

SEPTEMBER 16, 2016

LAPOP'S RESPONSE TO DAVID ROSNICK, ALEXANDER MAIN, AND LAURA JUNG

"HAVE US-FUNDED CARSI PROGRAMS REDUCED CRIME AND VIOLENCE IN CENTRAL AMERICA? AN EXAMINATION OF LAPOP'S IMPACT ASSESSMENT OF US VIOLENCE PREVENTION PROGRAMS IN CENTRAL AMERICA." WASHINGTON, D. C.: CENTER FOR ECONOMIC AND POLICY RESEARCH, SEPTEMBER, 2016.

LATIN AMERICAN PUBLIC OPINION PROJECT (LAPOP)

VANDERBILT UNIVERSITY

Nashville, TN



LAPOP

The funds to carry out the impact evaluation study came from the United States Agency for International Development (USAID); however, this report was produced by LAPOP. The opinions expressed in this report are those of the authors and do not necessarily reflect the point of view of USAID.

TABLE OF CONTENTS

Executive Summary	1
Background.....	1
Initial Flaws in the RMJ Critique	2
Overstatements in the CEPR Report.....	3
The Alleged “Major Problem” with the LAPOP study.....	4
Flaws in the RMJ Methodology	6
Conclusion	12

TABLE OF FIGURES

Figure 1: Perceptions of Robberies in Treatment and Control Areas (Model Predictions)	9
Figure 2: Perceptions of Robberies in Treatment and Control Areas (Model Predictions) with RMJ Counterfactual	10
Figure 3: Perceptions of Insecurity in Treatment and Control Areas (Model Predictions).....	12

EXECUTIVE SUMMARY

A report issued earlier this month from the Center for Economic and Policy Research (CEPR) is critical of a five-year long, multi-country, multi-method impact evaluation carried out by Vanderbilt University's Latin American Public Opinion Project (LAPOP) of a USAID violence and crime prevention project in Central America (related to the broader CARSJ program). LAPOP responds to this critique by demonstrating that the CEPR report's claim of having uncovered a central problem with the research is based on a misreading of the nature of the research design, a misguided decision to analyze the data in ways that are not supported by the study design, and an unjustified decision to overlook secondary evidence provided in the broader study. Further, the statistical approach used in the critique is flawed; it departs not only from the original study design but also from the solid basis of the gold-standard approach used by the LAPOP study.

BACKGROUND

The Latin American Public Opinion Project (LAPOP), based at Vanderbilt University, is a research institute that unites scholars, students, universities, and think tanks across the Americas. LAPOP regularly carries out surveys of public opinion on democracy and governance via its AmericasBarometer survey, which covers 34 countries and includes interviews with more than 225,000 individual respondents. Its data and related studies and reports are available for free download and online analysis at its [website](#).

In addition to the AmericasBarometer, LAPOP also conducts a number of specialized studies on particular topics, based on the interests of its research and donor communities. Its multi-year, multi-country, multi-method (quantitative and qualitative) [Impact Evaluation of USAID's Community-Based Crime and Violence Prevention Approach in Central America: Regional Report for El Salvador, Guatemala, Honduras and Panama](#) was published in 2014, and disseminated at national and international meetings and events.

As a group of dedicated scholars, LAPOP is always pleased when our research stimulates interest and follow-up analysis. The recent Center for Economic and Policy Research (CEPR) [report](#) by David Rosnick, Alexander Main, and Laura Jung, hereafter cited as RMJ, is one such effort. While we do not know if the authors' contention that "[t]o date, only one in-depth assessment of a CARSJ program has been published" is correct, we encourage more studies and analysis of the impact of public policies; ordinary citizens and government officials alike need to know whether public funds are achieving their objectives.

Since February 2015, LAPOP has responded to well over a dozen queries from David Rosnick, the lead author of the RMJ report. Over that year and a half, LAPOP provided him with details of our study and sent computer code beyond that which is posted on our website to facilitate his replication of our work. Since replication lies at the heart of scientific validity and reliability, we were pleased to read in Rosnick's email, dated July 29 this year, that,

This has been very helpful. I've been able to replicate the with-ids results reasonably well based on the available data. So thank you for that. I'm hoping that soon I can send you a document outlining what I see.

In other words, his analysis and our analysis came up with essentially the same results. Based on his email, we expected to engage in further dialogue with Rosnick and his team.

We are disappointed, therefore, that Rosnick and his team shortly thereafter published their critique without giving us the opportunity to comment on their work. Such a dialogue would have provided us an opportunity to discuss, for example, how our study design and method of analysis attended to standard threats to internal validity.

The unfortunate result is that RMJ contains a number of errors that, if corrected, would have substantially altered its conclusions. Most importantly, a significant component of the report's critique rests on what RMJ claim is a "main problem" with our study (see RMJ pp. 4), when in fact that putative "problem" only emerges as a result of their misreading and mischaracterization of a central element of our study.

INITIAL FLAWS IN THE RMJ CRITIQUE

RMJ misrepresent our research, revealing a systematic pattern of errors and omissions. For example, they state (p. 3), "The study, conducted in 2013...", when in fact the research effort took place over a five-year period, 2010-2014. This error is not trivial, for two reasons. First, our study was longitudinal, aiming to examine impact over a range of years rather than the typical one-shot evaluations that dominate the field. We were looking for impact *over time* as the central driving element of our research design, and as consistent with the pre-test, post-test "true experimental design" heralded as the strongest approach to experiment-based research.¹ It is hard to imagine how RMJ could have missed or discounted this design feature. The very first table (Table 1, p. 27) of our study lists the month and year for every round of the fieldwork in the 2010-2014 period, for each country. In addition, every graph showing impact includes three points in time: pre-treatment (round 1), mid-term (round 2) and final (round 3).

Further, while RMJ do note that we based our conclusions on a study of 127 communities in four countries, they fail to note that the conclusions were based on an unusually large sample involving 29,000 respondents. They also make no note of the multi-method nature of our evaluation, ignoring the fact that we complement our quantitative research with qualitative research involving 848 qualitative stakeholder interviews and 44 focus groups. While the summary report, of necessity, had to be brief, and therefore included only a small slice of the qualitative evidence, the far more detailed country reports, available on the same web page as the summary report, and thus easily accessible to RMJ, contain extensive corroborating qualitative information (see <http://www.vanderbilt.edu/lapop/carsi-study.php>). Those reports are between 200 and 300 pages each. The mixed method approach our study employed and the evidence from each component is important, given the consensus in the modern academic community that multi-methodology, sometimes referred to as "triangulation", nearly always results in a more robust assessment than relying on a single approach.² Interestingly, other reports are reaching conclusions

¹ Campbell, Donald T., and Julian C. Stanley. 1963. *Experimental and Quasi-Experimental Designs for Research*. Boston, MA: Houghton Mifflin.

² See, e.g., John W. Creswell and Vicki L. Plano Clark. 2011. *Designing and Conducting Mixed Methods Research* (2nd Ed.). New York: Sage Manuscripts.

similar to those we presented in our study; see, for example, last month's *New York Times* story that details a number of outcomes and reports that comport with our conclusions.³ In short, RMJ omits reference to our broader study design and the vast body of evidence we brought to bear on the neighborhood model of crime and violence prevention; their narrow assessment of the research question runs counter to current best practices.

A further problem with the RMJ study is its lack of transparency. We pride ourselves on the transparency of our research and posted online the study's datasets and corresponding computer code for anyone to download and analyze, which is precisely what RMJ did. In contrast, when RMJ published their report they provided no link to their analysis and replication code for us and others to examine. Without this information, their claims cannot be fully verified.

OVERSTATEMENTS IN THE CEPR REPORT

RMJ also fail to note in their executive summary, conclusions, or press release the mitigating language used in their more detailed analysis. For example, in reference to their own analyses they state (p. 9):

Though this analysis cannot rule out the possibility that there is no effect from intervention, the sample size has been reduced greatly from the thousands surveyed. The test may simply lack the power to detect a small effect.

In other words, using their approach, which reduced the sample size dramatically, they are unable to conclude that there is no effect (positive or negative) of the CARS community-based crime and violence prevention approach. We are troubled that RMJ do not highlight these limitations to their approach in their summary conclusions.

Further, on September 9, some days after RMJ published their study, the following wording appeared on the CEPR website at the end of their original statement:

This report was corrected on September 9 to note that the authors find that the study, rather than necessarily the data, cannot support the conclusion that the areas subject to treatment in the CARS programs showed better results than those areas that were not.

(<http://cepr.net/publications/reports/have-us-funded-carsi-programs-reduced-crime-and-violence>)

Similar qualifying language was buried in the accompanying CEPR press release:

The paper notes that in some treatment areas, "Statistically, the possibility that intervention had no effect on reported robberies cannot be ruled out."

³ Sonia Nazario, "How the Most Dangerous Place on Earth Got Safer: Programs funded by the United States are helping transform Honduras. Who says American power is dead?" *New York Times*, Sunday, August 11, 2016. <http://www.nytimes.com/2016/08/14/opinion/sunday/how-the-most-dangerous-place-on-earth-got-a-little-bit-safer.html> Accessed September 13, 2016.

In other words, rather than critique our overall conclusions, CEPR’s report delimits sharply its broader, headline-grabbing statement, “Study Doesn’t Show that Areas Subject to Treatment in CARSI Programs Have Better Results,” to a comment on a subset of treatment areas and, principally, one particular outcome indicator.

LAPOP’s study was never designed to examine the impact of CARSI⁴ on a community-by-community basis. It would have been pointless to even attempt to do so given that we were working in over 100 communities spread among 4 countries. We would naturally expect variation at such a micro-level. Rather, the entire LAPOP effort was directed using multiple methods and indicators in order to answer the overriding, policy-relevant question, which was whether or not the CARSI approach to crime and violence reduction showed evidence of effectiveness. It did.

Perhaps most telling about the weakness of the RMJ critique is that their conclusion makes only the weakest of claims. They do not refute our evidence, but merely claim that they cannot “rule out that the intervention had no effect.” Given the low power that their statistical analyses have, as they themselves note in the report, they also cannot confidently rule out that it did have an effect.⁵ When the data are analyzed in ways consistent with the original design (as we do in our study, and as they did in their replication of our original work), the statistical power is higher and significant effects are detected.

AN ALLEGED “MAJOR PROBLEM” WITH THE LAPOP STUDY

The CEPR report repeatedly points to the “nonrandom” nature of the selection of treatment vs. control communities. Beginning on page 1, the authors state that “[t]his report identifies major problems with the LAPOP study, namely, the nonrandomness of the selection of treatment versus control areas...”. That point is repeated several times, with the authors asserting that “[t]he main problem in the LAPOP study is the nonrandomness of the selection of treatment and control areas” (p. 4). If that is alone or jointly the main problem, then the critique’s main issue with our research evaporates, since in fact the treatment and control communities *were indeed selected at random*, as our report stated clearly and repeatedly (e.g., see pages 23, 24, 59, and the appendix to the main report). Our researchers spent months gathering extant data to define and characterize each location, and then did what very few USAID evaluations have ever done: Vanderbilt *randomly* selected treatment and control communities, and told USAID where their community-based interventions should occur (treatment) and where they should not (control).

Random selection of treatment and control communities is a critical feature of experimental design, because it wards against eight principal threats to internal validity (that is, the ability to infer that the

⁴ Note that, as a convenient shorthand, we will refer to the bundle of projects (interventions) that we assessed in our study as the CARSI approach, intervention, or program; in practice, CARSI’s scope is far broader.

⁵ It may be that preconceived notions on CARSI motivated RMJ’s approach. It concerns us that a member of the team presented forceful statements against U.S. policy in Central America and CARSI *prior* to their publication of their analysis of our data. For example, in a 2015 [paper](#) published in *NACLA*, Alexander Main, one of the three authors, critiques in the broadest terms not only CARSI, but U.S. efforts since the 1980s to aid political and economic development in Central America, and does so without citing supporting evidence.

treatment caused the observed effect). One of these threats is regression to the mean, a key focus of RMJ report. As Campbell and Stanley's (1963) classic work on experimental methods explains, regression to the mean is a significant problem when groups are "selected for their extremity," which is exactly why our design opted for random assignment. As Campbell and Stanley state, selecting treatment and control groups randomly means that non-treatment sources of variance (such as regression effects) "have been allowed to affect the individual scores in both directions" (and for both the control and treatment conditions) and thus threats to inference from regression effects are mitigated even when "a group *selected for independent reasons* turns out to have an extreme mean."⁶

The municipalities from which the treated and control communities were randomly selected, as we explain in detail in our study, were ones chosen for a very specific reason: high crime. It would have been inappropriate to draw a set of municipalities from the thousands, in total, that comprise these four countries. El Salvador alone has 263, many of those quite rural with relatively low levels of crime. USAID and the host governments identified for us a list of municipalities that had serious crime problems and were ones in which they could justify spending scarce U.S. and host country resources to address crime and violence problems. We verified this selection at the outset of our study by carrying out an entirely separate study of non-at-risk areas, which affirmed the assumption that the preselected communities had characteristics consistent with their "at-risk" designation. The implementation of our design went smoothly in El Salvador and Guatemala. But, as we note in the report, it was imperfect in Honduras and Panama.

In the case of Honduras, where our study was delayed by the 2009 coup, USAID had already selected the treatment communities by the time we were ready to begin, so random selection was no longer possible. However, we used the *propensity score matching* technique to select control communities that were as closely matched to the treatment communities as we could make them. Propensity score matching is a well-accepted approach to approximate a true experimental design.⁷ We faced a different problem in Panama, where, after we had begun our study, which randomly assigned treatment and control, USAID began the process of closing its mission to that country and halting work there. As a result, some of the treatment communities we had selected at random did not get treated, leaving us with too few treated communities for the country as a whole to draw proper inferences. While our central findings refer to the region as a whole, in order to check on the possible impact of these deviations from our original results, at one point in our published analysis, we report the data by country, so that one can compare Guatemala and El Salvador to Honduras, and collectively with and without Panama. As the study report notes, we found similar patterns of impact across the region.

What the RMJ report fails to recognize in its discussion is that once a set of at-risk municipalities had been identified non-randomly, the selection of treatment and control communities *within* those municipalities was indeed conducted randomly. In other words, it seems RMJ confused "municipality" with "community." Communities were small neighborhoods within municipalities. To be clear, it would not have made any sense to implement or study crime prevention and mitigation in municipalities that did

⁶ Campbell and Stanley (1963), pp. 11-12. Italics are from the original.

⁷ Matching increases confidence that observed effects are due to the treatment, by creating two groups – treatment and control – that are as similar as possible on observable variables.

not have major crime problems. However, once those high-crime municipalities were identified, the *selection of treatment and control communities for this community-based approach to crime and violence was random*. Barring the exceptions noted above, the comparison that we made is between randomly selected treatment and randomly selected control communities in the selected high crime municipalities in Central America.

In sum, on the issue of this supposed “main problem,” RMJ are simply wrong. We did randomize our selection of communities in Guatemala and El Salvador and did a propensity score matching in Honduras. In Panama, where we also did random selection, the number of communities that were eventually treated was too small for us to report individual country results, so we provided regional results with Panama included and excluded. The RMJ “main problem” critique on this point is without substance.

FLAWS IN THE RMJ METHODOLOGY

Without explicitly stating so, the RMJ critique agrees with the central model used in the LAPOP analysis—the Difference-in-Difference model described in depth in our study (see the section, “How to Interpret our Results”, pp. 29-30). The model shows that the treated areas, in the region as a whole, improved after the CARSI intervention at a level not seen in the control areas. LAPOP and RMJ, however, diverge on two issues: 1) the level of analysis; and 2) the standard by which to judge the improvements in the treatment area.

The first issue is the level at which the data should be analyzed. RMJ chose to aggregate and analyze the results at the level of the municipality. There is no justification or foundation for doing this. The LAPOP study was specifically designed *not* to draw inferences at the sub-national level. The number of municipalities was simply too small to do so. Our study (as presented in the methodology and analyses in our report, and briefly described below) was undertaken primarily at the Central American *regional* level. In our initial research design presented to USAID, we argued that significant outcomes (if present) might only be detected at the level of the region, pooling all of the communities, municipalities and countries.⁸

Since there was intense interest in the country-level results, we reported them at that level as well, and indeed took note of the evidence of impact there as well, an indication of the robustness of the findings. We did not, however, report at the level of the individual municipalities nor communities because, based on the design of the study, inferences could not be made at the sub-national level; to do so would have required considerably more data collection (e.g., increasing the number of treated and control communities per municipality and/or increasing the number of municipalities from which communities could be drawn) and expanding the programs to additional treatment areas. Our statistical power analysis and sampling strategy were designed for inferences at the regional level. Subsetting the sample into municipal sub-units could be misleading because the small sample could produce “Type II errors” (i.e., failing to detect significant effects that actually exist in the data). We suspect this is part of the reason RMJ report impacts that are not statistically significant.

⁸ This is because we expected to be working with a community sample of only 100 (though we slightly exceeded that number as a result of adding Honduras, initially excluded because of the June 2009 coup).

Our concern with focusing the analysis on subnational units is that there would not be a large enough sample of *communities* (as opposed to individuals, of whom there were many) in each subnational unit for the “law of large numbers” that probabilistic models depend on to smooth out the differences in sampled groups. With only a small number of communities sampled in each municipality (fewer than 10, on average), it is not surprising that the pre-treatment averages for each group are different. Indeed, we state this on page 24 of our study: “[d]ifferences in the averages for any given variable in the starting level of the treatment and control communities were expected for this study, and in fact that is what was found in the study.”

Focusing on the robberies measure analyzed by RMJ, there is only a small difference of less than two percentage points in the raw averages between the treatment and control groups in the first year (43.6 and 41.7, respectively) *for the region as a whole*. The difference between the groups is larger if we use our statistical model, which controls for socio-economic variables (46.2 and 38.5). RMJ do not offer a standard of what they consider to be “elevated” or “abnormally” high levels of reported robberies. Yet, given the range of averages reported by RMJ at the municipality level, neither of these rates appears to be particularly high.⁹ Furthermore, best practices caution against the types of inferential errors that can seep into research when one analyzes a subset of the entire treatment, or control, group (see Campbell and Stanley 1963).

The second issue that distinguishes our study from the CEPR report relates to how one assesses improvements in the treatment area. Both studies attempt to construct a counterfactual scenario with which to compare the treatment area result. This unobservable state is the level of robberies – to take the one variable that the CEPR report examines in detail – that *would have* been observed *if* there had been *no* CARSI program in the treatment area. The counterfactual used by LAPOP was the one developed in the initial design of the study, prior to the collection of any data, and one that we therefore retained throughout the years of our work. It is based on the Difference-in-Difference estimation strategy. This approach provides a well-established standard, determined *a priori*; the method has a very long history of use in economics and other social sciences, as we point out in our country-level reports. In those reports we state:

A Difference-In-Difference estimator (DID) is a widely used econometric technique in the impact evaluation field, and while its use dates back over decades, since the work thirty years ago by Ashenfelter and Card (1985)¹⁰, it has become one of the “gold standards” in the field.

A significant advantage of DID is that it is able to control for what statisticians call “omitted variable bias.” DID is able to take into account changes in factors such as the national economy, employment and overall crime rates, which impact the control and treatment groups equally. In contrast, if one only observes the

⁹ Specifically, the values for municipalities in the control areas in the first year of the study range from a minimum of 15.7 to a maximum of 76.3. The mean of these values is 38. Both regional means of 46.2 and 38.5 fall between the 60th and 70th percentiles of the municipality means in the control areas. The distribution of municipality means is similar in the treatment areas in the first year. The mean is 43.1 with a minimum of 18.6 and a maximum of 76.3.


¹⁰ Orley Ashenfelter and David Card. 1985. “Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs,” *The Review of Economics and Statistics* Vol. 67, No. 4, pp. 648-660.

treatment group before and after treatment, then one would not know whether the treatment or these alternative factors were responsible for the observed changes in the treatment group. The DID approach estimates both the difference between the treatment and control groups and the difference within each group over time. The treatment is considered effective if outcomes in the treatment group improve over time significantly more (or decline less) than they did in the control group. One of the great advantages of this approach is that it does not require that the treatment and control groups have identical starting values, a basic fact that is not addressed by RMJ, since it is the *trend* in each group and the difference between those trends that is of interest.¹¹ We can estimate a counterfactual case from the model to simulate the outcome as if no treatment had been administered. The counterfactual is the change in the outcome measure of the control group over time, but beginning at the same level as the treatment group. The difference between the counterfactual and the outcome of the treatment group is considered the *treatment effect*: the impact that can be said to have been caused by the intervention.

A clear illustration of the counterfactual used by LAPOP is shown in each of the published reports. The published results on robberies from the LAPOP report (page 32, essentially reproduced by RMJ in the Appendix, Table 1A) is reproduced below in Figure 1. Here we see the predictions from the DID model. The solid green line shows the model predictions for the treatment area, the blue line shows the model predictions for the control area, and the dashed green line shows the counterfactual from the DID model. The counterfactual is what we would expect to have occurred if the treatment group had not received the treatment. We would expect to observe the same trend in the treatment area as in the control area, if there were no treatment. In the control area, we see an increase in the second year, followed by a decline in the final year. So, the DID approach tells us that if there had been no treatment we would expect to observe the same trend in the treatment area as in the control area. The difference between the counterfactual in the third year and the treatment in the third year (7.9 percentage points) is, according to the DID approach, considered the *treatment effect*—that is, the predicted effect of the treatment. This difference is statistically significant, and this supports a conclusion that the programs under evaluation had a positive effect on the communities in the region.

¹¹ One challenge to the DID is when only one of the two groups is affected by something other than the treatment. Since our study relies on many neighborhoods, spread across several municipalities, and in multiple countries, the chances of a systematic occurrence such as this are extremely low. Our extensive qualitative evidence and field report offer no evidence of any such confounding effects. Another challenge to the DID is when the composition of treatment or control groups changes systematically over time. We analyzed the data for this, and found that the composition of the groups did not change on key observable variables (education, wealth, age, etc.) over the course of the experiment. In short, our study design minimized challenges to causal inference and we carefully examined the data and found no evidence of threats to inference. As in all social science work, such threats cannot be ruled out entirely, with the greatest challenge being that of external validity. While we have strong evidence to conclude that the treatment was effective in the area we studied, we cannot guarantee that similar interventions will succeed elsewhere. Still, our results were remarkably consistent across several countries, suggesting that external validity may be less of a concern in our study than in a typical single-country study.

Figure 1: Perceptions of Robberies in Treatment and Control Areas (Model Predictions)




RMJ, in contrast, argue that the presumed improvements are instead a result of treatment groups, when aggregated to the level of the municipality, showing higher than normal rates of robberies because treatment and control groups were allegedly not randomly assigned. The improvements, they argue, are a return to the “normal” level of robberies rather than due to the CARSI treatment.

Setting aside their erroneous assertion of nonrandom assignment, this claim is based on a flawed approach. There is no clear standard of what is the “normal” level of robberies or any other of the outcome indicators we studied. In any social scientific study, we need to establish *a priori* a counterfactual grounded in scientific theory. The authors of the CEPR report do not establish an explicit counterfactual with which to compare outcomes in the treatment area. Our careful reading of their methodology suggests that they implicitly assume that the “normal” level of robberies is the average municipal level of robberies in control areas in the first time period (prior to the implementation of the CARSI intervention in the treatment areas). As we explain below, such a counterfactual lacks any basis in scientific rigor.

Figure 2 illustrates the counterfactual implicit in the CEPR report methodology (gray dashed line) in the context of the original results of the study of robberies: the predicted level of robberies for the control area prior to intervention.

Figure 2: Perceptions of Robberies in Treatment and Control Areas (Model Predictions) with RMJ Counterfactual



As Figure 2 shows, even if we use such an unconventional counterfactual, we still observe an effect of the CARSI treatment (a difference of 5.9 percentage points between the treatment in the third year and the mean of the control area prior to treatment). While this effect is somewhat smaller than the 7.9 percentage points we report, even using the RMJ standard, we still observe a statistically significant improvement in robberies that would be appropriately attributed to the treatment, the CARSI programs.

Still, applying this RMJ-inspired counterfactual violates two important foundations of experimental designs in general and the Difference-in-Difference approach more specifically. First, the RMJ approach is only established after looking at the data, rather than setting a standard *a priori* and applying it throughout the study, as we did. Moreover, this approach uses only the control data in the first year of

the study, ignoring trends in the control areas that are informative and help to avoid the problem of omitted variables.


One could assert an infinite number of possible “normal” levels after a study is conducted. After looking at the trends in the data, one can decide arbitrarily that some areas have exceeded the threshold for “normal.” But establishing such a standard after the fact risks choosing a standard designed to guarantee certain results. In contrast, LAPOP established a counterfactual in the initial study design, *prior* to the collection of data without making any assumptions about the outcome. Moreover, the LAPOP standard is based on the well-established Difference-in-Difference methodology used in “gold standard” impact evaluation studies worldwide.

Implicit in the CEPR report methodology is its reliance on the outcome measure for the control group in the first time period. Using only pre-intervention control area data to evaluate the treatment outcome in a longitudinal study goes against established experimental design. Doing so risks what statisticians call *omitted variable bias*. Without taking into account trends in the control group over the study period, we cannot distinguish treatment effects from the influences of national or regional factors such as economic changes or critical events that affect outcomes in both the control and treatment groups. The DID approach is designed to identify and account for such factors. An excellent illustration of this is the case of El Salvador. In El Salvador, a truce between the major gangs in March 2012 affected crime levels throughout the country. Had our analysis not taken into account the changes seen in the control area, as well as changes in the treatment area, our conclusions would have erroneously attributed all improvements to CARSI instead of identifying the effect of the gang truce. This is an illustration where the LAPOP approach clearly avoided overstating the impact of the CARSI program by potentially wrongly attributing it to the program rather than the gang truce.

In sum, the CEPR report methodology is seriously flawed. Had we followed its approach, we would have departed from the original study design and best practices in ways that lack justification.

While RMJ focus primarily on a single indicator of crime and security in their critique—robberies – it is important to note that several of the indicators reported in the LAPOP study have starting values in the treatment and control areas that are statistically indistinguishable, including perceptions of insecurity, sale of drugs, and interpersonal trust. Figure 3 shows the published graphic for the indicator: *citizens’ perceptions of insecurity*. For this measure, “good” values are low (greater security) and “bad” values are high (greater *insecurity*). It is clear that the starting points are nearly identical before any treatment, yet perceptions of insecurity diverge considerably in subsequent years in the control and treatment areas. The treatment is shown to have a small, but statistically significant effect. Focusing principally on a single indicator is consistent with the tendency in the RMJ report to overlook quantitative and qualitative analyses in the reports and datasets that comport with the general conclusion drawn in our report.

Figure 3: Perceptions of Insecurity in Treatment and Control Areas (Model Predictions)



CONCLUSION

We at LAPOP stand behind the study design, analysis, and conclusions published in our report, *Impact Evaluation of USAID's Community-Based Crime and Violence Prevention Approach in Central America: Regional Report for El Salvador, Guatemala, Honduras and Panama (2014)*. Our full research design and analyses, both quantitative and qualitative, are detailed in the study reports and accompanying material. While we welcome an open scholarly debate about our research and findings, we have serious concerns about the rigor and methodology of the report by Rosnick, Main, and Jung recently published by CEPR.