to the overall impact on income (Table 17). While income percentage shares are very similar at baseline, they appear to change in different ways over time in MV and CV areas. The data show that the largest impact is on agricultural incomes. Based on the results in Table 17, the project has the effect of doubling farming incomes. The impacts on livestock activities and income from micro-enterprises are similar to each other and more limited in size. There is no impact on employment income. The project has a negative impact on wage income as if employment moved from wage to household farming activities, but the effect is not statistically significant. Interestingly, the project appears to have a positive impact on transfers and remittances. Overall, these results suggest that much of the observed impact on per capita income is driven by an increase in income from farming. Farming makes the largest share of per capita income and is the income source that increased the most after project implementation.

Table 17. Difference-in-difference impact of MV on different income sources (sub-classification, fixed effects)

	Average DD effect	DD effect second year	DD effect third year	DD effect fourth year
Agricultural income	0.517***	0.398**	0.430**	0.739**
	(0.000)	(0.002)	(0.005)	(0.002)
Livestock income	0.277***	0.211**	0.416***	0.204
	(0.000)	(0.026)	(0.000)	(0.147)
Business income	0.267**	0.423**	0.288**	0.088
	(0.002)	(0.010)	(0.028)	(0.425)
Employment income	-0.036	-0.059	-0.073	0.026
	(0.539)	(0.255)	(0.332)	(0.779)
Transfers income	0.217**	0.134*	0.201**	0.317**
	(0.004)	(0.074)	(0.041)	(0.013)

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

The large observed increase in per capita income is in stark contrast to the stability of consumption and absence of impact on consumption and poverty observed earlier. Reconciling data on consumption and income changes was discussed in the Midterm Report and is further analysed here to advance hypotheses that might explain the simultaneous occurrence of an increase in income and stability of consumption (Section 4.1).

3.4 Impact on education

The second output of the MVP aims to achieve "enhanced access to quality primary education."[7] By improving access to primary education, the project expects to reach intermediate outcomes that "ensure universal primary education, increase the quality of education overall and increase access to secondary education, especially for girls".[8] By reference to quality education, it is assumed that the project anticipates having a positive impact on learning outcomes for children in northern Ghana.

The MVP aims to achieve these goals through five output areas:[9]

1. Improving education quality
2. Increasing primary school enrolment

[7] 2014, Annual Report on the Millennium Villages Project in Northern Ghana, p. 17.
[8] 2016, Logframe for the Millennium Villages Accountable Grant Programme, DFID.
[9] 2012, Annual Report on the Millennium Villages Project in Northern Ghana, p. 10.

3. Increasing participation in secondary education
4. Improving gender parity
5. Engaging communities in education

The MVP anticipates that several changes will result from the activities to improve learning outcomes in the MVs. By implementing multiple interventions generating impacts at various levels, it is expected that there will be a range of outputs, such as: greater demand for education by children and parents, schools managed more effectively, additional and better quality resources to facilitate learning are available, and parents have greater incentives to send children to school – particularly girls. As a result of these outputs, improving children's access to formal education will be achieved and motivate parents to educate their daughters. At the same time, the education that children access will be better quality both in terms of teaching and the learning environment. These results are expected to increase both enrolment and classroom attendance, which will have a positive impact on improving learning outcomes in the MVs, and therefore the achievement of the MDGs.

The MVP aims to achieve the anticipated results through a range of activities delivered across the education sector. In the first year of operation, project staff conducted several needs assessments with the communities, parent-teacher associations (PTAs), school management committees (SMCs), and district education directorates. The meetings revealed the scale and variety of problems faced by the education system in the north: inadequate buildings and teaching materials, teacher absenteeism, poor teacher qualifications, high teacher turnover, language barriers to learning, economic and social constraints to school attendance such as long distances to school, absence of toilets for girls, and the value parents place on schooling.

The project devised an overall strategy to tackle these problems with the main goal of increasing school attendance. The strategy is based on delivering **activities** within three main pillars: (a) improving school quality; (b) sensitising communities and parents; and (c) enrolling more girls in school. Additional interventions aimed at bringing more children to school and monitoring children's/students' learning were attempted, but on a much smaller scale. Figure 10 summarises the MVP's education causal chain.

Figure 10. Theory of change of education interventions

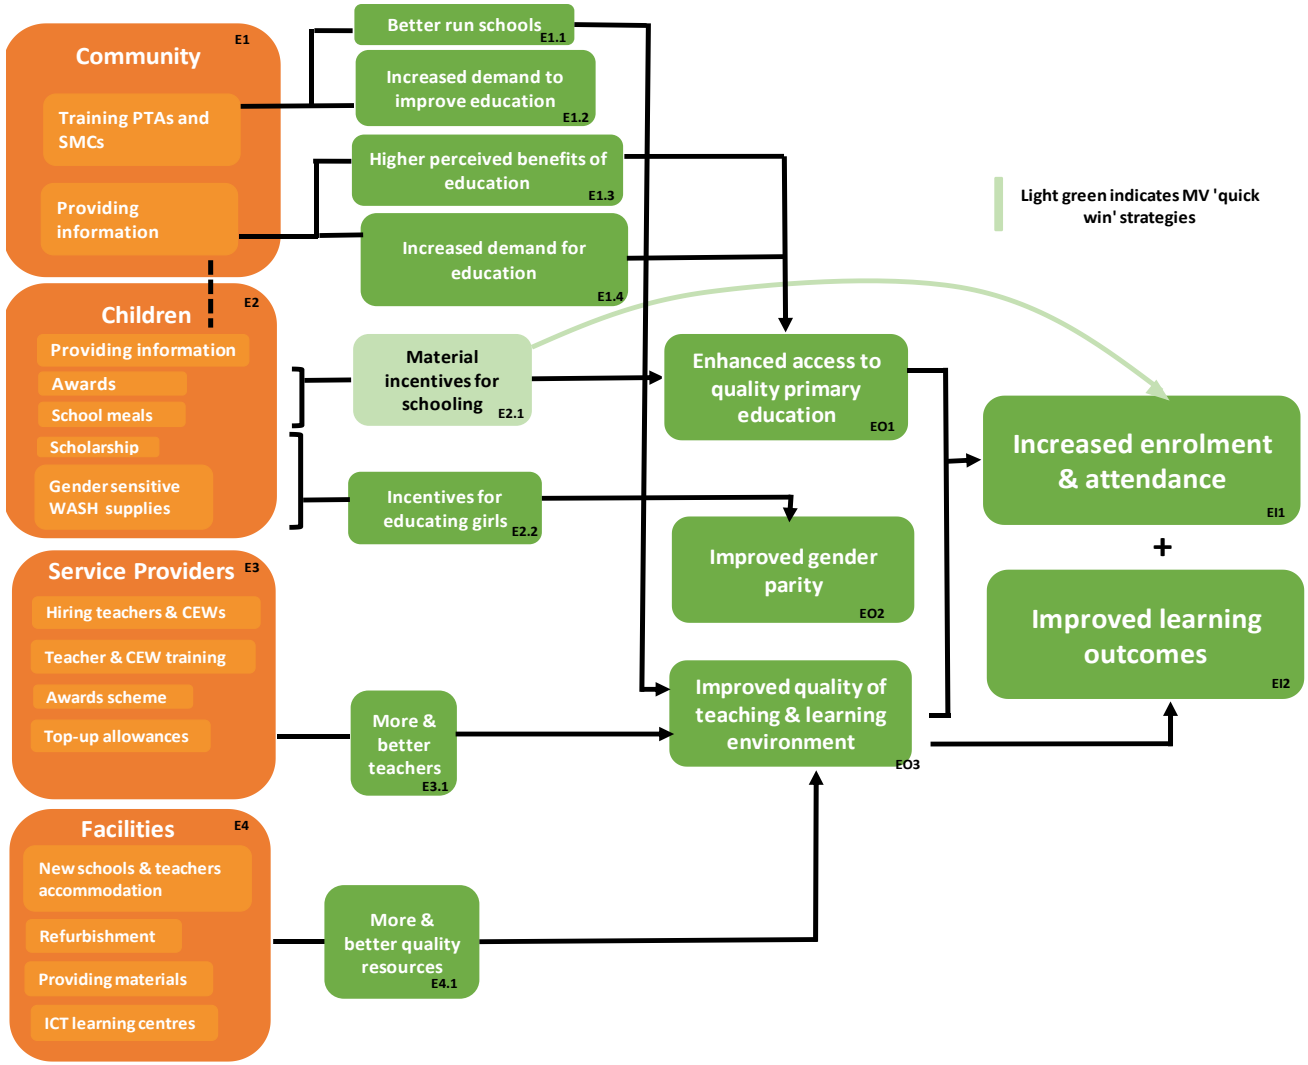

- **Schooling quality:** It was thought that one of the main factors behind low school attendance was the poor quality of instruction. This in turn was the result of poor school infrastructure and poor teaching (including the intimidation of children). Hence, the project invests heavily in the construction and rehabilitation of classrooms, school toilets and playgrounds, and refurbishes schools with sporting equipment, teaching materials, books and computers. In order to increase the *quality of teaching*, the project builds teacher quarters and provides other incentives for teachers to live in the communities. The project trains teachers on teaching methods and provides salary top-ups to staff of the Ghana Education Service (GES) to supervise teachers' work.
- **Community**: The aim of the community sensitisation work is to strengthen communities' understanding of the role they can play in advancing children's education (e.g. by ensuring children get to school on time, holding schools, head teachers and teachers accountable for children's performance, the school holding the community members accountable for their responsibilities to children, etc.). The project hires and trains community education workers (CEWs) with the goal that they will hold meetings and workshops with the communities, PTAs and SMCs to sensitise parents about the benefits of education. In addition, CEWs visit families of children not attending school and families of children who dropped out of school to get more children in school.
- **Gender parity**: In order to boost girls' school attendance, the MVP implemented a varied set of initiatives, including: school toilets for girls, delivery of sanitary pads to prevent absence from school during menstruation, community and parents sensitisation on the benefits of girls' schooling and scholarships for girls attending senior secondary schools.

In addition to these broad packages of initiatives, the project also tried to increase school attendance directly by supporting the provision of school meals[10] and establishing a real-time monitoring system in schools to improve learning. Real-time monitoring was performed by CEWs using mobile technology. CEWs would assess students' reading skills on a regular basis to inform project staff and education authorities about progress being made and establish areas where remedial education was needed.

The MVP's education interventions are delivered through various settings or "levels" based on the barriers each activity aims to address. In the case of Ghana's SADA (Savannah Accelerated Development Authority)-MVP, the intervention levels are grouped into child (children) level, school (facility) level, teacher (service provider) level and community level. At the child level, interventions focus on improving children's ability to benefit from schooling or their incentives for investing time and resources in their own education (e.g. activities such as school feeding or merit-based scholarships). School-level interventions aim to improve the quality of the learning environment in schools and classrooms (e.g. resource provision, building new schools or infrastructure, refurbishing existing infrastructure). Teacher-level interventions target teachers directly (e.g. hiring teachers directly, hiring CEWs, decreasing pupil-teacher ratios, provide teachers with new skills, and provide performance related incentives).[11] Finally, at the community level, interventions focus on creating demand for education and supporting parental engagement in school management functions through SMCs and PTAs.

A summary of the MVP's known educational activities and achievements is detailed in Table 18. It should be noted that not all the activities may fit within the allocated categorisation since each output

[10] The provision of school meals was under the existing school feeding programme. It is not a programme specific to MVP, but rather supported by the project.

[11] Stevenson et al., 2015, *Interventions for improving learning outcomes and access to education in low and middle-income countries: a systematic review*.

may contribute to multiple educational outcomes. Despite the complexities in specifying which outputs lead to various outcomes within the intervention logic, the table attempts to group activities together according to where they are most relevant.

Table 18. Summary of MVP's education activities[12]

Community	Children		Service providers	Facilities	
Community sensitisation	**Gender parity**	**Resource provision**	**Teaching quality**	**Infrastructure**	**Resource provision**
Training PTAs and SMCs	Gender-based violence awareness and Gender clubs in schools (482 school girls trained on sexual and reproductive health)	School meals	Teacher training (120 teachers trained in jolly phonic; 174 teachers trained in numeracy, literacy, etc.)	Provision of school furniture	Providing school supplies (10,000 exercise books supplied; 3,000 pens and pencils; 16,800 text books)
Radio programmes	CEW sensitise on girl education	CEW visit out of school children	Awards scheme for teachers and students	School construction	Student scholarships
Funding CEWs	Provision of sanitary pads for girls (3,200 sanitary pads in 20 MV schools)		CEWs trained as teachers (45 CEWs teach in primary schools)	Classroom construction & refurbishment (174 classrooms constructed in 20 cluster schools)	Provision of sports materials
	Running mentor girls camps		Lobbying GES for investment in teachers	Playground construction	Learning materials distributed
	Gender sensitive water, sanitation and hygiene (WASH) facilities		Solar lamps	Construction of early child development centres	
			CEWs provided with solar	Lobbying District Assemblies for school repairs	
			Curriculum printed	Building teacher accommodation (12 teachers quarters updated)	
			Top up allowances (10 GES staff)		
			Curriculum training for head teachers, circuit supervisors, gender facilitators, and district supervisors		

[12] Numbers in brackets are the indicative number of achievements over the course of the project documented in the 2016 SADA Mid-Year Report and 2015 SADA Annual Report. These numbers need to be verified with SADA-MVP.

The education interventions described above soon ran into a major difficulty. The education directorate lacked the resources to hire new teachers and indeed struggled with maintaining the existing teaching stock because of a lack of government funds. As a result, the MVP ended up employing the CEWs as teachers to supplement the absence of government teachers. This, however, reduced the CEWs' sensitisation work considerably and resources had to be invested in training CEWs rather than teachers.

Attempts to bring more girls to school were limited in scope and partially misguided. Early data collected in the area revealed that far more girls than boys attend school. While bringing more girls to school can only be beneficial, gender parity in education is surely not an issue in the area if more girls are attending primary, middle and secondary schools. In addition, the MV interventions were limited in scope. Only about 30 scholarships were given to girls attending secondary school through the externally funded "Connect to Learn" project, and sanitary pads were distributed to only a fraction of the female school population.

Other initiatives were weakly implemented. School feeding was supported by training cooks, but the intervention was, for the most part, implemented by the World Food Programme with its own resources and only in a small number of pre-selected communities. Monitoring real-time literacy achievements alerted project staff and the education directorate to the lack of progress in learning, but did not result in remedial education activities and did little more than stimulate debates on the right pedagogical approaches to follow.

In the end, the project relied heavily on the first pillar of interventions. It constructed and rehabilitated a large number of classrooms and provided teaching and learning material. However, less effort was spent on increasing the demand for schooling in other ways through teacher training, community sensitisation, promotion of gender parity and school feeding. Given that teaching was conducted by newly trained CEWs, this package of interventions would be expected to have a moderate impact on attendance and very little impact on students' learning outcomes

The assumptions underpinning the education interventions include:

- Access to schools, infrastructure and school materials are necessary to improve education quality
- There are sufficient untapped resources to ensure that teacher shortages can be addressed
- Children and young people are willing to engage in the educational programmes envisaged for example mentoring, enrichment, life skills or problem solving programmes
- Teachers are receptive to new teaching methods, programmes and meeting the needs of different vulnerable groups
- The Ministry of Education and the Ghana Education Service are willing partners, alongside NGOs and partners
- The community is open to changing attitudes towards pre-existing gender norms – for example early marriage for females, role of education for females, etc.

The project impact on school attendance was indeed small or nil. This is clear from the charts in Figure 11 which show net attendance ratios over the four survey rounds in MV and CV areas at different school levels. Attendance rates are nearly identical over time for the three school levels in MV and CV areas: primary, junior secondary school and senior secondary school. Except small year-on-year fluctuations, attendance is stable over time for all school levels and attendance rates of junior and senior secondary are extremely low (below 20% and 10% respectively). No obvious difference in trends in the MV and CV areas is visible.

Figure 11. Net attendance rates over time in MV and CV areas

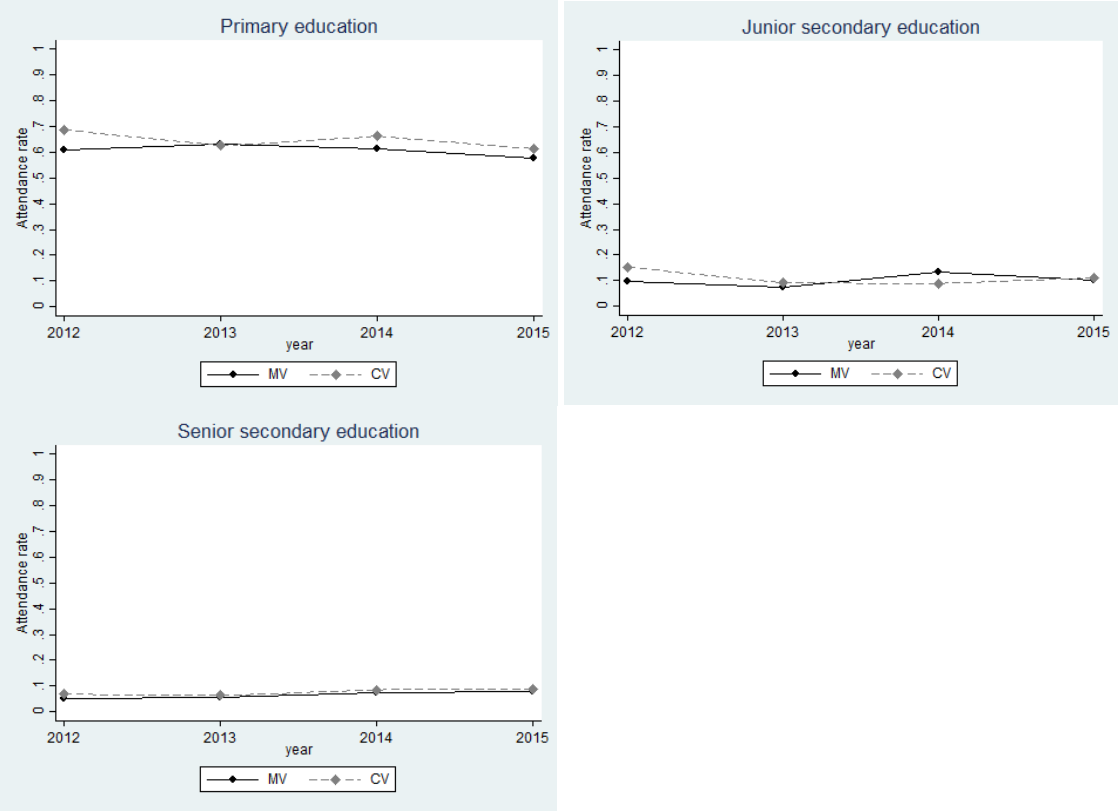

Regression analysis (Table 19) confirms the visual impressions generated by Figure 11. Some impact on primary school attendance (an increase by 10%) is found in the first year of project implementation, and an increase in junior secondary school attendance by 8% is found at the midterm. But the overall impact on basic education is modest and there is no impact on senior secondary school attendance.

Table 19. Impact on net attendance ratios

	Primary	Junior secondary	Senior secondary
Average DD effect	0.061*	0.048	0.011
	(0.057)	(0.156)	(0.568)
DD effect second year (2013)	0.103**	0.030	0.008
	(0.007)	(0.371)	(0.587)
DD effect third year (2014)	0.040	0.078*	0.012
	(0.285)	(0.069)	(0.703)
DD effect fourth year (2015)	0.043	0.034	0.006
	(0.232)	(0.380)	(0.508)
Sample size	10,903	4,369	5,475

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications

Late entrance (children starting school after the official school age), early entrance, year repetition, and recurrent spells in and out of school are common in the studied area. There is therefore a risk that age-specific attendance rates, as the ones reported in Table 19, present a distorted picture of the impact of MV on children school attendance. In particular, the project is working with older children who have dropped out of school or have never been to school. Because of their older age, these children do not appear in either the numerator or denominator of the age-specific attendance rates

used in Table 19. In order to account for the general impact of the intervention on school attendance regardless of age, we consider two further net attendance ratios. The first is the primary school attendance rate among children age 5 to 18. The second is the attendance rate of any school level (primary, junior secondary (JSS), and secondary) among children of age 5 to 18. The results are presented in Table 20. The impact of MV on primary school attendance of children age 5 to 18 is modest and not statistically significant. However, MV appears to increase school attendance of any grade among children of age 5 to 18 by 5%. The impact is similar and robust across years. The impact is not very large but, as a way of comparison, the much hailed PROGRESA conditional cash transfer programme increased primary school attendance by about 1 percentage point and secondary school attendance in a range between 3.5 and 9.3 percentage points in rural Mexico (Skoufias, 2005).

Table 20. Impact on net attendance ratios of children of age 5 to 18

	Primary attendance	Attendance of any school
Average DD effect	0.026	0.054**
	(0.237)	(0.018)
DD effect second year (2013)	0.046	0.052**
	(0.070)	(0.035)
DD effect third year (2014)	0.015	0.050*
	(0.563)	(0.051)
DD effect fourth year (2015)	0.019	0.059**
	(0.463)	(0.016)
Sample size	22,757	22,757

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications

3.5 Impact on migration

Theoretically, the project could have an impact on migration in either direction. The project could attract household members from relatives living in neighbouring areas or through marriages. In addition, migration from MV areas could decrease because households are tied to their communities by the benefits offered by the project. On the other hand, the project might be able to encourage migration outside the MV area by increasing incomes of the poor, improving their education and reducing labour demand as agricultural productivity increases.

The data give some support to the first hypothesis about migration into the MV areas (Tables 21 and 22).[13] The MVP increases in-migration and reduces out-migration. To calculate migration flows information is used on reported new household members and on members who moved outside the household over the 12 months before the household interview. In addition, we only consider household members individuals who have resided in the household for at least six months. The results are reported in Tables 21 and 22. The dependent variable is the number of migrants so that the coefficients measure the difference in the number of people migrating. The effects on in-migration are consistently positive though they are rarely statistically significant. The impact is clearly small. Similarly, the sign of the impact on out-migration is mostly negative though the coefficients are never

[13] Note that the results reported in Tables 21 and 22 are not fully comparable to the population changes previously reported in Table 6. Table 6 reported percent population changes in the sample, while Tables 21 and 22 report differences in changes in units of in-migrants and out-migrants. However, while the size of the changes is not comparable, the sign of the changes is. That is, if Table 6 reports a per cent increase in in-migration, Table 20 should also report a positive change. But note that while Table 6 reports an increase in out-migration in MV areas compared to CV areas, the opposite is reported in Table 22, where out-migration from MV areas is found to decrease. The difference is very small, as the coefficients of the models reported in Table 22 are close to zero and not statistically significant. The difference in the results of Table 6 and Table 22 is the result of using different samples of observations. Table 6 is based on the full sample of households interviewed, while Tables 21 and 22 are trimmed samples and the observations are weighted by the propensity score.

statistically significant. The project does not seem to increase migration in either direction in a substantive way.

Table 21. Impact on in-migration

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.158	0.123	0.204**
	(0.163)	(0.299)	(0.028)
DD effect second year	0.110	0.084	0.156
	(0.438)	(0.565)	(0.222)
DD effect third year	0.107	0.066	0.154
	(0.558)	(0.723)	(0.392)
DD effect fourth year	0.257	0.221	0.301*
	(0.171)	(0.255)	(0.067)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

Table 22. Impact on out-migration

	Cross-section	Fixed effects	Lagged model
Average DD effect	-0.043	-0.051	-0.020
	(0.542)	(0.462)	(0.605)
DD effect second year	-0.076	-0.083	-0.054
	(0.349)	(0.308)	(0.271)
DD effect third year	0.019	0.010	0.042
	(0.808)	(0.905)	(0.529)
DD effect fourth year	-0.073	-0.082	-0.046
	(0.364)	(0.297)	(0.335)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

Since the MV appears to have an opposite impact on migration in and out of households, we also look at the combined impact of individuals moving in and out of the household (Table 23). The combined impact simply consists of the number of individuals that migrated into the household, minus the number of household members that migrated out of the household. The impact on net migration (in-migrants minus out-migrants) is positive though it is statistically significant only in the lagged model specification. The MV appears to increase household size by 1/5 of a member on average as a result of in- and out-migration alone.

Table 23. Impact on net migration

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.198	0.172	0.218***
	(0.152)	(0.231)	(0.008)
DD effect second year	0.186	0.166	0.211**
	(0.309)	(0.381)	(0.088)
DD effect third year	0.087	0.055	0.106
	(0.638)	(0.769)	(0.527)
DD effect fourth year	0.322	0.295	0.337
	(0.119)	(0.216)	(0.051)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

Since the impact of MVP on net migration is positive and since it is likely that the programme increases survival chances. The midterm data did indeed show lower mortality rates among children under 5, though not statistically significantly different from those in CV areas. We were not able to assess the impact on mothers' mortality (which requires huge amount of data to be calculated) but the expectation is that the package of health interventions in place should produce a reduction however small. An impact of MV on household size would have interesting consequences in the estimation of some of the project outcomes. Income and expenditure are measured on a per capita and per adult equivalent basis, in such a way that an increase in household size can have the effect of reducing the observed impact. We therefore assess the impact of the intervention on the size of the household (Table 24). MV has a positive impact on household size. However, the size of the effect is less than 0.1 of a household member on average and is never statistically significant with the exception for the fourth round of data in one specification of the model. This implies that if household size is increasing in MV areas it is increasing at a similar rate in CV areas. Since household size is increasing by less than 1.5% in the MV areas, this is unlikely to bias the estimation of per capita income and expenditure, but the impact on household size will be monitored and assessed again with the final round of data.

Table 24. Impact on household size

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.083	0.083	0.091
	(0.521)	(0.521)	(0.421)
DD effect second year	-0.127	-0.120	-0.116
	(0.363)	(0.389)	(0.332)
DD effect third year	0.155	0.133	0.146
	(0.283)	(0.358)	(0.284)
DD effect fourth year	0.224	0.243	0.247*
	(0.167)	(0.143)	(0.092)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

4. Impact on savings and household wealth

The following section addresses an apparent inconsistency in the intervention's effects on expenditure and income observed in the previous section. In Section 3, we showed that the intervention did not have an impact on consumption and poverty while it had a consistent large impact on incomes, particularly incomes generated by agriculture. The inconsistency arises because it would be expected that extremely poor households immediately spend any income gain brought about by the MVP to satisfy basic needs. Indeed, poverty reduction is the ultimate goal of the MVP. In this section we discuss the theoretical foundations of the relationship between income, consumption and savings. We then illustrate how households save and accumulate wealth in the area. Finally, we estimate the impact of the MVP on savings and wealth.

4.1 Theories of consumption, income and savings

Three theories of demand and savings are relevant and applicable to the context under study: the standard Keynesian consumption function, the Friedman-Modigliani permanent income hypothesis and Deaton's precautionary savings model. We briefly describe the theories and how they fit in with the observed data and the expectations they predict about households' saving behaviour.

The simple Keynesian consumption function predicts that consumption is a proportion of income (Keynes, 1936). As income increases, a proportion of the increase is consumed, but not all, for a number of reasons, mainly as a precaution against future negative income shocks. In this model, consumption tracks income closely and savings increase over time.

According to the Friedman-Modigliani permanent income hypothesis (Friedman, 1957; Modigliani, 1966), consumption is a function of "lifetime" or "permanent" income. Households consume a proportion of their permanent, not current, income. In this model, as income increases, consumption does not unless the change in income is permanent. Thus, an increase in income can generate two opposite consumption behaviours. If the increase is perceived as a positive one-off income shock (e.g. winning a small lottery amount), the increase is entirely saved and consumption does not increase. However, if the increase in income is perceived as a permanent increase in income (e.g. salary increase following a promotion) consumption does change.

In Deaton's precautionary saving model (Deaton, 1991), households' consumption decisions are determined by: (a) a compulsion to spend any income to satisfy immediate needs ("impatience"); and (b) a need to save income for precautionary reasons to fight off any future negative income shock. Consumption is therefore the result of these two opposing forces whose outcome will depend on the psychology of the individual and the riskiness of the context. The model predicts that consumption tracks income as in the Keynesian model, but not too closely because of the precautionary motive for saving. In this model, as income increases consumption also increases, though to a lesser extent that will depend on the level of "impatience" and riskiness of the environment.

Our data show a consistent and large increase in income followed by a very small or nil change in consumption. The Keynesian model fails to explain this behaviour since it predicts a proportional increase in consumption. The permanent income theory explains the behaviour in the following way: households perceive the increase in incomes generated by the MVP as temporary rather than permanent. They expect their incomes to return to previous levels once the MVP ends and thus they save rather than spend any income gain generated by the project. The precautionary model fails to explain the observed behaviour. Extremely poor households, as those in our sample, should be very "impatient" and the model expects them to spend significant portions of additional income. The model fits the data only by assuming an extremely high uncertainty of incomes leading to much

precautionary saving behaviour and a low level of "impatience". High income uncertainty is a reasonable assumption in the MVP context, but low impatience is not. In summary, the permanent income theory appears to be the best candidate model fitting the observed household consumption behaviour under the hypothesis that households perceive income changes brought about by the intervention as temporary rather than permanent.

Other theories of consumption are proposed in the literature such as the relative income hypothesis originally formulated by Duesenberry (1949). However, these alternative models do not apply to poor economies and no other theory of consumption is relevant to the northern Ghana context than the three presented above. It should be noted that at first sight, none of these theories explains the consumption behaviour observed in a fully satisfactory way. An additional limitation of these theories is that they do not seem to contemplate the possibility that households accumulate wealth by saving income. Saving is mostly seen as a precautionary behaviour. The models seem to split societies into rich and poor where the rich save and the poor spend all of their income. The models do not seem to contemplate the possibility that the poor become rich by saving; thus, they are in a way all poverty trap models of consumption behaviour.

4.2 Savings and wealth in the area

In this section, we describe household saving and wealth using baseline household data from the MV and CV areas combined. Households in the area store their wealth in different ways (Table 28). In particular savings, livestock, assets, debts, land and housing are considered. Savings include cash holdings with banks and informal saving groups (*susu*). Households might be reluctant to share information about their financial savings or be unaware of their size since normally only one respondent is selected during household interviews and they may not know the details of each household member's saving account. Based on the available data, only 20% of households report savings with banks or *susu*. The amount of financial savings is overall negligible and represents an average 1% of current household expenditure.

Livestock is by far the largest component of household wealth. Our definition of livestock includes all animals: bulls, cows, goats, horses, donkeys, pigs, sheep, ducks, chickens and guinea fowls. More than 95% of households report holding livestock and the average value of animal holdings amounts to 70% of current annual expenditure. On average (mean), households could fund 70% of their current annual consumption by selling all their animals. No other wealth component stores as much value as livestock and all other components of wealth are very small in comparison. Note also that the distribution of livestock across households is skewed. A large majority of households own a small number of animals and the unevenness of the distribution is clear by the histogram in Figure 12 and the difference between mean and median value in Table 25. Animals are also a very liquid form of storing wealth since they can be easily sold and purchased in local markets.

Table 25. Components of households' wealth

	% holding	Median	Mean	Ratio of expenditure
Savings (bank and susu)	19.5	0	56.7	0.01
Livestock (Cedis 2013)	95.6	848	3008	0.71
Assets (Cedis 2013)	97.4	153	274	0.06
Land (hectares)	96.1	2.8	4.6	
Agricultural profits per hectare (Cedis 2013)		168	230	
Housing (rent)		183	195	0.05
Outstanding debt		0	8.4	0.04
Household expenditure (Cedis 2013)		3239	4263	

Note: stock figures (savings, livestock, assets, land and outstanding debt) refer to values reported (or imputed) at the time of the interview. Flow figures (agricultural profits, housing and expenditure) refers to values reported (or imputed) in relation to 12 months preceding the interview.

Household assets are a less liquid type of household wealth that is not easily sold or purchased. Our definition of household assets includes home furniture such as chairs, tables and beds; small appliances such as radios, refrigerators, TVs and lamps; personal appliances such as watches, computers and cameras; and vehicles such as carts, bikes or other motor vehicles. More than 97% of households have some assets, mostly in the form of basic furniture or small appliances (e.g. radios) and means of transport (e.g. bicycles). The average value of assets held is very small and amounts to little more than 5% of annual household expenditure. The distribution of assets is skewed though not as much as the distribution of livestock (Figure 12).

Figure 12. Distribution of livestock, assets and land

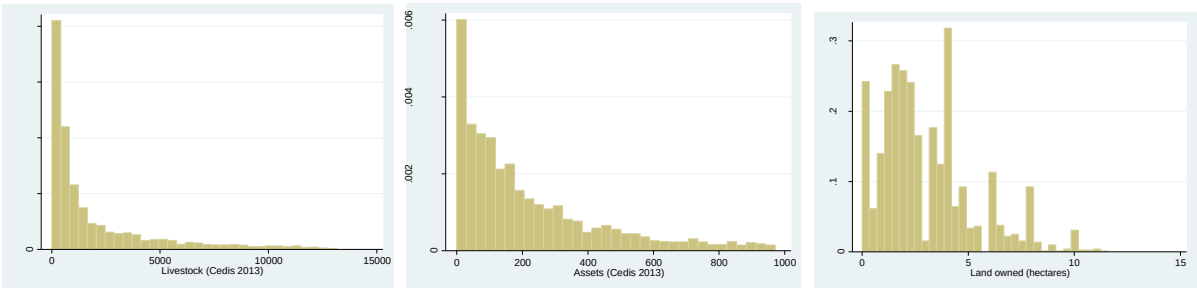

The charts are histograms. The y axis represents the frequency of observations with values falling in the range represented by the "bin" – the horizontal bars are reported on the x axis.

Most households (96%) own some land. Rights to land are rarely formalised through titles and are mostly customary. Land is assigned by local chiefs to families, and rights to land are inherited. Land is not scarce in the area and the average holding size is slightly less than five hectares per household. Much of the land is not cultivated and not always of good quality. Land does not operate as a storage of wealth because it cannot be bought or sold and a market for land does not exist.

All households own some dwelling, which is normally of very poor quality and without a toilet or access to water and electricity. Using national level data collected by the GSS through household surveys[14] we used dwelling characteristics (e.g. number of rooms, availability of water, electricity and latrines, and information on rental values) to predict rental costs of properties in the area. Home rents calculated in this way amount to 5% of total household expenditure, which reflects the very small value of properties. Rental of properties is unknown in the area and a housing market does not seem to exist.

[14] Ghana Living Standards Survey 6.

Since land and house properties cannot be used to store wealth, the accumulation of wealth and savings must occur through financial savings, livestock and assets. Home assets, as defined in our study, are not very liquid and therefore savings are likely to be made either through financial savings or animals.

4.3 Impact on savings and wealth

The project has a large impact on financial savings and a considerable impact on animal holdings (Tables 26 and 27). These, as well as the other project impacts reported in this section, are DD estimates of the difference between the changes in MV areas and the changes in CV areas. Financial savings in MV areas increased on average by more than 30 percentage points in comparison to CV areas. The impact was particularly strong in the second and third year of the intervention. Since financial savings represent a small component of overall wealth, this increase is likely to absorb only a fraction of any income increase that is not converted into consumption.

Table 26. Impact on household savings

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.396**	0.361**	0.562***
	(0.032)	(0.043)	(0.000)
DD effect second year	0.180	0.148	0.362**
	(0.344)	(0.419)	(0.021)
DD effect third year	0.470**	0.447**	0.616***
	(0.040)	(0.044)	(0.000)
DD effect fourth year	0.534**	0.498**	0.714
	(0.015)	(0.022)	(0.797)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

The project increased the value of livestock holdings by 20 to 30% in the MV areas (Table 27). The impact was particularly pronounced during the second year of the intervention when the impact might have been compounded by a positive income shock. Note that since on average livestock holdings are valued as much as 70% of annual expenditure, they could absorb a good portion of unspent income changes. The conversion of incomes into livestock holdings could happen directly through the purchase of animals or indirectly by feeding existing stock with the excess of food production generated by the increase in agricultural production.

Table 27. Impact on livestock holdings (sub-classification)

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.230	0.310	0.380**
	(0.261)	(0.208)	(0.004)
DD effect second year	0.235	0.259	0.292*
	(0.206)	(0.240)	(0.054)
DD effect third year	0.492**	0.679**	0.596***
	(0.021)	(0.010)	(0.000)
DD effect fourth year	0.179	0.293	0.255
	(0.493)	(0.323)	(0.196)
Sample size	11,658	11,658	5,829

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications

The project also increased the value of household assets (Table 28). Even though the percentage change does not appear to be very large, it is statistically significant in one of the three estimation models employed. The charts in Figure 13 show a clear shift to the right in the distribution of the overall value of assets in MV areas.

Table 28. Impact on assets holdings (sub-classification)

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.079 (0.554)	0.072 (0.564)	0.297** (0.001)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

Figure 13. Assets holding at baseline and midterm

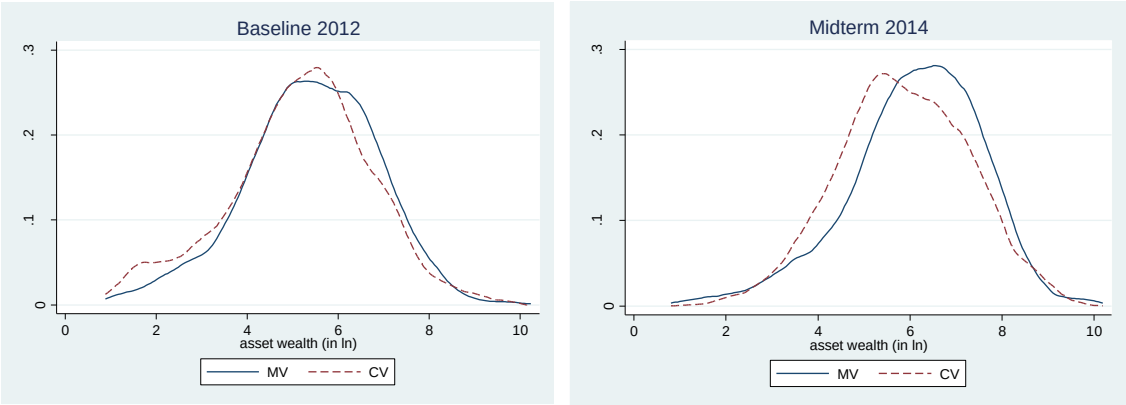

A better way of looking at changes in assets is to review the impact on the purchases of durable goods. The purchases of durable goods is excluded from the computations of total household expenditure. Durable goods are relatively infrequent and the value of the use of the goods is spread over a large number of years when performing computations of consumption expenditure. In addition, some durables are better described as investments and are not included in household expenditure since they do not increase welfare in any way. Our definition of durables, which is based on the data collected by the household surveys, includes home repairs and maintenance, power equipment, computers, mobile phones, household furniture, agriculture tools and equipment, motorised and non-motorised vehicles, dowries, and marriage and funeral expenses.

The project increased the purchase of durable goods by about 50% (Table 29). The impact was relatively small in the first year, extraordinary in the second year and negligible in the third year of the intervention. The extraordinarily large impact during the second year may suggest that the MVP's impact was compounded by some positive income shock in the area.

Table 29. Impact on consumption of durable goods (sub-classification)

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.524**	0.529**	0.358***
	(0.017)	(0.014)	(0.000)
DD effect second year	0.244	0.255	0.077
	(0.284)	(0.256)	(0.559)
DD effect third year	1.196***	1.189***	1.033***
	(0.000)	(0.000)	(0.000)
DD effect fourth year	0.134	0.144	-0.031
	(0.606)	(0.574)	(0.797)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

We also looked at household outstanding debt, measured as the difference between the amount of contracted principal debt minus repayments made over the 12 months before the survey. The project appears to have a positive impact on indebtedness in the first year of intervention and particularly in the second year, while no effect is found in the third year (Table 30). The impact might be the result of agricultural loans promoted by the intervention, which could have been invested in the purchase of livestock and other durable goods.

Table 30. Impact on household net borrowing

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.145**	0.142**	0.207***
	(0.024)	(0.029)	(0.000)
DD effect second year	0.161*	0.158*	0.224**
	(0.056)	(0.067)	(0.004)
DD effect third year	0.348**	0.338**	0.413**
	(0.008)	(0.010)	(0.002)
DD effect fourth year	-0.082	-0.080	-0.019
	(0.341)	(0.337)	(0.772)
Sample size	7,847	7,847	5,860

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

The analysis conducted in this section suggests that the observed discrepancy between changes in incomes and consumption might be explained by changes in savings and wealth. All figures relating to incomes, expenditure, assets and livestock are measured with considerable error and a full reconciliation of the figures by adding-up cannot be worked out. However, the analysis of the sign and relative size of the changes brought about by the intervention allows inferences about what appears to be rather puzzling household behaviour. In line with the permanent income hypothesis, households have interpreted the increase in agricultural incomes as a temporary event and have exploited the sudden increase in incomes to purchase livestock, assets and durable goods. Purchases have occurred with a precautionary saving motive (e.g. to protect against future income shocks) or with the goal of accumulating wealth and increasing income capacity. The impact on savings and wealth accumulation might have been compounded by generous agricultural loans made by the project in the first and second year of operation. Project impact on purchase of durable goods and livestock growth suggest that loans might have encouraged the purchase of animals and other goods.

Our analysis was conducted within the theoretical framework of major theories of consumption and saving behaviour. The data appear to fit more closely the permanent income theory of consumption, which suggests households save temporary income gains to smooth future adverse shocks. One implication of this analysis is that the impact of MV consists merely of making households more resilient to shocks, but as accumulated assets are sold off over time in the face of shocks, this impact is only temporary. It should be noted, however, that this conclusion, although supported by the data, is valid within the theoretical framework of the permanent income hypothesis. But households could invest the savings made and become permanently richer. This is not an outcome predicted by the theory and it is not currently supported by the data, but it cannot be excluded that it could occur in the future, particularly if shocks do not happen in the immediate term and people who save for a shock begin to see these savings accumulate.

5. Impact on farm productivity

This section further investigates one of the most remarkable impacts of the intervention: the large increase in household incomes. As discussed in previous sections, the increase in household income is likely to be driven by an increase in agricultural incomes and in this section we analyse more closely the determinants of this change. First, a causal change for the MVP package of agricultural interventions is presented. We then outline a methodology to analyse the observed impact on agricultural outputs resulting from changes in inputs, productivity, etc.

5.1 A causal chain of MV agricultural interventions

The stated goal of the MVP agricultural interventions is to improve agricultural incomes, food security and the development of an agricultural value chain.[15] Food security is not specifically defined in the project reports and we interpret food security as stable and sufficient availability of food at the household level. Food security is perceived by the project reports as a by-product of increased food production and better management of post-harvest losses. Training and cooperatives also may offer partial protection against food scarcity at any particular time. By improving food and nutrition security, the project expects to "increase incomes … of farming households in the SADA-MVP cluster".[16] The MVP aims to achieve these goals by investing in five key areas in the agriculture sector:[17]

1. Improving delivery of agricultural extension services
2. Improving access to physical agricultural inputs
3. Enhancing agronomic practices
4. Increasing access to agricultural credit
5. Strengthening farmer-based organisations and their linkages to markets

The Business Case for the northern Ghana MV identified the following agricultural problems: crop yields, prices and market access are subject to uncertainty and seasonal variability; low agricultural productivity; poor access to markets; farmers are unable to increase their income through agricultural inputs because they lack finance of their own and it is not possible to access credit from commercial lenders; and lack of inputs leading to soil nutrients becoming depleted after repeated cropping (environmental degradation).[18]

The MVP anticipates several changes will result from the activities to improve livelihoods, agriculture and food security in the MVs. By implementing multiple interventions generating impacts at various levels, a range of outputs are expected, including: ensuring the availability of water for the dry season; providing agricultural extension services; increasing access to farming inputs; strengthening market linkages; reducing pre and post-harvest losses; building strong cooperatives; and ensuring access to financial services. As a result of these outputs, several results will be achieved that will improve crop production, develop markets for agricultural products, and build the capacity of agro dealers, storage facility owners and civil society organisations to provide reliable products and services to all actors in the agricultural value chain. These in turn will increase agricultural output and improve value chains, which will have a positive impact on food security and increasing incomes in the MVs. The agricultural interventions, integrated with activities delivered across other sectors to generate synergies, will then enable people in rural areas to save and accumulate wealth, stimulating investment and diversification into non-farm work.[19]

[15] 2016, Logframe for the Millennium Villages Accountable Grant Programme, DFID.
[16] 2015, Mid-Year Report on the Millennium Villages Project in Northern Ghana, p. 7.
[17] 2016, Logframe for the Millennium Villages Accountable Grant Programme, DFID.
[18] 2011, Business Case: Millennium Village in Northern Ghana, DFID.
[19] Ibid.

To tackle the agricultural problems faced in northern Ghana, the MVP employed a generic set of activities aimed at achieving "quick wins" by delivering inputs, subsidising improved seeds of high yielding crop varieties or hybrids, training farmers on agronomic practices to eliminate "hunger months", forming cooperatives, and developing food storage options and markets.[20] The interventions are not connected to each other in a causal chain form, though they do contribute to improving profits in various ways. Each intervention affects profits under particular conditions and for particular farmer groups. However, we limit our analysis to the general expected effects of each intervention on agricultural profits. Agricultural profits are expressed as output prices (p), times output quantities produced (q), minus input prices (w), times input quantities used (x). Figure 14 illustrates how the different interventions affect specific components of agricultural profits.

The evaluation collected data on food security consisting of perceptions of food availability and seasonal food shortages. These data were collected at the baseline, at the midterm and will be collected again at the endline. The midterm data found a positive impact of the intervention on perceptions of food availability. In this report we focus on crop production as the main channel to increase agricultural incomes and food availability. The goal of the agricultural interventions is to increase agricultural productivity and crop production through a package including: input delivery or subsidisation of seeds and fertiliser, farmers' training, cooperative formation, food storage and market development. We discuss these tools one-by-one and Figures 14 and 15 illustrate a schematic theory of change of the interventions and main relationships.

- **Resource provision:** The MVP provides farmers with agricultural inputs such as seeds, fertiliser, water, tractor services and land preparation. In particular, the project provides improved seeds for the production of maize, soy and rice. Seeds and fertiliser are either donated or provided through loans made on concessional terms or with very low repayment rates (i.e. heavily subsidised). Therefore, they affect the quantity of output produced directly (q), while mildly affecting the quantities purchased (x) and costs. The provision of other inputs that are complementary to seed and fertiliser, like water and tractor services, was less systematic and less successful. Farmers are allowed to rent small tractors at below market rates, but only 10 tractors are available to loan and the subsidy amounts to 20% of the market cost. The project helps farmers prepare land for cultivation, but only for a limited number of rice plots. Finally, studies were conducted for micro-irrigation projects but never implemented because they were made redundant by government plans to build a new dam on the White Volta river, which will positively affect the whole area under the MVP. Additionally, to ensure agricultural extension agents (AEAs) can access the rural communities they serve, the MVP provides motorbikes and fuel stipends.[21]

- **Information provision and extension:** The MVP helped hire eight new AEAs, adding to the 14 that were already employed by the government. In addition, AEAs are given training and basic tools. They are supervised to increase efficiency and the time actually spent in the communities. AEAs work through more than 150 "lead farmers," who are selected in each community based on skills and motivation and are in charge of managing farmer groups of 15–20 members. Lead farmers are equipped with tools and training and charged with the task of training their farmers' group. Training relies heavily on farm visits and demonstration plots and includes sessions on planting, land preparation, weed control, harvesting, integrated soil fertility management and post-harvest management. The MVP expects training to increase profits by increasing farm productivity through increased production quantities (q).

[20] 2011, Business Case: Millennium Village in Northern Ghana, DFID.
[21] 2014, Annual Report on the Millennium Villages Project in Northern Ghana.

- **Cooperative development:** The MVP conducted a number of studies on agricultural systems and value chains to inform the selection of promising new crops and to improve market access. Large buyers for farmers' produce were identified and farmers receive training on market quality standards and requirements. Mango, maize, millet and acacia were identified as promising new crops and farmers are given saplings and training to grow them. Market development initiatives are expected to improve profits by giving farmers access to better prices and promoting the production of higher value crops (p). As a vehicle to achieving this objective, the project has given great attention to organising farmers through the formation and capacity building cooperatives. Farmers' cooperatives were formed in each community at the onset of the project and two cooperative officers were hired to support them. Cooperatives are formed following Ghanaian legislation on cooperatives. Cooperative members are trained by the project and the cooperative structure is used to channel agricultural loans to farmers. Many benefits may be generated by cooperatives, including: the opportunity for farmers to spread agricultural risk among members, increased access to credit via collective responsibility of loans, and increased negotiating power with traders in determining input and output prices. Cooperatives can help increase profits in several ways.

- **Infrastructure development:** The project addresses how to minimise post-harvest losses. Part of the losses are resolved by training farmers on proper harvest times. However, other losses are the result of improper storage methods or the absence of storage facilities. The project therefore rehabilitates warehouses or builds entirely new storage facilities. Improved storage has an immediate impact on quantities of output sold, as losses are reduced, but also helps prices as farmers have the opportunity to sell their produce when prices are more favourable. This intervention therefore helps profits by positively affecting output quantities and prices.

Finally, it should be noted that the MVP consists of a varied package of interventions in multiple sectors including education, health and infrastructure. Investments in other areas may benefit agricultural profits indirectly. For example, improvements in health of the population may result in a reduction of working time lost to sickness and a general increase of workers' productivity. Similarly, the rehabilitation of roads or construction of new roads can improve access to markets and help change prices to farmers' advantage. Other interventions may improve all components of profits by increasing total farm productivity (q), changing prices favourably for farmers (p and w) and providing better access to inputs (x).

Figure 14. Theory of change for agricultural interventions

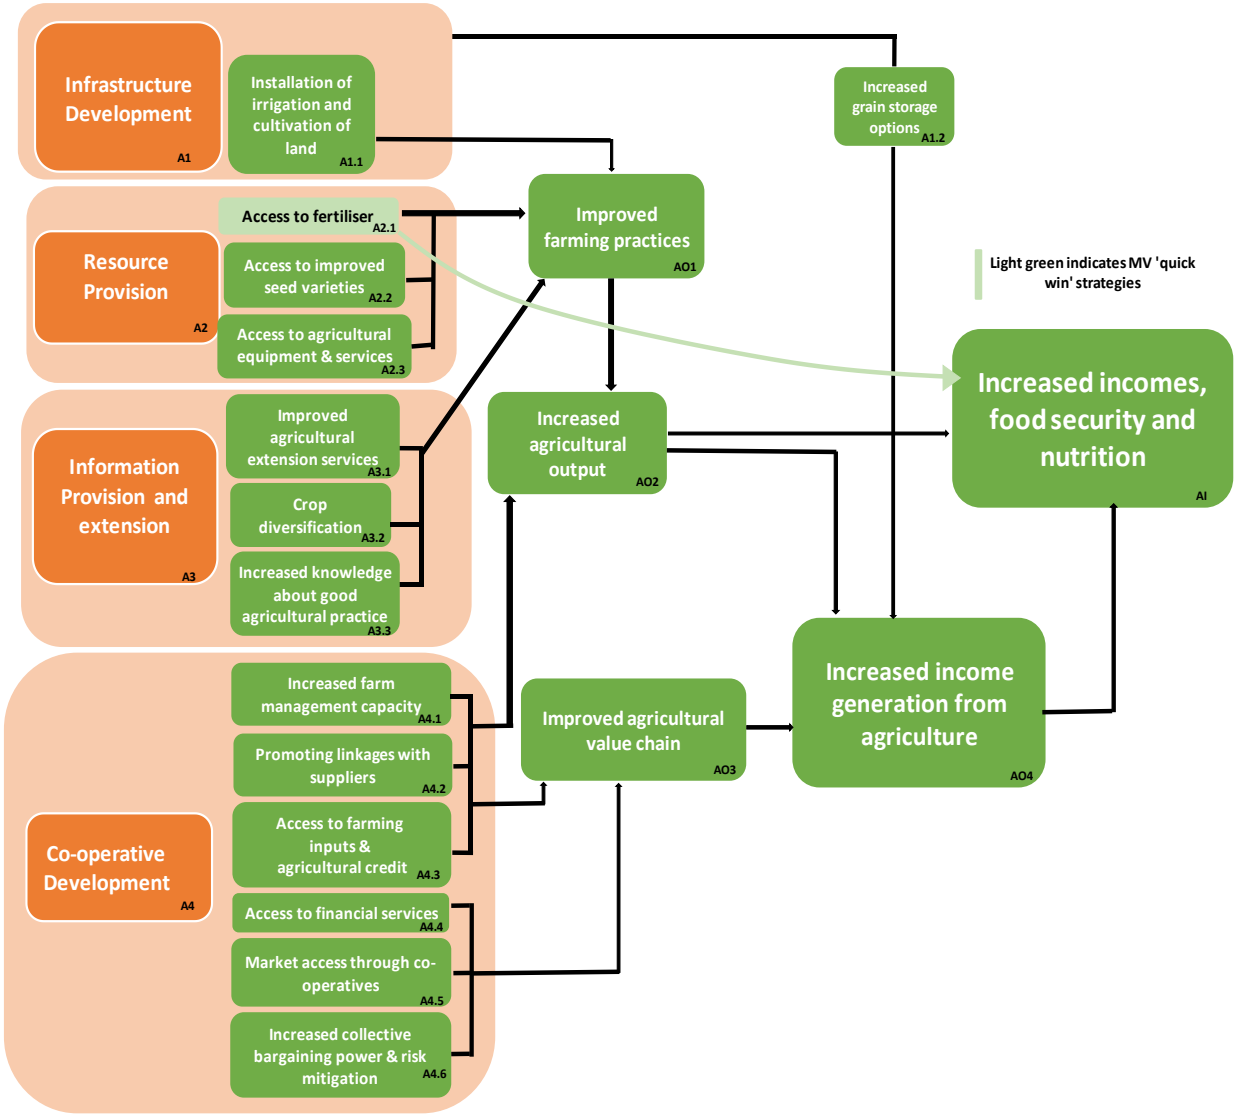

Figure 15. In-depth review of profits within the agricultural causal chain

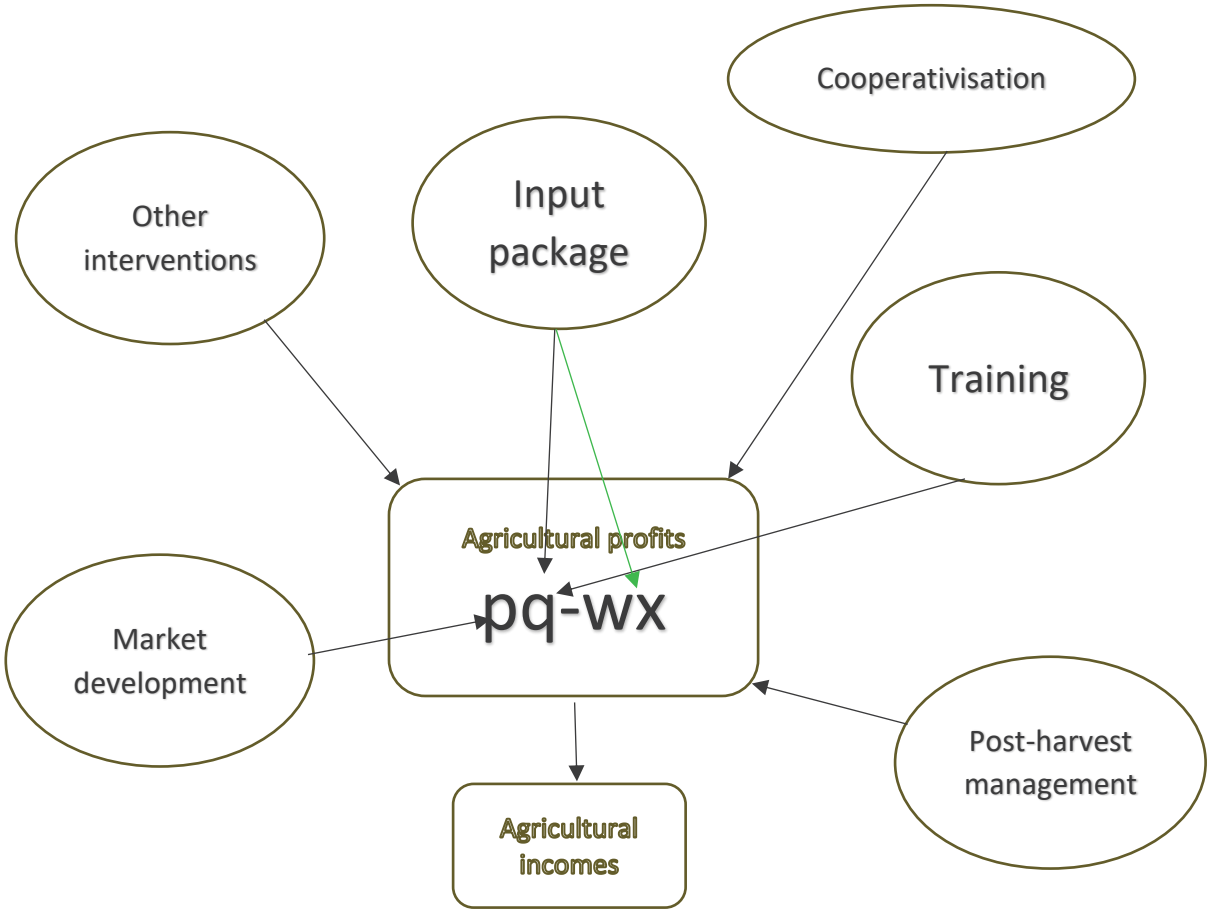

A summary of the MVP's known agricultural activities and achievements is detailed in Table 34. It should be noted that not all the activities may fit within the allocated categorisation since each output may contribute to multiple agricultural outcomes. Despite the complexities in specifying which outputs lead to various outcomes within the intervention logic, Table 31 attempts to group activities together according to where they are most relevant.

Table 31. Summary of MVP's agricultural activities to date[22]

Infrastructure	Resource provision	Information provision	Cooperative development
Demonstration plots (400)	Seeds	Outreach services	Cooperatives for farmer groups
Grain storage and warehouses	Fertilisers	Training (e.g. on financial management, conservation, using improved seeds, husbandry) (400 lead farmer trained)	Training (e.g. on market quality standards) (10 community entrepreneurs trained)
Land irrigation	Tractors		Setting up Village Savings and Loan Associations (50 groups with membership of 1,425)

[22] Numbers in brackets are the indicative number of achievements over the course of the project that have been documented in the 2016 SADA Mid-Year Report. These numbers need to be verified with SADA-MVP.

Infrastructure	Resource provision	Information provision	Cooperative development
Land cultivation (200 hectare paddy rice cultivated)	Motorbikes and fuel stipends for AEAs		

The assumptions underpinning the agricultural interventions include:

- Farmers and communities are open to new ideas and ways of working and living
- Increased income gained from increased production, greater market access, etc. will translate into reinvestment to sustain income increases into the next years
- Training, incentives, structures and systems created will be used effectively and taken up by the host communities in sufficient volume to realise change
- Interventions will empower the communities to participate and run groups, projects, businesses, be involved in decision-making, etc. both during and after the intervention
- The new crops, other agricultural activities, ways of working promoted by MVP are appropriate to the local context, farming systems, cultural norms, etc.
- Targeted communities are motivated to concentrate their efforts on agriculture as opposed to other income-generating activities

Heterogeneous effects

The estimation of project effects across groups may allow for the testing of links implicit to the causal chain depicted in the theory of change. In this section we make some suggestions on how the agricultural theory of change can be tested through sub-group analysis. Note, however, that here, as well as in other sections of the report, this approach is purely exploratory and we do not suggest that sub-group analysis provides causal explanation for the following reasons.

1. Subdividing the population by groups increases the probability of finding "false" statistically significant effects.
2. Sub-group analysis that was not defined in the Analysis Plan is by definition only exploratory and makes no causality claims. The Analysis Plan considered three key sub-group analysis of the data: sex of beneficiary, district of intervention (Builsa and West Mamprusi), and distance from the MV area ("near" and "far" control villages).
3. Sub-group analysis can generate very small groups with very limited statistical power and noisy estimations.
4. It should be noted that the sample of MV and CV villages belong to a cluster of villages located in a very small geographic area. Households and villages present a high level of homogeneity, thus limiting the scope for building different groups of observations. For example, nearly all households are cultivating land plots of a similar size and there is no difference in land property rights.

We identify the following potential sub-groups for further exploratory analysis to be conducted in the final evaluation report. First, the project is explicitly promoting some specific crops (maize and rice). We can disaggregate farmers based on the main crop produced at baseline and assess the impact of MV on different crop yields. Differences in impacts across crops may reveal differences in input delivery or in structural production constraints. Second, the sample of MV and CV areas can be subdivided in localities near and far from input and output markets. Farmers closer to markets may enjoy more favourable prices. These effects can be detected by estimating impacts on the size of agricultural marketed surplus and total crop sales. Finally, impacts can be disaggregated by sex of beneficiary in the following ways: female-headed households and productivity of land plots owned by

women. The analysis should reveal whether women are targeted by the intervention and whether women face constraints in transforming the project outputs into outcomes.

5.2 Estimation of the production function

To assess the impact of MVP on agricultural production we first investigate the impact on the quantity produced (*q*) rather than on profits for two main reasons: (a) most of the impact on profits should occur through changes in the quantities produced. The input package increases input use, while training and other interventions increase the productivity of inputs. Impacts on profits occurring through prices are less likely because they mostly rely on other interventions and on market development initiatives, which were not particularly successful; and (b) the estimation of a profit function requires good data on input prices, and our data on wages and other inputs are incomplete and inaccurate. Local community-level prices were not collected through market surveys but by interviewing "knowledgeable" individuals in the communities. One alternative is calculating average community-level prices from expenditures and quantities reported by farmers in the household questionnaire, which we will conduct in a separate exercise and report on this at endline analysis.

The focus on quantities produced rather than profits has some limitations. As observed in the previous section, the project promotes higher value crops and even common crops, such as maize, that can be produced at higher quality standards and therefore sold at higher prices. On the other hand, local market prices can be affected by local supply so that, for example, a successful increase in production may result in a reduction of prices. As stated above, based on our understanding of how the project operated in the field, we expect price effects to be relatively small. However, we reserve to conduct an analysis of the impact on profits at a later stage and investigate here the primal impact on quantity produced and productivities.

The simplest specification for the agricultural production function is the Cobb-Douglas form:

$$q = Ax^{\alpha}z^{\beta}$$

where *q* is the quantity of output produced, *x*'s are variable inputs such as fertiliser, seeds, and labour, and *z's* are fixed inputs such as land and productive capital. The α and *β* parameters measure production elasticities, that is the per cent increase in quantity produced for a per cent change in the quantity of input used. Finally, *A* is total factor productivity or "disembodied" technical efficiency, that is any contributions to production that are not embodied in the inputs included in *x* and *z*.

One advantage of the Cobb-Douglas form is that it can be easily estimated with ordinary least squares using a logarithmic transformation:

$$\ln q_i = \ln A + \sum_{j}^{n} \alpha_j \ln x_{ji} + \sum_{m}^{k} \beta_m \ln z_{mi} + \varepsilon$$

This specification allows for the separate estimate of the contributions to production of variable and fixed inputs (parameters αand β), and the contribution of any other factor (parameter A) or total factor productivity.

However, we cannot estimate agricultural production as the sum of quantities produced of each crop. Our farmers produce a variety of agricultural goods which cannot be simply added up unless they are transformed in values by multiplying quantities produced by their prices. The dependent variable is therefore the value of the agricultural production, that is, quantities of each crop produced multiplied by its price (*v*=*qp*). Production values are not quantities and make the dependent variable sensitive to price variations, as a higher production value may simply reflect the production of a higher value crop. To remove price effects we normalise the output value by a price index calculated at the household

level. The price index is a geometric mean of median village-level crop prices weighted by the farm-level production (value) share of each crop. The dependent variable of our production function is a normalised production value. We divide production values by the price index *P*, which is calculated in the following way using 22 village-level crop prices:

$$P_i = \prod_j^n P_j^{\frac{v_{ji}}{\sum v_{ji}}}$$

Variable inputs (*x*) are quantities of inputs used over the previous 12 month agricultural year. They include kilograms of seeds, chemical fertiliser, herbicides and pesticides used in all cultivated plots, and the number of days of own labour and hired labour in agriculture over the previous 12 months. Fixed inputs include land and capital. Land is measured in hectares of cultivated land while capital assets are measured as the value of the sum of the following production assets: animals (oxen, horses and donkeys), animal-drawn cart, tractor, plough, hoe, axel, shovel, spraying machine, sickle and power tiller.

We now show how an application of the standard Oaxaca decomposition to a DD analysis of the production function can help us understand how the MVP affects agricultural production. After this exercise it will be possible to separate the observed project impact on agricultural output in (a) a component resulting from changes in input used; (b) a component resulting from changes in returns to inputs (productivities); and (c) an otherwise unexplained component. It should be recalled that based on our understanding of the way the project operates that was described in the previous section: changes in inputs are mostly determined by the delivery of the package of inputs; changes in productivities are mainly obtained through farmer training; and changes in total factor productivity are determined by changes in other MV interventions. In order to explain impacts resulting from changes in input use by the delivery of the input package, impacts determined by changes in returns to input can be attributed to farmer training; any other unexplained impact can be attributed to an overall impact of the project on productivity via improvements in health or access to markets.

To illustrate our application of the Oaxaca decomposition to DD analysis, we start by taking first differences of the production function over two periods, thus relating changes in agricultural output to changes in inputs. We do this separately for the project (1) and control observations (0). In order to simplify notation we consider only a change over two periods, we ignore logarithms and we consider a single input (*x*). We employ the difference operator *d* to express changes in variables from one period to the next. The "differenced" production functions in MV and CV areas respectively are:

$$dq_1 = A_1 + \beta_1 dx_1$$
$$dq_0 = A_0 + \beta_0 dx_0$$

Subtracting the first expression from the second gives the DD estimator of programme impact:

$$dq_1 - dq_0 = A_1 + \beta_1 dx_1 - A_0 + \beta_0 dx_0$$

By adding and subtracting $\beta_0 dx_1$ we obtain:

$$dq_1 - dq_0 = A_1 + \beta_1 dx_1 - A_0 + \beta_0 dx_0 + \beta_0 dx_1 - \beta_0 dx_1$$

which simplifies to the familiar Oaxaca decomposition:

$$dq_1 - dq_0 = (A_1 - A_0) + \beta_0(dx_1 - dx_0) + (\beta_1 - \beta_0)dx_1$$

which decomposes the DD effect into (a) a component brought about by the difference in input changes $\beta_0(dx_1 - dx_0)$; (b) a component brought about by the difference in input productivities$(\beta_1 - \beta_0)dx_1$; and (c) a component otherwise unexplained resulting from changes determined by the project $(A_1 - A_0)$. Estimation over multiple periods simply requires the inclusion

of time variables for each survey round, which capture otherwise unexplained changes from one round to the other. In our analysis we use four survey rounds, three year-to-year changes and therefore include two time dummy variables in the estimated regressions.

We start by estimating the impact of the MVP intervention on agricultural output using the same DD methodology based on sub-classification that was adopted in the previous sections (Table 32). The outcome considered is the value of agricultural production normalised by a price index. The project has a large impact on agricultural output. The average effect is 0.36, meaning that, on average, every year agricultural output is 36% higher than at baseline in MV areas in comparison with CV areas. The disaggregation of the impact by year shows that the impact of the intervention was particularly strong in the third year and similarly good in the second and the fourth year.

Table 32. Impact on consumption of durable goods (sub-classification)

	Cross-section	Fixed effects	Lagged model
Average DD effect	0.358***	0.342***	0.259**
	(0.000)	(0.000)	(0.004)
DD effect second year	0.321***	0.286***	0.203**
	(0.000)	(0.000)	(0.014)
DD effect third year	0.485***	0.467***	0.385**
	(0.000)	(0.000)	(0.002)
DD effect fourth year	0.270**	0.271**	0.188*
	(0.020)	(0.015)	(0.094)
Sample size	7,847	7,628	5,721

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

We then estimate the baseline production function described above separately for MV and CV areas (Table 33). We test the difference in the coefficients of the production functions of MV and CV and find that the production functions in the two areas are structurally equivalent at the baseline. None of the T-statistics testing the equality of the coefficients across the two equations are significantly different at the 5% level (third column of Table 28) and a F-test of joint equality of all coefficients is not rejected.

Table 33. Agricultural production function in MV and CV areas at baseline

	MV areas	CV areas	Test of difference
Fertiliser	0.020	0.036***	1.12
	(0.115)	(0.000)	(0.264)
Seeds	0.092**	0.166***	1.70
	(0.003)	(0.000)	(0.092)
Herbicides	0.030**	0.031**	0.06
	(0.001)	(0.013)	(0.950)
Pesticides	0.035**	0.041***	0.31
	(0.039)	(0.000)	(0.756)
Labour	0.314**	0.308***	0.06
	(0.003)	(0.000)	(0.955)
Land	0.104**	0.099**	0.07
	(0.042)	(0.006)	(0.942)
Capital	0.017	0.009	0.47
	(0.259)	(0.281)	(0.639)
Constant	4.698	4.664	0.05
	(0.259)	(0.281)	(0.958)
R-square	0.415	0.466	
Observations	665	1,242	

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. F-test of joint significance of all joint coefficients was rejected (F=0.98, P-value=0.453)

Finally we decompose the impact of the intervention on agricultural output using the standard Oaxaca decomposition approach described above. The results are shown in Table 34. The first column shows the impact of the intervention on agricultural output without including any control variables and estimates an average impact of the project using a first difference estimator in the same way as the estimators of Table 33 above. We find an average impact of 0.35 of the intervention. Since the dependent variable is in logarithms, this is roughly equivalent to a 35% increase.

The second and the third column of Table 34 estimate changes in the production function over time for the MV areas and the CV areas separately. The results of these regressions are the basis for performing the decomposition of effects, which are presented in the fourth and fifth columns. The fifth column shows the changes in agricultural output that are explained by changes in input use. These effects represent how much the output in CV areas would increase if farmers in CV areas increased the use of inputs by the same amount of the observed amounts among farmers in the MV areas. The majority of input coefficients are statistically significant and the input increases together explain 76% of the overall observed change in agricultural output. Fertiliser, seeds and land appear to make the largest contributions to the increase in agricultural output in the MV areas. The fourth column shows the changes in output that are not explained by changes in inputs. The Oaxaca decomposition decomposes “unexplained” changes that are not explained by changes in inputs in changes in returns to inputs (fifth column of Table 34) and other unexplained changes (MV and time dummy variables in the fourth column). In terms of the production function equation above, these changes correspond respectively to changes resulting from differences in the parameters (the production elasticities α and β) and changes resulting from differences in total factor productivity (parameter *A*). Changes in returns to factors are very small, often of negative sign (pointing to a decrease in input productivity in MV areas) and never statistically significant. Other unexplained changes occurring over time have a relatively small impact on the increase in agricultural output and the coefficients are not statistically significant.

This analysis suggests that much, if not all, the observed improvement in agricultural output observed in the MV areas is the result of an increase in input use, which in turn is likely to be a result of the

package of inputs (fertiliser and seed) freely distributed or heavily subsidised by the project. A good fraction of the increase in production appears to be the result of an increase in the use of land and labour. Land use is regulated by customary rights in the area and is available to expand cultivation to most farmers. Family labour is also abundant in the area given the large family sizes and the few employment opportunities available outside agriculture. This result suggests that the provision of inputs leads to a large use of land and labour, which in turn implies that labour and land productivity are unlikely to increase as they are used more extensively rather than more productively.

Farmer training and other MVP interventions could potentially increase farmers' productivity. However, we do not observe an improvement in the productivity of inputs. None of the input productivities in the MV areas increase in comparison to CV areas. The sign of changes in productivities are negative for fertiliser, seeds and land, suggesting that these inputs might not always be used efficiently. Finally, there is a residual unexplained positive effect of the project on agricultural output that is not embodied in changes in inputs. These are changes attributable to other MV interventions, such as roads or irrigation. The effect, however, is not statistically significant and not very large.

The fact that the increase in agricultural output is mainly the result of an increase in the use of fertiliser and seeds provided by the project raises some important questions about the sustainability of the intervention. It is not certain that high levels of input use will be maintained by farmers once free delivery or subsidising of inputs by the project is discontinued.

Table 34. Oaxaca decomposition of the impact of MV interventions on agricultural output

	DD effect without control variables	Production function in MV areas	Production function in CV areas	Unexplained changes	Changes in inputs
	$\beta_1 MV$	$A_1 + \beta_{21}T_{21} + \beta_{31}T_{31} + \sum \beta_1\, dx_1$	$A_0 + \beta_{20}T_{20} + \beta_{30}T_{30} + \sum \beta_0\, dx_0$	$(A_1 - A_0) + \sum (\beta_1 - \beta_0) dx_1$	$\sum \beta_0 (dx_1 - dx_0)$
MV total difference	0.352***	0.002	-0.088**	0.089	
	(0.000)	(0.976)	(0.008)	(0.367)	
T2	-0.205***	-0.065	-0.250***	0.062	
	(0.000)	(0.345)	(0.000)	(0.109)	
T3	-0.184***	-0.101	-0.064	-0.012	
	(0.000)	(0.145)	(0.165)	(0.675)	
Fertiliser		-0.000	0.014***	-0.032	0.037**
		(0.926)	(0.000)	(0.171)	(0.001)
Seeds		0.105***	0.159***	-0.006	0.074**
		(0.000)	(0.000)	(0.436)	(0.029)
Herbicides		0.009	0.007*	0.006	0.010
		(0.151)	(0.085)	(0.898)	(0.210)
Pesticides		0.035***	0.049***	-0.017	0.032*
		(0.000)	(0.000)	(0.531)	(0.082)
Labour		0.303***	0.261***	0.008	0.034*
		(0.000)	(0.000)	(0.698)	(0.077)
Land		0.105***	0.126***	-0.020	0.083**
		(0.000)	(0.000)	(0.730)	(0.030)
Capital		-0.001	-0.005	0.005	-0.001
		(0.926)	(0.439)	(0.851)	(0.711)
Component of difference				0.083	0.269***
				(0.333)	(0.000)
R-square	0.017	0.343	0.384		
Observations	5,721	1,995	3,726		

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance.

6. Impact on female-headed households

In this section we conduct a gender analysis of the results presented so far by disaggregating households into male-headed and female-headed households. Household head changes over time and therefore we consider in our analysis the head status at baseline. The observed impact is therefore the impact of MV on those households that were female-headed before the intervention. Our definition of female-headed households includes: single widows, polygamous widows, female heads whose husband is away and whose husband is present, and heads who are either unmarried, single or living with a co-worker.[23] There were 10.3% such households at the baseline, with a slight majority in the comparison villages (11.0%) than in the project villages (8.9%).

Since female-headed households are a relatively small sub-sample of the group, we are unable to use the sub-classification methods of analysis employed in the previous sections. Instead, we estimate project effects using inverse probability weights based on the propensity score by running regressions with a female-headed dummy and interaction terms for different years of the intervention. The coefficient estimates reported in Table 35 represent differences in impact for female-headed households in comparison to male-headed households. To simplify descriptions we use only the cross-sectional model specification and we estimate the effects for: per adult equivalent expenditure, per capita income, school attendance ratios and in- and out-migration. Note that since the estimations reported in Table 35 are based on inverse probability weighting, the results often differ in size from those presented in the previous sections.

The introduction of a large number of statistical tests generates a multiple testing problem, whereby some statistical difference may emerge as a result of chance and should be interpreted with caution. We mainly focus therefore on the aggregate average difference result reported in the top row of each table. Expenditures and incomes of female-headed households are not affected in any particular way in comparison to male-headed households. In the case of school attendance, the positive impact of MV on primary school attendance does not appear to differentially affect children of female-headed household. On the other hand, there seems to be a positive impact on attendance of junior secondary school by children of female-headed households, but these effects are not statistically significant. There are clear differences in movements in and out of the households that depend on the particular demographic structure of female-headed households. By definition, female-headed households have a smaller number of male adults and therefore less likely to have male adult migrants. Households headed by females in the project areas report larger in-migration and smaller out-migration.

Table 35. Impact on income and expenditure by sex of head of household

	Per adult eq. expenditure		Per capita income	
	Male-headed	**Female-headed**	**Male-headed**	**Female-headed**
Average DD effect	-0.021	0.037	0.465***	-0.008
	(0.732)	(0.516)	(0.000)	(0.971)
DD effect second year	-0.016	0.065	0.342***	0.460**
	(0.825)	(0.360)	(0.000)	(0.032)
DD effect third year	0.065	0.064	0.489***	-0.070
	(0.414)	(0.534)	(0.003)	(0.811)
DD effect fourth year	-0.114*	-0.020	0.565***	-0.420
	(0.085)	(0.766)	(0.005)	(0.228)

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

[23] These are the definitions used in the Household Survey instrument.

Table 36. Impact on age-specific school attendance by sex of head of household

	Primary attendance		JSS attendance		SSS attendance	
	Male–head	Female-head	Male–head	Female-head	Male–head	Female-head
Average DD effect	0.065*	-0.092	0.044	0.087	0.008	-0.016
	(0.064)	(0.134)	(0.190)	(0.181)	(0.636)	(0.710)
DD effect second year	0.109**	-0.023	0.020	0.003	0.009	-0.047**
	(0.012)	(0.773)	(0.535)	(0.958)	(0.624)	(0.032)
DD effect third year	0.039	-0.144**	0.090**	0.135	0.010	0.030
	(0.317)	(0.047)	(0.042)	(0.193)	(0.722)	(0.689)
DD effect fourth year	0.051	-0.091	0.023	0.130	0.007	-0.033
	(0.203)	(0.237)	(0.559)	(0.147)	(0.736)	(0.498)

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

Table 37. Impact on in-migration and out-migration by sex of head of household

	In-migration		Out-migration	
	Male-headed	Female-headed	Male-headed	Female-headed
Average DD effect	0.205	0.346	0.114	-0.334**
	(0.105)	(0.327)	(0.333)	(0.020)
DD effect second year	0.141	-0.358**	0.036	-0.367***
	(0.387)	(0.049)	(0.761)	(0.003)
DD effect third year	0.309	0.175	0.078	-0.213
	(0.117)	(0.716)	(0.597)	(0.397)
DD effect fourth year	0.166	1.235*	0.230*	-0.419**
	(0.390)	(0.075)	(0.099)	(0.016)

P-values in parentheses based on cluster-adjusted standard errors. *** is statistical significance at 1%, ** is 5% significance and * is 10% significance. Standard errors are calculated using 500 bootstrap replications.

Conclusions

The fourth round of quantitative data collection show that the MVP has made considerable impacts in some intervention areas and achieved less in others. Areas where the data show that there has been positive change are summarised below.

- Impacts on agricultural incomes are the largest and per capita income has increased on average by at least 50%.
- Income gains were mostly saved and invested, suggesting that households are accumulating wealth. There has been a large impact on financial savings and a considerable impact on animal holdings. Financial savings increased by more than 30% and livestock holdings increased by 20 to 30%.
- The project has successfully increased efforts to promote social mobilisation in the MVs, leading to the formation of new groups such as women's group, savings groups and to an increase in overall participation in other groups.
- The project has also been successful in providing health services and reaching a large majority of the intended population. These include services that were provided less successfully during previous years, such as deworming and vitamin A.

Areas where the project has had little or no impact are summarised below.

- There was no impact on consumption and poverty. Economic theories predict that extremely poor households would immediately spend any income gain brought about by the intervention to satisfy basic needs, but this has not been the case.
- There was no significant change in either poverty or extreme poverty as defined by the GSS.
- The impact on primary school attendance (an increase by 10%) was found in the first year of project implementation, and an increase in junior secondary school attendance by 8% was found at the midterm. The overall impact on basic education is nevertheless modest and there is no impact on senior secondary school attendance.
- In terms of migration, the MVP increased in-migration and reduces out-migration. It does not, however, seem to significantly increase migration in either direction.

References

DEATON, A. 1991. Saving and Liquidity Constraints. *Econometrica,* 59**,** 1221-1248.

DUESENBERRY, J. S. 1949. *Income, Saving and the Theory of Consumer Behaviour,* Cambridge, Harvard University Press.

FRIEDMAN, M. 1957. *A Theory of the Consumption Function,* Princeton University Press.

HILL, P. 1986. *Development Economics on Trial: the Anthropological Case for a Prosecution,* New York, Cambridge University Press.

IMBENS, G. W. & RUBIN, D. B. 2015. *Causal Inference for Statistics, Social and Biomedical Sciences: an Introduction,* New York, Cambridge University Press.

KEYNES, J. M. 1936. *The General Theory of Employment, Interest and Money*.

MODIGLIANI, F. 1966. Life Cycle Hypothesis of Saving, Demand for Wealth and Supply of Capital. *Social Research,* 33**,** 160-217.

ODI, 1979. *Integrated Rural Development*, Briefing Paper No. 4 Dec 1979.

PRESTON, S. H., HEUVELINE, P. & GUILLOT, M. 2001. *Demography: Measuring and Modellling Population Processes,* Oxford, Blackwell Publishers Ltd.

SKOUFIAS, E. 2005. *PROGRESA and Its Impacts on the Welfare of Rural Households in Mexico*, International Food Policy Research Institute Research Report 139, Washington D.C.

STUART, E. A. 2010. Matching Methods for Causal Inference. *Statistical Science,* 25**,** 1-21.

Appendix 1. Price deflators

Table A1.1 Price deflators

Year	Builsa project	Builsa control	West Mamprusi project	West Mamprusi control
2011/2012	0.969	0.978	0.943	0.958
2012/2013	1.000	1.000	1.000	1.000
2013/2014	1.090	1.090	1.141	1.141
2014/2015	1.165	1.165	1.301	1.301

Appendix 2. Sensitivity analysis

Matching methods assume that adjusting for differences in pre-treatment values of determinants of the outcomes removes bias in outcomes comparisons between a project and a control group. This assumption is known as *unconfoundedness* and cannot be tested. However, we can perform some assessments that can make the hypothesis of unconfoundedness less credible and plausible. In other words, we can probe the validity of the hypothesis but we cannot prove it. We consider here two methods to assess unconfoundedness. If the data fail these assessments, then our assumptions about our ability of removing bias by adjusting for differences in the covariates distributions of the project and control group need to be revised.

A2.1 Estimating impact on pseudo-outcomes

The first method consists of estimating the impact of the intervention on outcomes that are known not to be affected by the intervention. Consider dividing the set of covariates X in potential outcomes Xp and remaining Xr covariates. This assessment consists of estimating project effects on the Xp covariates using the Xr covariates for the estimation of the propensity score and for regression adjustment. The assessment can be made more restrictive in two ways. The first approach consists of subdividing the analysis of the outcome variables on intervals of the value of the variable. For example, in the case of a binary variable, we could test that the impact of the intervention on the potential outcome at different quintiles of the distribution are jointly zero. The second approach consists of testing the equality of the averages of the potential outcomes in subpopulations, of the treatment and the control groups. For example, we could test the impact on pseudo-outcomes across sub-groups, such as, sex and region. Both approaches are sensitive to the multiple testing problem, as the probability of finding a statistically significant difference increases with the number of tests conducted. We do not adopt either of these two most restrictive versions of the assessment in our application.

One difficulty in implementing this approach with our data is that it is difficult to imagine a covariate that is not affected by the MV programme over time. The programme consists of a set of multiple interventions that affects different outcomes in different ways. All the baseline covariates used in building the propensity score could in principle be affected by the intervention. Therefore, we cannot use changes in baseline covariates over time as pseudo-outcomes. One exception is the household being affected by a weather shock such as a flood or drought, as this is clearly independent of the intervention. However, shocks are reported by households and the project could change the actual impact of weather shocks or the perception of the impact of such shocks. In addition, these events affect many households in localised areas and the impact on these events may not be efficiently adjusted by additional covariates.

The ideal setting for this type of assessment is when lagged values of the outcome variables are available. Our data contain lagged values of some variables based on retrospective questions. In particular, school attendance rates can be calculated in relation to the 12 months before the survey. The baseline survey also collected retrospective data for several components of household income in relation to one year and two years before the survey. In particular, retrospective data sets were collected for: livestock holdings, a sub-sample of agricultural crops, micro-enterprise profits and wages. These four income components were not collected at the same level of detail that would make them comparable to the same component of the full baseline income data. In our application, we look at three pseudo-impacts: (a) the impact on primary net attendance ratios one year before the survey; (b) the impact on the value of livestock holding over two years before the survey; and (c) the impact on micro-enterprise profits one year before the survey. The results of these analyses are reported in Table A2.1 below. The estimation of the project impact on pseudo-outcomes reveals no surprises. As expected, we do not find an impact of the intervention in the years before the intervention was

implemented for the three potential outcomes considered. None of the impact coefficient estimates is large in size and none is statistically significant.

Table A2.1 Pseudo-impact of the MV project (sub-classification method)

Outcome	Coefficient	Standard error	P-value
Net attendance ratio (primary)	0.01	0.02	0.427
Value of livestock holdings	-17.09	136.88	0.901
Micro-enterprise profits	5.66	5.70	0.321

A2.2 Assessing project effects on pseudo-treatments

The second method to assess unconfoundedness consists of testing the impact of interventions that are not implemented. The easiest way to conduct this assessment is by focusing on sub-groups of the control groups that do not receive the intervention. Consider three groups: project (p), control 1 (c1) and control 2 (c2). Households in G=c1, c2 receive the control treatment (p=0), while units in G=p receive the project treatment (p=1). We define group unconfoundedness as the group being independent of the outcomes given the covariates:

G _||_ Y1, Y0|X

Group unconfoundedness implies the testable restriction:

G _||_ Y|X, G includes (c1, c2)

The ideal set up for this assessment is to have two control groups that are systematically different and likely to show biases in comparisons that are not adjusted by the covariates. Since none of the two control groups received the treatment, the project should have no impact on the outcomes of the two groups. Finding no evidence of any difference between the two groups does not imply unconfoundedness, but makes it more plausible.

In our application, there is an obvious subdivision of the control group which consists of the far-away control villages in the Builsa and West Mamprusi districts. Villages located in the vicinity of the project areas cannot be used in this exercise because they are potentially affected by the intervention. On the other hand, far-away villages in the control group are sufficiently distant from the project areas, so as to rule out a significant impact of the intervention on most outcomes. At the same time, the two far-away control regions belong to two different districts that are different in many ways and likely to exhibit biases in treatment-control comparisons. In our application, we look at three key outcomes: (a) net attendance ratios in primary school; (b) per capita expenditure; and (c) per capita income. The results are reported in Table A2.2 below. As expected, we do not find an impact of the pseudo-intervention on school attendance and expenditure in the Builsa far-away control group, in comparison to the West Mamprusi far-away control group.

We do find, however, a surprising impact on per capita income. The impact is large and statistically significant at 10%. Note that this result does not imply that unconfoundedness does not hold. Nor does it imply that the observed impact of MV on income is biased, but it does raise questions on the interpretation of the observed impact of MV on income. The impact of the pseudo-intervention on income suggests that there are covariate factors affecting income at a geographic level, that are not adjusted for by the set of covariates included in our matching strategy. One obvious possibility is weather shocks. Weather shocks reported by households are included in the regression adjustment, but either they are poorly measured or they are not the only covariate shocks affecting income. Apart from weather shocks, other unobserved structural differences between “near” and “far” control

villages, such as for example higher land productivity, may be responsible for different time trends. In conclusion, this result gives some plausibility to the hypothesis that the impact of MV on incomes observed in our analysis, could be the result of a positive covariate shock independent of the intervention or unaccounted for structural differences between MV and CV communities. This hypothesis will be further explored in the analysis of the endline data with additional sensitivity analysis.

Table A2.2 Impact of pseudo-interventions

Outcome	Coefficient	Standard error	P-value
Net attendance ratio (primary)	0.04	0.04	0.342
Per adult equivalent expenditure	-0.02	0.12	0.873
Per capita income	-0.38*	0.20	0.051