MATHEMATICA
Policy Research, Inc.

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

MEMORANDUM

TO: Rebecca Tunstall

FROM: Ken Fortson; Anu Rangarajan

DATE: 9/12/2008
MCC-Armenia

SUBJECT: Rural Road Rehabilitation Project Evaluation Plan

The Millennium Challenge Account with Armenia (MCA) aims to increase household income and reduce poverty in rural Armenia through improved performance of the country's agricultural sector. Armenia plans to achieve this goal through an integrated, nationwide initiative to improve major components of the rural infrastructure, with a focus on lifeline roads and the irrigation system, supplementing these infrastructure improvements with training in agricultural practices. By improving living standards among rural residents, these investments can in turn lead to future economic growth in rural areas and throughout the country as a whole.

MCC has commissioned a rigorous impact evaluation to separately examine each of the three main components of the MCA-Armenia program. This memo describes our design for evaluating the Rural Road Rehabilitation Project (RRRP).

The implementation of the RRRP will expand the access of rural communities to agricultural markets and social infrastructure, as well as increase non-farm income opportunities by improving the condition of rural roads. The Compact was designed to include MCC funding to rehabilitate up to 943 km of the rural roads from 'very poor', or 'poor', to 'good' condition. These 943 km are distributed amongst 85 road links. The road links were divided into three packages based primarily on estimated traffic on the road, though several roads were reallocated into other packages with nearby roads to facilitate efficiency. Construction for the higher-traffic road package was originally scheduled to begin in the summer of 2008, followed by the second and third packages not long afterwards.

An economic rate of return (ERR) was estimated for each of the road links, and in order to be funded, a road's ERR must meet or exceed 12.5 percent. The ERR is calculated from several inputs, including the vehicular traffic, vehicle operating costs, and the cost of the project, among others. If a road's ERR meets or exceeds this threshold, it is eligible for rehabilitation; if not, it will not be funded by the Compact. **Our core research design exploits the ERR threshold by comparing roads that just passed the ERR threshold to those that fell just short of it.**

Two important practical considerations complicate this core design. First, due to budgetary constraints—primarily stemming from the appreciation of the Armenian dram—some roads that pass the ERR threshold may still not be funded. Second, other funding sources, such as the

Asian Development Bank or the Japanese Development Bank, may fund some of the other road projects originally planned by MCA. (Road package 1 is already being co-financed by the GoA.) These two general issues will necessitate adjustments to the core design; however, we do not yet know exactly what the circumstances will be.

We discuss our proposed research design in more detail in the sections that follow. The key research questions guiding our design of the RRRP evaluation are:

- Did the RRRP affect the quality of and accessibility to roads?
- Did the RRRP affect agricultural productivity and profits?
- Did the RRRP improve household well-being for communities served by these roads, especially income and poverty?

We first provide an overview of the regression discontinuity (RD) design that serves as the foundation of our analytic approach. We then discuss practical considerations that would likely require adjustments to the core design, and what those adjustments might be. We then discuss the main source of data, the Integrated Survey of Living Standards (ISLS). We next discuss in detail our econometric approach for estimating RRRP impacts, and alternative specifications that can be employed to further explore the key research questions. We conclude with an approximate timeline for the impact evaluation report.

I. OVERVIEW OF REGRESSION DISCONTINUITY APPROACH

An RD design is applicable when roads are selected for rehabilitation based on a "scoring" rule, where roads with eligibility "scores" above a threshold value are determined eligible for rehabilitation and those with scores below the threshold value are not (or vice versa). In this context, the ERR of each road would be the scoring variable. Roads with ERRs above 12.5 are eligible for rehabilitation, while those with ERRs less than 12.5 are not eligible. **The RD approach is most defensible if the scoring rule—in this case, the ERR threshold—is strictly followed.** The beauty of an RD design is that, unlike other comparison group designs, the selection rule for determining eligibility is *fully* known.

The most rigorous RD design would be to compare the mean outcomes of roads with ERRs *just below* the threshold value (that is, those who are just eligible for rehabilitation) to the mean outcomes of roads whose ERRs are *just above* the threshold value (that is, those who are just ineligible). However, if we were to limit the analysis to roads near the ERR threshold, we would not have a large enough sample to precisely detect the impacts of MCA's RRRP.

Instead, we have included all roads that MCA initially considered for rehabilitation, across the *full range* of the ERR distribution. In order to do this, we will estimate impacts using regression models where the outcomes of the treatment group (communities where roads are rehabilitated) are compared to those of the comparison group (communities whose roads are ultimately not funded), controlling for a smooth function of the ERR. If we assume that the outcome variables exhibit a stable continuous relationship with the ERR and that we have correctly modeled this relationship, the sharp discontinuity in the ERR at the cutoff value, conditional on this smooth function of the ERR, will identify the RRRP's impacts.

Figure 1 demonstrates this estimation approach graphically, where hypothetical post-rehabilitation poverty data are plotted against hypothetical ERRs. The figure also displays the fitted regression lines for the funded roads (the treatment group) and unfunded roads (the comparison group), where for simplicity, the slopes of the two regression lines are assumed to be the same (although this condition can be relaxed). The estimated impact is the vertical difference between the two regression lines at the ERR threshold value of 12.5, that is, at the point of discontinuity.

Figure 1. The RD Method Visually

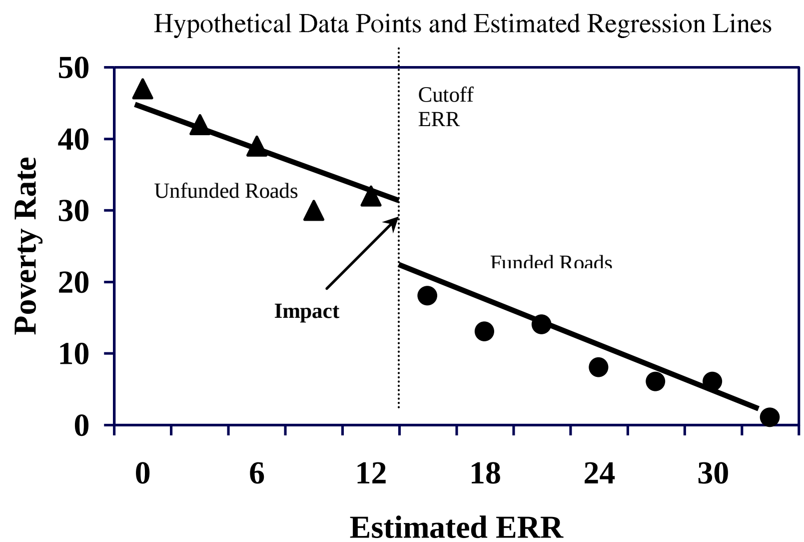

The RD approach is designed to estimate impacts on roads *near the ERR threshold*. Thus, an important interpretation issue is that an RD design only provides heuristic evidence of RRRP impacts for the *average* road.

II. ADJUSTMENTS TO THE RD DESIGN

The RD design described in the previous section will most likely form the basis of our evaluation design. However, a handful of practical complications preclude implementing this approach without some adjustments:

- ***Appreciation of the Armenian dram relative to the US dollar, general inflation and increases to unit costs for road construction materials have reduced the budget available for road rehabilitation.*** As a consequence, MCA anticipated that almost no roads can be rehabilitated beyond the first package. Current estimates project that several roads in the later two packages will pass the ERR threshold, but funds will probably not be available through the Compact to pay for these roads.

- ***Roads in the later packages may have higher ERR thresholds.*** While MCA was projecting that almost none of the roads from the second and third packages will be funded, this could change. If the dram/dollar exchange rate were to increase, for example, more roads could be rehabilitated. Unless funding greatly increased, however, it is unlikely that *all* roads meeting the ERR threshold could be afforded—only a subset of them would be rehabilitated.[1]

- ***Other donors may rehabilitate unfunded MCA roads.*** The Asian Development Bank and Japanese Development Bank are also funding other road projects, and could possibly fund some of the MCA roads that cannot be afforded with the Compact budget. Completed designs for all RRRP package 2 and 3 roads are available to other donor sources.

- ***Some communities are served by more than one road.*** A few villages are served by two different roads. Sometimes both are be funded by MCA; other times, one may be funded by another donor.

Because nearly all roads in the first package met the ERR threshold, an analysis focusing exclusively on those roads would not be sufficient to estimate program impacts. We would need to incorporate other roads in the analysis. We have identified two general approaches, and we discuss these options below along with the conditions under which each option is feasible.

---

[1] Additionally, the ERRs calculated for later packages may not be comparable to the ERRs for Package 1 roads. This complicates the empirical model slightly but does not harm the validity of the design.

Incorporating Roads from the Second and Third Packages. If all roads in the second and third packages that meet the ERR threshold were funded, incorporating roads from these packages would be seamless. More likely, however, only a subset of the eligible roads could be afforded. **These roads could still be incorporated into the RD design so long as the roads with the highest ERRs were the ones selected for rehabilitation.** In other words, the threshold for roads in the later packages funded by MCA would be *higher* than the 12.5 percent ERR threshold for roads in the first package. In this case, our analysis would essentially pool together two analyses: one based on roads in the first package, with a discontinuity at 12.5 percent, and another based on roads in the other two packages, with a discontinuity at some higher threshold.

Incorporating Additional GoA Roads. If the Government of Armenia funds some of the roads from the second and third MCA packages in their 2009 budget, these roads could still potentially be included in the analysis if the GoA follows a strict ERR rule for selecting road projects to fund. **As before, an RD design would only be feasible if the roads with the highest ERRs were the ones the GoA chooses to fund.** The analytic approach that incorporates later-package roads funded by the GoA would be the same as above. Strictly speaking, the interpretation of the impact estimates would differ slightly; instead of estimating the impact of Compact-funded roads, we would be estimating the impact of roads that were *originally planned* by MCA. This also applies to the Package 1 roads co-financed by the GoA.

Communities Served by Two Roads. The main complication that could arise if a single community is served by more than one road is that one road could be funded (making it a treatment group road) while another road serving the same community is not (making it a comparison group road). In these cases, we will base the analysis on the *higher* ERR of the two roads serving that community, as it is the higher ERR that determines whether the community is served by any treatment group road. This does slightly change the interpretation of the impact estimates, however, as we would be estimating the impact of being served by *any* project road that was rehabilitated, rather than just one.

Another complication is that a community served by a potential MCA project road may also be served by another road that is rehabilitated by another source. This would not invalidate our analysis, but would change the interpretation of the impact estimates. In this case, some of the communities served by comparison group roads might also be served by an outside-funded road, but the same would be true of the treatment group roads as well. Therefore, the impact estimates would tell us the impact of MCA project roads *relative to any other activities.* As an additional sensitivity check, we can exclude communities served by outside-funded roads (for example) to check whether they dilute our estimated impacts. However, we must be careful that excluding these roads does not leave a comparison group that is systematically different.

III. COMPARISON GROUP DESIGN

If an RD design is not possible, we will explore using a comparison group design instead. With a comparison group design, impacts would be estimated by comparing treatment group communities, where roads were rehabilitated, to a comparison group of communities where roads were not rehabilitated. Unlike an RD design, however, the division of communities into treatment and comparison groups does not follow a clear, strict decision rule (such as an ERR threshold). Therefore, it is an inferior option to an RD design because we cannot conclusively disentangle the impacts of the RRRP from the effects of other unobserved community characteristics.

Nevertheless, a careful comparison group design can still be a powerful tool for estimating impacts if RD were not possible. RD could be rendered infeasible for several possible reasons:

- Only roads in the first package are funded, and none (or very few) roads from the second and third packages;
- The ERRs are not strictly followed for later packages (whether funded by MCA or GoA); or
- Other donors fund roads that did not meet ERR thresholds.

The composition of the treatment and comparison communities would depend on the exact circumstances. The treatment communities would most likely include roads from the original MCA packages, whether they were ultimately funded by MCA or another source. The comparison roads would likely be those that were included in the original MCA packages but were not eventually rehabilitated by any donor.

We could conceivably supplement the treatment group communities in the analysis with roads that were rehabilitated by other sources and were never in the MCA packages. For example, the GoA and the Asian Development Bank have both funded other road projects. However, our data come from the Integrated Survey of Living Standards (ISLS), so the only other GoA or ADB roads would be those that were selected into the sample by chance. In the 2008 ISLS, only eight of the communities with GoA-funded roads in 2008 were selected, and only three of the communities with ADB-funded roads in 2008. Moreover, we do not have pre-intervention data for the majority of those GoA and ADB roads—only one community from each was in both the 2007 and 2008 ISLS.

Because of the limited number of GoA and ADB roads in the sample, it may be necessary to revise the ISLS sampling approach for 2009 to ensure more of these communities are included. This decision will need to be made by the end of October at the latest so we can adjust the ISLS sampling designs in time for NSS to draw next year's sample. We would probably still lack pre-intervention data for these communities. In this case, we could consider supplementing the baseline data with communities that were selected by chance into earlier ISLS samples and were included in GoA-funded road projects later. This would still be a small number of communities, however—only seven communities with GoA-funded roads in 2008 were selected for the 2007 ISLS, for example.

IV. INTEGRATED SURVEY OF LIVING STANDARDS

Armenia's National Statistical Service is fielding the ISLS annually throughout the Compact. The ISLS is a household survey on a broad range of topical domains and will serve as the key data source for the RRRP impact evaluation. The core sample of the ISLS includes 768 enumeration areas, each containing 8 households for a total sample of approximately 6100 households. Additionally, MCA is funding an oversample of 216 more enumeration areas in rural communities, an increase of approximately 1700 households.

The oversample is exclusively dedicated to communities served by MCA project roads.[2] Some other communities are selected by chance into the core sample, and we expect that a total of approximately 2200 households served by MCA project roads will be in the sample each year. Communities served by 82 of the 85 project roads are represented in the sample. (The three other project roads had to be excluded to preserve the marz proportions in the ISLS oversample.)

Intermediate Outcomes. While most of the outcomes of primary interest to MCA and MCC are longer-term outcomes, such as economic improvements, these outcomes may not be immediately observable. Consequently, we will closely examine intermediate outcomes through which the road rehabilitation projects are intended to improve household income. We would expect an impact on households' income only if we observe that a substantial proportion of the targeted communities are actually experiencing improvements in their roads. Examining the intermediate outcomes also establishes the counterfactual—how the quality of and access to roads would have changed even in the absence of the road projects. Table 1 summarizes the key intermediate that can be examined using the ISLS data.

---

[2] The ISLS sampling approach only includes communities *directly* served by project roads; however, other nearby communities may benefit as well. Measuring the total impact of the RRRP would require summing the impacts on all affected communities. However, including some communities that are indirectly served by project roads would come at the cost of communities that are directly served. Because impacts are likely largest in communities directly served by project roads, the impact evaluation is geared toward maximizing the statistical precision for those communities. Hence, the interpretation of the impact analysis will be the average impact on households in communities that are directly served by project roads.

Table 1. Intermediate Outcome Measures

Intermediate Outcome Measures	Time Frame
Utilization of Local Roads. Whether (and how often) road transportation is typically used to purchase agricultural supplies; to sell agricultural produce; to access employment outside the community; or for other purposes	Typical Month
Perceived Quality of Roads. Both roads within the village and roads connecting the village to other communities, as well as the quality of local public transportation.	As of Survey Date
Availability of Transportation. What modes of transportation are available and commonly used, especially to access key community services.	As of Survey Date
Access to Social Infrastructure. How far (distance and time) are key services, including health facilities, schools, community centers, and markets.	As of Survey Date
Agricultural Challenges. We will pay particular attention to whether transportation or access to markets were prevalent barriers.	Last Agricultural Season

Final Outcomes. The ultimate goal of the RRRP is to increase household income in rural Armenia, and hence, these outcomes are an important focus of our analysis. The primary domains we will focus on are agricultural revenue, costs, and profits from agriculture and employment by household members. We can also use the average sale price of specific crops for other farmers in the village to monetize crops that are consumed by the household or bartered. Table 2 summarizes the key final outcomes that can be examined using the ISLS data.

Table 2. Final Outcome Measures

Final Outcome Measures	Time Frame
Agricultural Production**.** Total amount of specific crops grown; amount of crops grown per square meter; total value of all crops cultivated; total amount consumed.	Last Agricultural Season
Food Production**.** The total amount of specific foods produced, such as bread, cheese, and milk; total value of all food produced; total amount consumed.	Last Agricultural Season
Livestock. Number of cows, pigs, sheep, and other livestock owned as of the survey date; number bought and sold in last 12 months.	As of Survey Date / Last Year
Revenue from Agricultural or Food Production. Value of crops or food sold; total value of all crops or food (including those sold, bartered, or consumed).	Last Agricultural Season
Agricultural Costs. Expenditures on fertilizers, pesticides, irrigation water, hired labor, equipment, and taxes (individually and in total).	Last Agricultural Season
Profit from Agricultural and Food Production. Revenues minus costs—the income from agricultural activities and food production. This could be measured two ways: One would be the purely monetary profit, while the other would monetize the amount consumed by the household.	Last Agricultural Season
Income from Employment. Whether household head, spouse, and any grown children were employed; total earnings from employment.	Last Week
Income from Remittances.	Last Year
Household Income. Can sum agricultural profits, employment income, and remittances to estimate total income, which in turn can be used to estimate the poverty rate.	Last Year

Community Survey. MCC is funding a community survey that will complement the ISLS, beginning in 2009. The community survey will be fielded in the same rural villages included in the ISLS sample. The community survey will be based on interviews with community leaders, such as mayors and will include village-level information. The content for the survey instrument is currently being finalized, but it provide a different perspective on several topical areas, including perceived quality of roads, transportation utilization, and access to social infrastructure. It will also cover community information that individual households would not necessarily know, such as new infrastructure that has been built or programs funded by other donors—providing valuable contextual information.

V. ESTIMATING PROGRAM IMPACTS

In this section we describe the empirical models that will be used to estimate RRRP impacts. We focus on the RD models, but touch on how the models would be specified if we were to instead use a comparison group approach.

RD Specification. The survey data will be cross-sectional, with a new cross-section of respondents drawn each year. Given this data structure, our econometric specification is designed to compare how villages served by treatment group roads changed over time to how villages served by roads in the comparison group changed over time. In its simplest form, impacts under the RD design can be estimated mathematically using the following regression model:

$$y_{irt} = \beta' x_{ir} + \phi T_r + \gamma ERR_r + \lambda T_r \times F_t + \delta ERR_r \times F_t + \theta F_t + \eta_r + \varepsilon_{irt}$$

where y_{irt} is the outcome of interest for household i served by road r at time t; x_{ir} is a vector of time-invariant characteristics of household i served by road r; F_t is an indicator variable equal to 0 in the baseline year and 1 in the follow-up year; T_r is an indicator equal to one if road r is in the treatment group and zero if it is in the comparison group; ERR_r is the estimated economic rate of return that divides roads into those that are funded (treatment) or unfunded (comparison); η_r is a road-specific error term (a road "random effect"); ε_{irt} is a random error term for household i served by road r observed at time t; and β, ϕ, γ, λ, δ, and θ are parameters to be estimated.

In this formulation, the estimate of the parameter λ is the regression-adjusted impact estimate, and represents the difference between the intercepts of the fitted regression lines for the treatment and comparison groups in the follow-up year. Standard t-tests can be used to gauge the statistical significance of the impact estimates.

The vector of baseline characteristics x_{ir} will include both household and village characteristics. At a minimum, we will control for village characteristics such as the geographic region. We will also control for household size and composition, and characteristics of the household head, namely, education level, gender, and age. In the framework of a repeated cross-sectional model, however, the characteristics that are included must be restricted to those that are unaffected by the road rehabilitation project.

The model above is designed to answer the general research question, "How have villages in the treatment group changed from the baseline year to the follow-up year, relative to villages in the comparison group?" This core model can be tweaked in a variety of ways to explore alternative specifications. One critical specification check will be to examine whether the impact estimates are robust when the sample is limited to the roads with ERRs close to the threshold. Although the sample size would be too small to use this as the core specification, the RD design hinges on sharp changes for observation just above and just below the discontinuity, warranting an investigation into whether any impacts observed in the full model are truly concentrated near the point of the discontinuity.

We will also explore alternative sample weights which could be modified to either give each road equal weight, or weights equal to the populations of the villages served by those roads. Such explorations would not change the general interpretation of the impact estimate, but they can provide insights on how robust the impact estimates are to these alternative specifications.

Comparison Group Specification. The general structure of the specification will be the same if we must resort to a comparison group specification instead of an RD approach; however, the comparison group design would not require the interaction term for *ERR* and F_t. (We could probably include this term if the ERR threshold is followed for most roads, but not all; this would constitute a so-called "fuzzy" RD design.) The main difference between the two is simply that the RD approach more persuasively identifies program impacts.

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

Precision of the Impact Estimates. A limitation of the impact analysis is that our estimates will not be very precise statistically. Even with optimistic assumptions about the roll-out, with at least half of the projects funded, we estimate that the minimum detectable impact on poverty is 11 percentage points.[3] Two main factors contribute to this imprecision. First, the clustered nature of the intervention described above reduces statistical power relative to an unclustered sample design. Instead of having 2200 households drawn independently from many different communities, they are drawn from just 82 road projects. Second, because of the substantial correlation between *T* and the *ERR* in the model above, impact estimates are less precise under a nonexperimental RD design than under an otherwise comparable random assignment design.

[3] These calculations assume a two-tailed t-test with 80 percent power and a 5 percent significance level; 41 treatment roads, 41 comparison roads, and 2200 total households responding; an intraclass correlation of 0.062; and an RD design effect of 3.

This design effect is typically between 2 and 4; in other words, the sample under an RD design must be two to four times larger than under a random assignment (RA) design to yield impact estimates with the same levels of statistical precision.

Tables. The report will include a variety of tables and figures with descriptive statistics on the data; however, the focus will be on estimates of RRRP impacts and their statistical significance. We will also report regression-adjusted means for the treatment and comparisons villages—that is, the means for the two groups if they had identical village and household characteristics. Table 3 provides an example of the structure of these tables.

Table 3. Example of Impact Estimates Table

	Treatment Mean	*Comparison Mean*	*Program Impact*	*p-Value of Impact*
Household Income				
Poverty Rate				
Etc.				

Notes: */**/*** indicate statistical significance at the .10/.05/.01 level.

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

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

$$y_{irt} = \beta' x_{ir} + \phi T_r + \gamma ERR_r + \lambda_{S=1} T_r \times F_t \times (S_{ir} = 1) + \lambda_{S=0} T_r \times F_t \times (S_{ir} = 0) + \delta ERR_r \times F_t + \theta F_t + \eta_r + \varepsilon_{irt}$$

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

Distributional Effects. The implicit focus of the analysis plan outlined above is on examining differences on the mean household. In conducting the analysis, it is also important to examine whether the interventions' effects vary at different levels of the outcome distribution. For example, the impact on income for households with very low or very high income may differ from the impact on households at the mean. Specifically, road projects may benefit higher-income households more if, for example, higher income households are more likely to own vehicles that could utilize improved roads. Conversely, the poor in the community might benefit more than the wealthy if public transportation is greatly improved, or if the road projects increase demand for manual labor.

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

Estimating impacts for specified quantiles starts with the same regression model as our core specification. The difference is in the methodology for estimating the parameters, which in turn, affects the interpretation of those impact estimates. While a standard regression model compares the impact for the average household, a quantile regression instead compares the impact of the interventions for a specified percentile, such as the 25th or the 75th percentile. For example, looking at quantile regressions at the 25th percentile can inform us about how the road projects have affected poorer households. Quantile regressions at the 50th percentile, the median, are also more robust to the influence of extreme outliers in the data, and thus can serve to validate the findings from standard regression analysis.

VI. PLANS FOR REPORTING

The final impact evaluation is planned for 2010, based primarily on data from the 2009 ISLS. However, depending on when the road projects are actually implemented, it may be useful to push the final impact evaluation back one year to permit a longer follow-up period that uses the 2010 ISLS, since many of the outcomes are unlikely to be immediately observable. If we adopt this approach, a tentative timeline (pending plans for implementing road packages) would be:

2007: Baseline ISLS data collection. (The baseline may extend through 2008.)
2008 Q3-Q4: Construction begins for Package 1 roads
2009: Possible construction for Packages 2 and 3; ISLS data collection
2010: Final ISLS data collection
2011 Q2: Data processing completed
2011 Q4: Impact evaluation report completed

We could also consider an interim report focusing on impacts for intermediate outcomes, such as the quality of roads, in the year after the roads are rehabilitated. We have not budgeted for this but would be open to exploring this option if it is of interest to MCC and MCA. The value of this will depend largely on when the road packages are actually implemented; if most road packages are not completed until sometime in 2009, then 2010 would be the first year of post-intervention data. In this case, an interim impact evaluation would not be beneficial, and the final impact evaluation would concentrate more on the intermediate outcomes than the final outcomes, for which impacts are unlikely to have been realized so soon after implementation. We could also consider pushing the impact evaluation back another year (into 2012, using the 2011 ISLS) to permit a longer follow-up.

cc: Sonya Vartivarian; Dan Player; Randall Blair