"Illiteracy" Revisited: What Ortega and Rodríguez Read in the Household Survey

David Rosnick and Mark Weisbrot

May 2008

Center for Economic and Policy Research
1611 Connecticut Avenue, NW, Suite 400
Washington, D.C. 20009
202-293-5380
www.cepr.net

Contents

About the Authors

David Rosnick is a Research Associate and Mark Weisbrot is Co-Director at the Center for Economic and Policy Research in Washington, DC.

Acknowledgements

The authors would like to thank Luis Sandoval for research and also Dean Baker, John Schmitt and Dan Beeton for helpful comments.

Introduction

In the summer of 2003, the Venezuelan government began *Misión Robinson*, a program aimed at adult literacy. By late 2005, the government claimed victory; saying 1.4 million had been taught to read. By comparison, the 2001 census reported only 1.1 million illiterate adults.

In "Freed from Illiteracy? A Closer Look at Venezuela's *Robinson* Literacy Campaign,"[1] Daniel Ortega and Francisco Rodríguez use the Venezuelan Households survey to gauge the impact of the effort on the adult national literacy rate. Ortega and Rodríguez conclude, "most of our estimates of program impact represent qualitatively small and rarely statistically significant effects of *Robinson*, while some point estimates are actually negative."

However, as everyone agrees that *Misión Robinson* existed, the question is not at all whether or not the effect is statistically significant. The real question is whether or not the program was effective, and on what scale. Ortega and Rodríguez's use of the Household Survey has serious problems and their statistical tests lack power and are not robust. Even overlooking these issues, it is only by ignoring reasonable assumptions about the program's reach that they can claim to show that the program was not a large-scale effort. Correspondingly, we make no claim in this paper that the program was in fact a large-scale program—only that the data and methods of Ortega and Rodríguez do not invalidate such a claim.

The Household Survey

Most importantly, the Household Survey can easily fail to capture significant gains in literacy. The survey simply asks whoever is present at the time of the interview about another family member: "Does this family member know how to read or write?" A respondent may answer "yes" to the question, knowing that her brother or father can write his name. At the end of a literacy program, however, the person may have considerably improved reading skills, but would still show up in the survey as the same response. So this survey is much too crude to measure the results of a program such as *Misión Robinson*; it was never designed to measure literacy or reading skills. To estimate the impact of such a program, one would need some before-and-after test for the participants. It would be possible to see some sign of the program's impact in the Household Survey, but the effect that is visible in this crude survey could be a small fraction of the people who acquired or improved their reading skills.

A semi-literate worker who could not read a newspaper article before the program but can do so afterwards has a significant improvement in skills, but may show up as a "yes" in the Households Survey in both instances. Such gains should not be discounted. A person whose income grows from $1.40 to $1.85 a day is still poor under a $2.00 per day standard, despite a not-insignificant rise.[2]

An improper estimate of the population

[1] Rodríguez, Francisco and Daniel Ortega. 2006. "Freed from Illiteracy? A Closer Look at Venezuela's *Robinson* Literacy Program." Middletown, CT: Wesleyan University, Department of Economics.

[2] Rather than the *poverty rate*—the percentage of people who earn less than the poverty line—economists studying poverty frequently look at the *poverty density*—the average percent of the poverty line by which families fall short.

It is possible that a significant majority of *Robinson* graduates were functionally but not fully illiterate to start. At first blush, it may seem strange that a literacy program could wind up teaching Venezuelans reported as literate and largely miss the million who report as illiterate. Yet, functional illiteracy in Venezuela is surely more widespread than illiteracy reported in the Household Survey. According to a 1994 survey, more than 20 percent of adults in the United States failed to meet minimal literacy standards. According to the survey report:

> The easiest task in level 1… directs respondents to look at a medicine label to determine the "maximum number of days you should take this medicine". The label contains only one reference to number of days and this information is located under the heading "DOSAGE". The reader must go to this part of the label and locate the phrase "not longer than 7 days".[3]

The adult literacy rate in the United States is reported as 99 percent, and yet 20 percent of those same adults cannot read the simplest instructions off a medicine label. In Venezuela, a million self-reported illiterate adults should imply millions more who are functionally illiterate. Even if a 2001 rate that included functional illiteracy in Venezuela were 25 to 30 percent of adults, then there would exist a pool of 3.9-4.6 million in need of training. If *Robinson* drew from this pool proportionately among reported and functional (non-reported) illiterates, then no more than 400,000 of a program that included 1.4 million participants would show up in the Household Survey.

TABLE 1:
Proportionate Participation by Literacy Level

Total Illiteracy Rate[4]	Participation rate of Reported Illiterates[5]	Participation rate of Functionally Illiterate[6]	Reported Illiterates Participating	Increase in Survey Literacy Rate
25%	36%	36%	393,120	2.5%
30	30	30	327,600	2.1
35	26	26	280,800	1.8

Source: Authors' analysis.

This is shown in Table 1. For example, the middle line assumes a total illiteracy rate of 30 percent (1.1 million people reported in the Household Survey plus 3.5 million functionally illiterate). Assume that 30 percent of each group participates in the program. This gives us 1.4 million participants, but only 327,600 participants would show up in the Household Survey as having learned to read. This could be even less if a higher proportion are drawn from the functionally illiterate population.

In fact, participants in the *Robinson* program were not selected randomly. Participation was entirely voluntary. With illiteracy concentrated in the highest age groups, many illiterate had little motive to learn to read. On the other hand, younger semi-literate adults might have a much greater incentive as well as have an easier time completing the program than would a fully illiterate individual. If *Robinson* attracted disproportionately the functionally illiterate, far fewer would appear in the Household Survey as illiterate before the program and literate after.

[3] Kirsch, Irwin. 2001. "The International Adult Literacy Survey (IALS): Understanding What Was Measured." Princeton, NJ: Educational Testing Service.

[4] This includes both the 1.1 reported illiterate and the functionally illiterate

[5] These are the 1.1 million people reported as illiterate in the Household Survey

[6] These are people who are reported as literate in the Household Survey but are functionally illiterate.

TABLE 2:
Disproportionate Participation by the Functionally Illiterate

Total Illiteracy Rate	Participation rate of Reported Illiterates	Participation rate of Functionally Illiterate	Reported Illiterates Participating	Increase in Survey Literacy Rate
25%	21%	42%	228,660	1.5%
30	17	34	185,500	1.2
35	14	29	156,050	1.0

Source: Authors' analysis.

Table 2 assumes that participation rates for the functionally illiterate were twice the rate of the reported illiterates. In this case, the number of participants that would show up in the survey could be as low as 156,000, which would yield only a 1.0 percent increase in the survey literacy rate.

These calculations assume that the government did not exaggerate participation in the literacy program. Suppose, instead, that *Robinson* was still a very large program, but reached only 700,000 adults, or alternatively that 1.4 million adults participated in the program, but only half of the 1.4 million actually learned to read. In such a case, we must halve the estimated effects on measured literacy in Tables 1 and 2. In other words, the Household Survey might show a literacy gain of less than 1 percent of the adult population and yet, still be consistent with a successful large-scale literacy program.

Just as participation rates may differ between the reported and functionally illiterate, so may success rates. Suppose participation rates were equal between the two groups as in Table 1, but that the functionally illiterate participants were more likely to gain full literacy than were reported illiterate. In Table 3, we show the observed increase in the official literacy rate assuming only one in three reported illiterates participating became literate, but two in three functionally illiterate participants did as well.

TABLE 3:
Disproportionate Success by Literacy Level

Total Illiteracy Rate	Reach rate[7] of Reported Illiterates	Reach rate of Functionally Illiterate	Reported Illiterates Succeeding	Increase in Survey Literacy Rate
25%	12%	24%	131,040	0.8%
30	10	20	109,200	0.7
35	9	17	93,600	0.6

Source: Authors' analysis.

In each of the scenarios we assumed 1.4 million participants, but in Table 3 we posited that between 800,000 and 840,000 adults learned to read as a result of *Misión Robinson.* The assumptions of Table 3 are entirely consistent with a very large-scale literacy program, yet we expect under these assumptions that the official literacy rate increased less than one percentage point—representing perhaps as few as 100,000 adults. Faulty as the models of Ortega and Rodríguez may be, they are not inconsistent with a large scale and successful program.

The *Misiones* themselves may bias the Household data

The existence of a literacy program itself could even eliminate some positive responses in the Household Survey. An illiterate mother once might have responded that her son is literate because he can sign his name. The same mother having learned to read through *Robinson*, or knowing

[7] Reach rate is defined here as the participation rate times the success rate.

someone else who has, may have developed a different view of what it means to be literate and cease responding affirmatively.

Of course, the psychology could cut the other way amidst a literacy campaign, as respondents feel pressure to say they can read and write. However, it still remains that the Household Survey can and will be expected to fail to capture changes in important areas both demographically and socially at which *Robinson* likely aimed.

Furthermore, Ortega and Rodríguez draw on the survey data to estimate the change in the number of illiterate. While a survey of this type may be able to provide an accurate estimate of the literacy *rate*, it is far less able to provide an accurate headcount of illiteracy.[8] Any estimate of headcount illiteracy in recent years needs to take into account the healthcare and anti-poverty programs carried out over the same time. These programs may have reduced mortality rates among the elderly, where measured illiteracy is very highly concentrated. If so, we would expect to find more illiterate persons (in terms of headcount) in Venezuela than will show up in the survey data.

The bottom line is that a Household Survey with just one question about whether a family member can read cannot be expected to capture most of the effects of a large-scale literacy program such as *Misión Robinson*. The data at hand is simply unsuited to that task.

The Ortega and Rodríguez Model for Literacy

Even discounting issues with use of the Household Survey, Ortega and Rodríguez fall short. For simplicity we will focus on the simplest regressions, in which Ortega and Rodríguez examine the evolution of the adult (aged 15 and over) literacy rate over time. For example, illiteracy is related strongly to birth year. In the second half of 2005, the literacy rate among 15-19 year-olds was over 98 percent compared to 81 percent among those 65-69. This is largely on account of schooling when each cohort was young. As time passes, older, less literate Venezuelans pass on, thereby increasing the overall literacy rate. We therefore expect the literacy rate to rise over time, but the speed with which the literacy rate rises should be decreasing. In simple terms, Ortega and Rodríguez wish to test if the literacy rate since the start of *Misión Robinson* is above the trend, which should represent the evolution of these various factors unrelated to *Robinson*.

[8] In the United States, for example, the establishment survey of employers, which is much closer to a census, measuring employees on payroll, is a more reliable measure of current employment levels than the household survey, which is used to measure employment rates.

FIGURE 1
The Pre-*Robinson* Literacy Rates

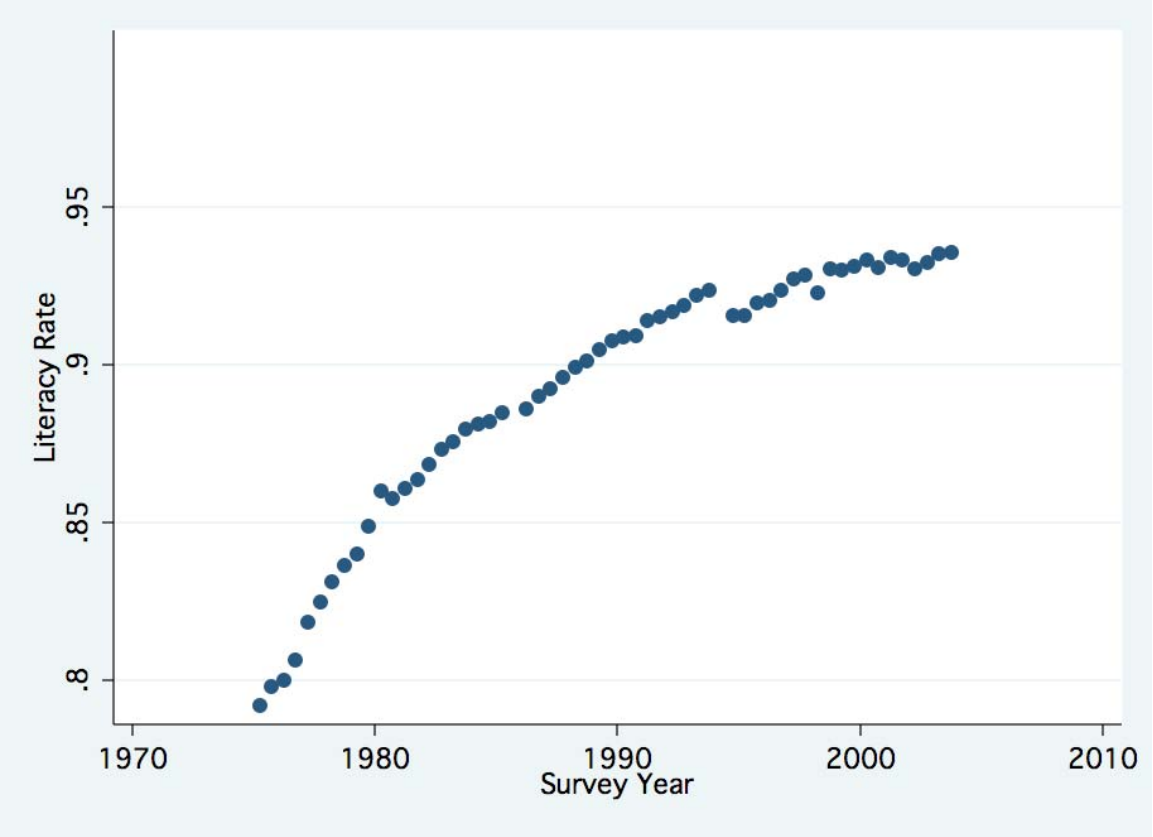

Source: Authors' analysis.

Figure 1 shows the pre-*Robinson* literacy data. There is a clear break in the data in 1994, and the speed of illiteracy reduction is definitely slowing over time. It appears that the improvements in literacy prior to *Robinson* have become so small that they are not easily detectable.

Figure 2 demonstrates more clearly the slowdown in progress leading up to the *Robinson* program. The figure shows the two-, three-, four-, and five-year average declines in the illiteracy rate. There are two gaps in the data—in 1985 and 1994—corresponding to missing survey results. Aside from a jump in 1985, the literacy gains become smaller over time and reach nearly zero just prior to the *Misión*. This is important information if we are to assess *Robinson*. If the trend in the literacy rate is slowing and very close to zero just prior to the program, then any observed gains in literacy during *Robinson* may be attributed to the *Misión*. If, on the other hand, we have reason to believe that gains in literacy are on an upward trend, then we can attribute very little of any literacy gains to *Robinson*. The data is ambiguous as to whether the literacy rate is increasing or decreasing just before the start of the *Misión*. However, the data do indicate that literacy gains are shrinking prior to *Robinson*.

FIGURE 2
Average Rates of Illiteracy Eradication over Time (Pre-*Robinson*)[/a]

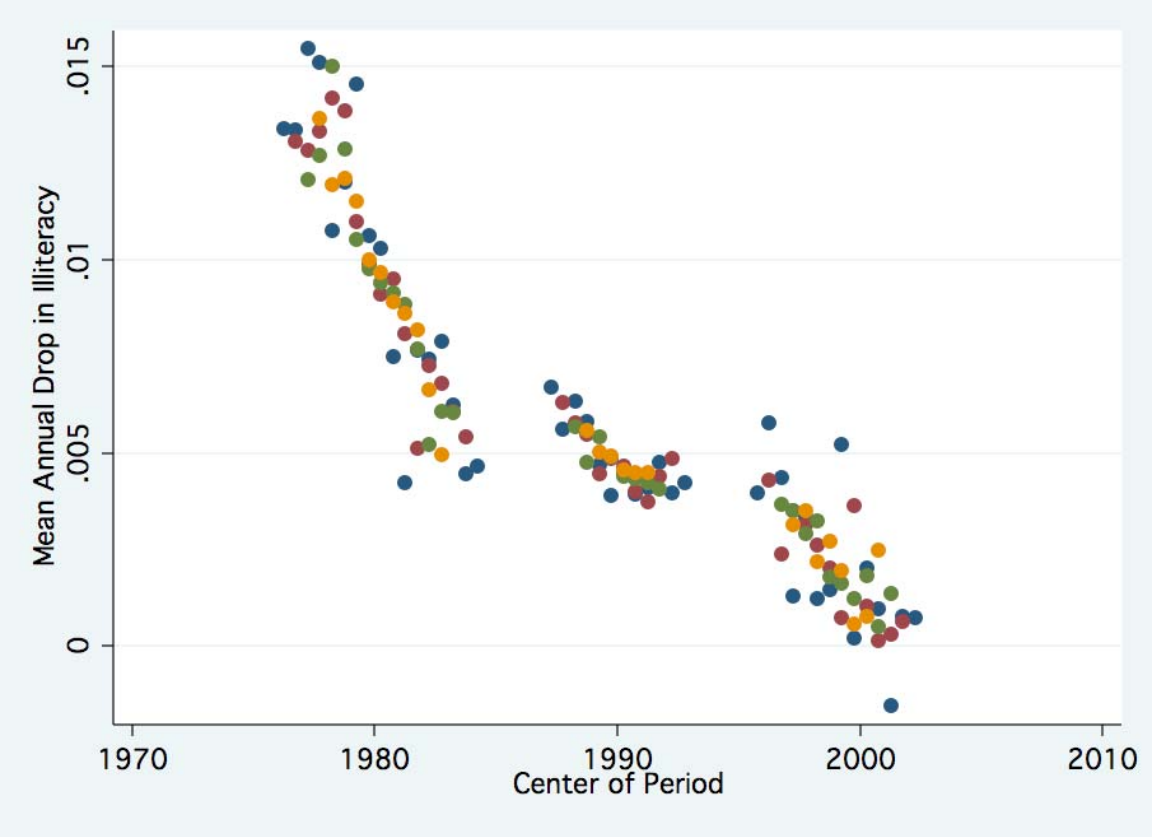

Source: Authors' analysis.
Notes:
a/ Two-year averages are shown in blue, three-year in red, four-year in green, and five-year averages are shown in orange.

Despite this obvious trend in the data, Ortega and Rodríguez at first assume the literacy rate at time t, L_t evolves as a cubic time trend. Ortega and Rodríguez do consider other time trends, but the vast majority of their reported regression results are based on their assumption of a cubic trend.

$$L_t = \alpha + \beta \cdot ROBINSON_t + \gamma_0 \cdot POST94 + \gamma_1 t + \gamma_2 t^2 + \gamma_3 t^3 + e_t$$
$$e_t = \rho e_{t-1} + \varepsilon_t$$

where $ROBINSON_t$, the variable of interest, takes on the value 1 starting in the second semester of 2003[9], and $POST94_t$ takes on the value 1 up until 1994[10] to account for a survey change at that time. Time is measured in semesters relative to the first semester of 2003. Finally, e_t represents a random error in each semester.[11]

Two concerns immediately come to mind. First, it is not obvious why *Robinson* would have a constant effect on literacy starting in the second half of 2003. Second, the existence of a cubic (or

[9] Note that this variable is lagged in certain models consistent with Ortega and Rodríguez, Table 3.

[10] This is *not* how the variable is described in the paper. The results reported in Ortega and Rodríguez, Table 2 used a pre-1994 indicator rather than a post-indicator despite the variable name and description in the paper. A pre-1994 dummy is consistent with the results and the description in Ortega and Rodríguez, Table 3.

[11] Ortega and Rodríguez's model is written out with errors $e_t = \rho \varepsilon_{t-1} + \varepsilon_t$; the discrepancy in the model description is simply an oversight in their paper.

indeed, any) polynomial time trend implies the literacy rate will at some time inevitably exceed 100 percent, or drop below 0 percent.

An underlying polynomial time trend

As discussed above, Ortega and Rodríguez use polynomial fits to the underlying time trend. Sooner or later, the rate will trend up past 100 percent or down below 0 percent. As it turns out, Ortega and Rodríguez's model implies that Venezuela's literacy rate should pass 100 percent by 2015, as in Figure 3 below. The trend and two-standard-error confidence bounds for literacy in the absence of *Misión Robinson* are shown. As can be seen from Figure 3, it would indeed be an impressive program that raised literacy rates faster than this trend.

FIGURE 3
Ortega and Rodríguez: Literacy Rate Fit using Cubic Time Trend

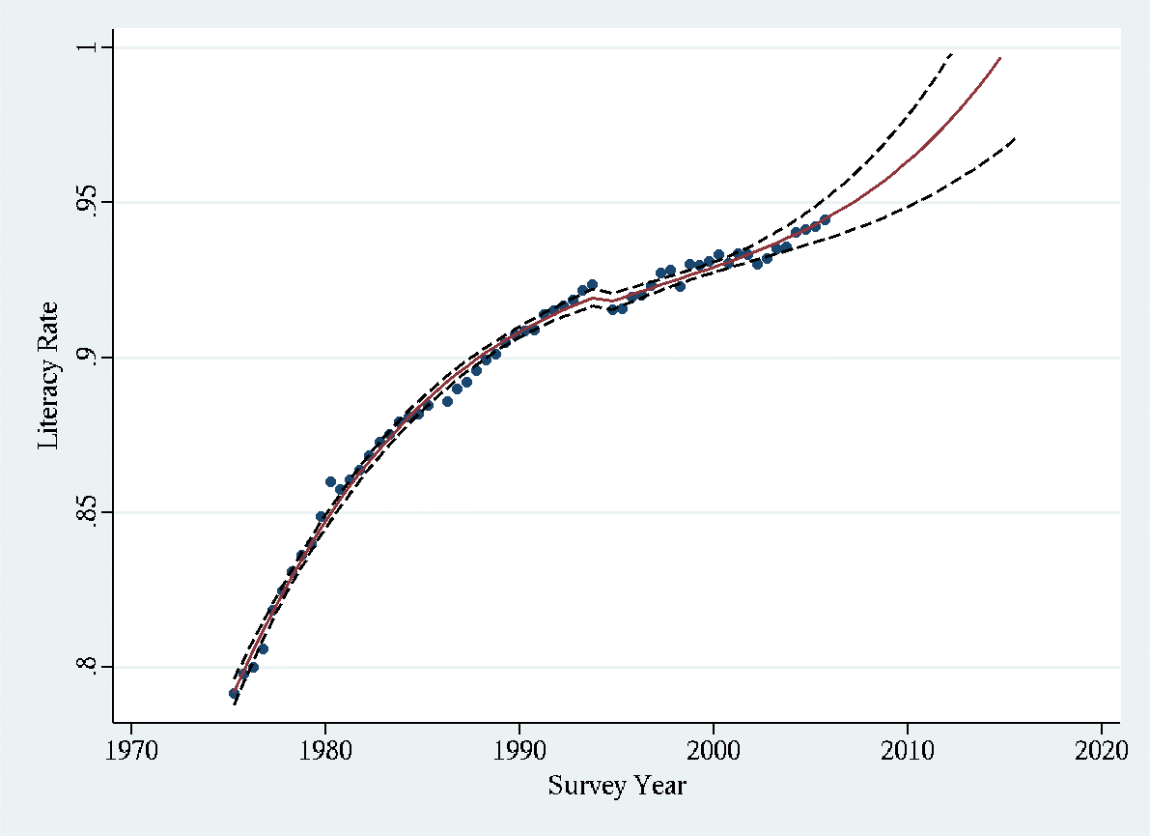

Source: Authors' analysis.

FIGURE 4
A Series of Polynomials Trends

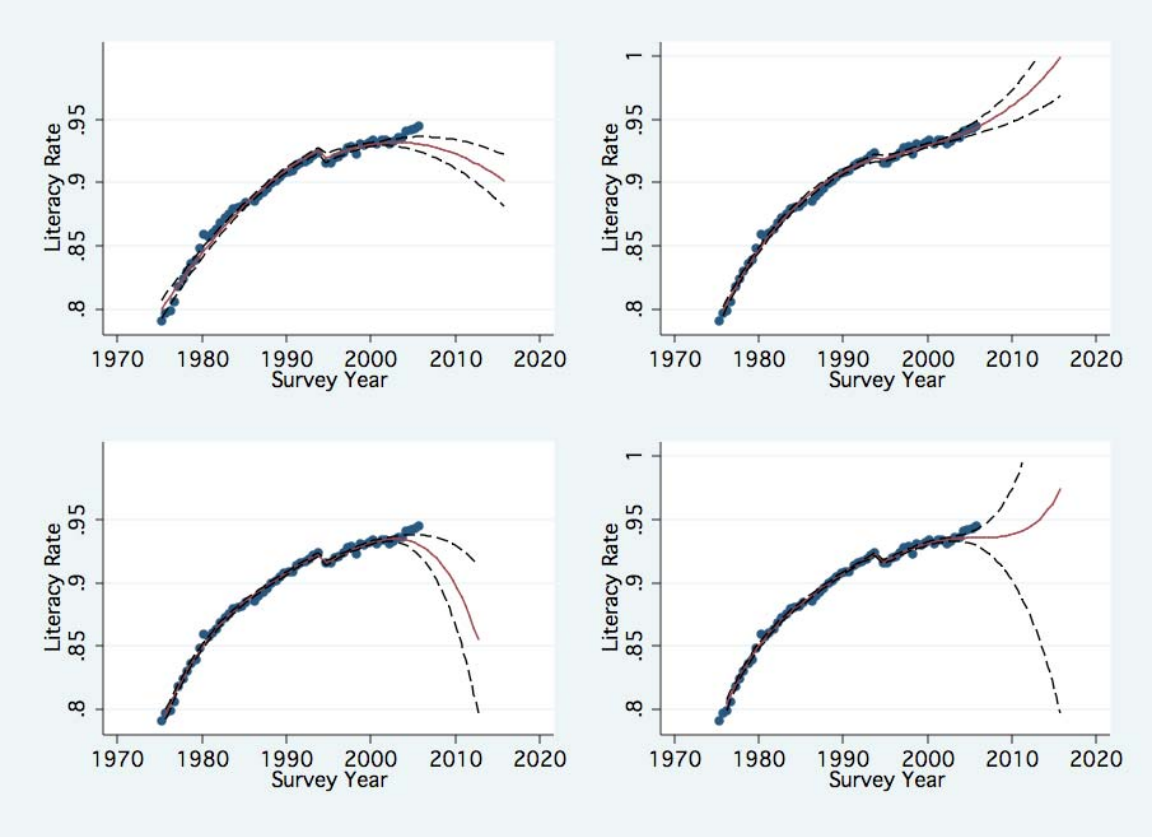

Source: Authors' analysis.

However, this rapid improvement in the trend is in no way robust. Figure 4 shows the quadratic, cubic, quartic and quintic polynomial trends.

Such a wide variety of polynomial trends that can be fit to the data serve as confirmation that the data will not help us evaluate *Robinson*. Indeed, these data are so statistically weak that it is not clear whether the literacy rate is rising or falling just prior to the start of the program. In addition, the cubic trend runs counter to what the data show. Figure 5 shows the rates of decline in illiteracy as predicted by Ortega and Rodríguez's cubic trend (as indicated by the solid blue line) in the absence of *Misión Robinson*. As in Figure 2, the observed average changes are shown as well.

FIGURE 5
Pre-*Robinson* Illiteracy Eradication—Predicted and Observed

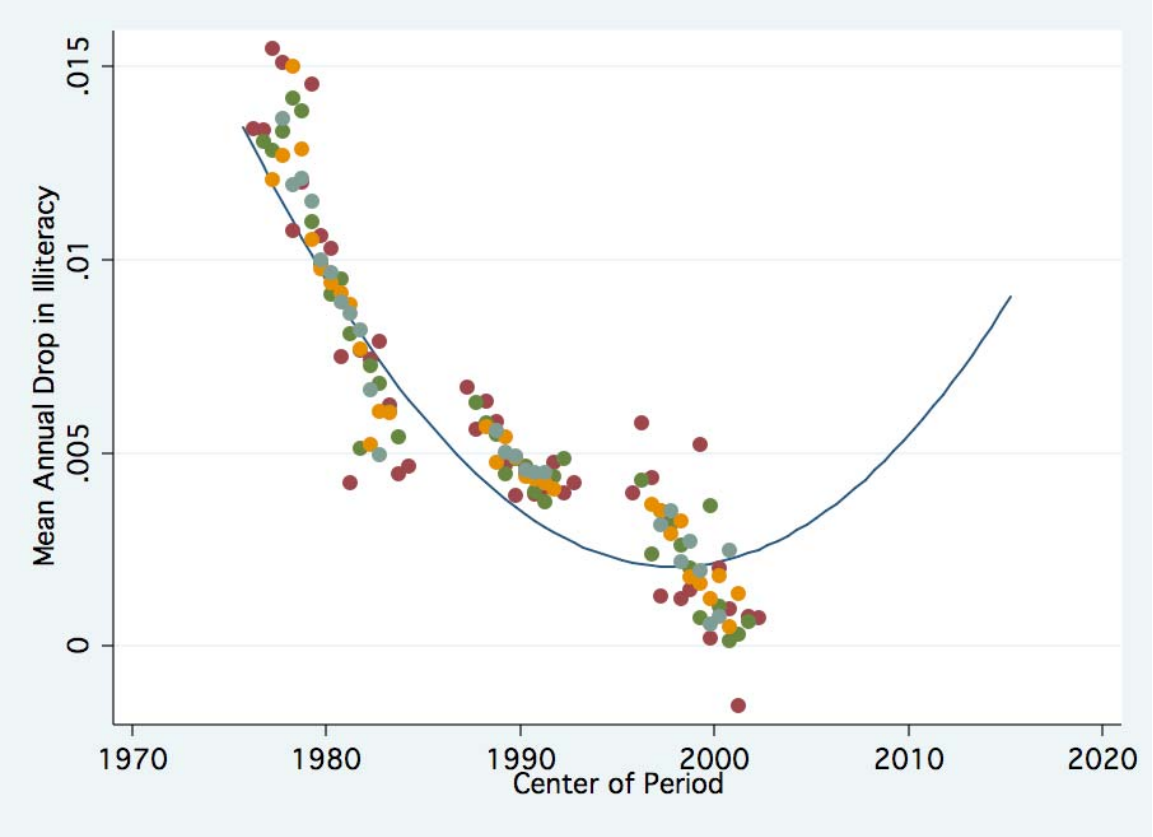

Source: Authors' analysis.

The cubic trend runs against what we observe in the data, counter-intuitively suggesting that large gains in literacy ought to appear in the absence of the literacy program. It is difficult to overstate the strangeness of this result. Putting aside breaks in the survey data, increases in literacy have dropped consistently for the entire time period for which we have data. Furthermore, the data shows literacy gains slowing where the cubic fit shows that by 1998 they hold steady and then begin growing faster. Unless there is some explanation for why these gains should show up, when the pre-*Robinson* history shows the opposite, the cubic time trend sets an unreasonably high bar for the *Robinson* program.

Similarly, the literacy gains have become so small that a quadratic time trend would show increasing illiteracy during the *Robinson* program, which hardly makes any sense. Simply put, the observed leveling off of the literacy rate in the early 2000s makes a tricky business of trying to pick out a reasonable trend.

A constant effect on literacy

Ortega and Rodríguez initially assume that *Misión Robinson* raises or lowers the adult literacy by a constant. That is, when the program starts, the literacy rate should move up by a fixed number of percentage points. More sensibly, one might imagine the *Misión* to have an increasing effect over time as more adults come through the system. This approach is discussed later in their paper, but since the government claims that most of the students entered early in the program, such an effect would be even more difficult to observe in the survey data. Regardless, any sort of complicated model of *Robinson*'s effect is necessarily constrained by the fact that there are few data points since the program's inception with which to work. Ortega and Rodríguez's data contains only five data

points for the *Robinson* period, one of which was based on data in part just as the program was gearing up.

Further complicating this model is that any misspecification of the *Robinson* effect will likely be picked up by the underlying time trend. Consequently, proper specification of the time trend becomes even more critical. For simplicity, we will just examine the constant effect. As the first students were going through the program while the second semester 2003 survey was underway, we will examine models with and without a lag of the *Robinson* dummy by one period. With respect to the untransformed, cubic time trend, lagging has no appreciable impact on the results.

Results

At first glance, Ortega and Rodríguez's results appear to be robust. Table 4 shows variations on the 'Ortega-Rodríguez' model, some of which are variations reported in the paper. As closely as possible, columns 1 and 2 replicate the first two columns of Ortega and Rodríguez, Table 2. Column 3 replicates the first result in Ortega and Rodríguez, Table 3, column 2. These three results are broadly similar in that none show a significant coefficient for the *Robinson* dummy (Rob).

Ortega and Rodríguez show the baseline cubic and sixth degree results (columns 1 and 2), and lagged cubic (column 3), but does *not* report results for the sixth-degree polynomial trend with the lagged *Robinson* dummy (column 4). As column 4 shows, the coefficient in that case is positive and significant. In fact, as Table 5 shows, when using a lagged *Robinson* dummy only the linear and cubic trends fail to show a significant positive effect on the official literacy rate.

TABLE 4
Near-OR Regressions

	Baseline		**Lagged**	
	(1)	(2)	(3)	(4)
Rob	-0.0006	0.004	0.001	0.005
	(0.003)	(0.002)	(0.002)	(0.002)***
1994	0.004	0.015	0.003	0.015
	(0.002)	(0.001)***	(0.002)	(0.001)***
Time	0.0014	0.0005	0.0012	0.0004
	(0.0003)***	(0.0003)*	(0.0002)***	(0.0002)*
Time^2	3E-5	1.3E-4	3E-5	1.1E-4
	(1E-5)**	(3E-5)***	(1E-5)**	(3E-5)
Time^3	1.0E-6	2.1E-5	9E-7	1.9E-5
	(2E-7)***	(3E-6)***	(1E-7)***	(4E-6)***
Time^4		9E-7		8E-7
		(1E-7)***		(2E-7)***
Time^5		1.5E-8		1.4E-8
		(3E-9)***		(3E-9)***
Time^6		9E-11		8E-11
		(2E-11)***		(2E-11)***
constant	0.937	0.933	0.936	0.934
	(0.002)***	(0.001)***	(0.001)***	(0.001)***

Source: Authors' analysis.
Notes:
Newey-West standard errors in parenthesis adjusted for heteroskedasticity and first-order autocorrelation.
* Significant at the 10% level
** Significant at the 5% level
*** Significant at the 1% level

TABLE 5
Significance of the *Robinson* dummy

Polynomial Degree	1	2	3	4	5	6	7	8	9	10
					Without Lag					
***Robinson* Effect†**	-2.2	1.1 ***	-0.1	0.7	0.3	0.4	0.2	0.5	0.3	0.1
					With Lag					
***Robinson* Effect†**	-2.1	1.1 ***	0.1	1.1 ***	0.7 ***	0.5 ***	0.5 ***	0.6 ***	0.4 ***	0.3 ***

Source: Authors' analysis.
Notes:
* Significantly larger at the 5% level
** Significantly larger at the 2.5% level
*** Significantly larger at the 0.5% level
† Regression point estimate in percentage points. Significance indicates regression result is inconsistent with a negative or no effect.

However, the significance or insignificance of these test results is not really relevant to Ortega and Rodríguez's argument. They are not trying to argue that there was no program at all and that nobody learned to read. Their argument is that the program was a small-scale (or very small) program, and not a large-scale one.[12] The appropriate statistical test for their argument is one that assesses whether the data are consistent with various increases in the literacy rate. Only then, given possible increases in literacy consistent with the household data, is there the additional question – as explained in Tables 1-3 – of whether the increase found in the household data is consistent with a large-scale program.

TABLE 6
TITLE

Polynomial Degree		1	2	3	4	5	6	7	8	9	10
						Without Lag					
Robinson Effect†		-2.2	1.1	-0.1	0.7	0.3	0.4	0.2	0.5	0.3	0.1
Change in	0.0%	***									
Official	0.5	***		**							*
Literacy††	0.75	***		***			*	**		*	***
	1.0	***		***		**	***	***	*	***	***
						With Lag					
Robinson Effect†		-2.1	1.1	0.1	1.1	0.7	0.5	0.5	0.6	0.4	0.3
Change in	0%	***									
Official	0.5	***		*							
Literacy††	0.75	***		***			*	***		**	***
	1.0	***		***			***	***	**	***	***

Source: Authors' analysis.
Notes:
* Significantly smaller at the 5% level
** Significantly smaller at the 2.5% level
*** Significantly smaller at the 0.5% level
† Regression point estimate in percentage points
†† Significance indicates that the observed effect is inconsistent with a change in official literacy of the indicated size.

In Table 6, we show the estimate of the *Robinson* effect and test if the change in literacy is at least zero, half, three-quarters, or one percentage point. In this table, asterisks indicate that the regression

[12] The results of Table 5 merely indicate that if the *Robinson* dummy is not lagged, then the corresponding statistical tests are not very powerful.

result is inconsistent with a *Robinson* effect of a particular size. For example, the first square in the "Without Lag" section shows that the linear (polynomial degree 1) point estimate (-2.2) is inconsistent with any positive effect of the program on the measured literacy rate. On the other hand, the square next to that (polynomial degree 2) point estimate of 1.1 is consistent with change in the measured literacy rate of between 0 and 1.0 percent. Thus, a cubic trend is only consistent with a *Robinson* effect of less than half a percentage point of official literacy gains, but the quadratic and quartic are consistent with gains of more than a full percentage point.

We note once again that the linear and cubic trends run counter to all the other polynomials. Table 6 shows that, except for the linear trend (which Ortega and Rodríguez acknowledge is not appropriate) and the cubic trend, all the other regression specifications are consistent with an observed increase in the measured literacy rate of at least 0.5 percentage points. This is consistent, under reasonable assumptions as described above, with a large-scale program.

In other words, regression analysis of the household data – even allowing the flaws in Ortega and Rodríguez's methodology described above – show results that are consistent with a large scale program, given that these results would only be partly seen in the Households Survey data, for the reasons discussed in the first section of this paper. The results are even consistent with a large-scale program of the size reported by the government, with 1.4 million participants.

Conclusion

This paper shows that the econometric analysis of Daniel Ortega and Francisco Rodríguez does not provide statistical evidence regarding the size of the *Misión Robinson* national literacy program in Venezuela. As explained in the first section, the Household Survey data on which they relied are too crude a measure to make such an estimate. In addition, the statistical tests they used are not appropriate to test (falsify) the hypothesis that the program was a large-scale program.

In demonstrating this result, we make no estimate of the actual size of the program. We simply demonstrate that the statistical analysis of Ortega and Rodríguez does not shed any light on this matter.