

The Economic Consequences of Hugo Chavez: A Synthetic Control Analysis

Kevin Grier* and Norman Maynard†

June 2014

Abstract

We use the synthetic control method to perform a case study of the impact of Hugo Chavez on the Venezuelan economy. We compare outcomes under Chavez's leadership and policies against a counterfactual of "business as usual". We find that, relative to our control, per capita income fell dramatically. While poverty, health, and inequality outcomes all improved during the Chavez administration, these outcomes also improved in each of the control cases and thus we cannot attribute the improvements to Chavismo. We conclude that the overall economic consequences of the Chavez administration were bleak.

*Department of Economics, University of Oklahoma. email: angus@ou.edu

†Department of Economics and Finance, College of Charleston. email: maynardna@cofc.edu

1 Introduction

Despite decades of effort, we still do not fully understand why some countries grow rich and others do not. One intriguing potential answer is that the national leader is crucial to the performance of the economy. It seems intuitive that, for all his faults, Park Chung-hee was essential to South Korea's growth miracle. On the other hand, it seems equally intuitive that Kim Il-Sung and his progeny have been very bad for the economic progress of North Korea.

But intuition is a poor guide for establishing causality. To say that a leader caused an outcome is to say that without that particular leader, things would have been significantly different, and establishing the counterfactual can be a tricky process. Recently, Abadie and Gardeazabal (2003) and Abadie et al. (2010) have proposed an approach to case studies which creates a statistical synthetic control for comparison to an observed treatment.¹ They then identify the effect of a treatment in a single region by comparing outcomes there to those in the control constructed from similar regions and states that did not experience the same treatment.

Here we use the method to study the effects of one particular leader, Hugo Chavez, on his country, Venezuela. Chavez is a controversial figure who won 3 elections, re-wrote Venezuela's constitution, re-structured its Supreme Court, and survived coup and recall attempts. Poverty fell, inequality fell, clinics opened, and Chavez was hailed as a hero by many. At the same time, Venezuela is a major oil exporter and oil prices boomed during much of Chavez's tenure and many have questioned whether any positive outcomes under Chavez's tenure as president were really due to his leadership and policies. To answer this question, we need to know how Chavez's Venezuela fared relative to how it would have fared if politics and policies in Venezuela had remained "business as usual", and we measure business as usual by a weighted average of control countries that best fit Venezuela's experience in the 30 years before Chavez. Using Abadie's terminology, we create a synthetic Venezuela and compare its outcomes to those that occurred under Chavez, specifically investigating average incomes, health & poverty outcomes, and inequality.

We find that average income rose significantly slower under Chavez than it would have without him, as shown by the performance of the synthetic control. The gap between the control and the treated unit is quite large, on the order of thousands of dollars per capita.

¹Abadie and Gardeazabal (2003) developed the method in order to identify the effect of terrorism in the Basque Country in Spain. Abadie et al. (2010) explain the technique in detail and analyze the effect of a change in California's tobacco regulations. Gautier et al. (2009) and Montalvo (2011) provide additional applications to the analysis of terrorism. And of course the synthetic control method has seen numerous recent applications including affirmative action (Hinrichs, 2012), compulsory voting (Fowler, 2013), economic liberalization (Billmeier and Nannicini, 2013), and natural disasters Cavallo et al. (2013).

We also find that while health and poverty outcomes improved under Chavez, they are not particularly different than what occurred in the synthetic control. Finally, we see little evidence that Chavez lowered inequality beyond what occurs in the synthetic control. Thus, under Chavez, Venezuela experienced lower average incomes, but did not produce any greater levels of health improvement, poverty abatement, or inequality reduction, as it likely would have without his leadership and policies.

Our paper is related to other recent work on the impact of national leaders. Jones and Olken (2005) address causality between leaders and economic growth by focusing on leaders who died in office due to natural causes or accidents. They find that leaders are highly influential, especially in autocratic systems with few constraints on the executive. Besley et al. (2011) use an expanded version of the same dataset to identify the effect of leaders' level of education on growth, showing that having a more educated leader leads to higher growth rates. Easterly and Pennings (2014) take a growth accounting approach and argue that "only a small fraction of the variation in growth in autocracies can be explained by variation in leader quality."

The paper most closely related to ours is Garcia Ribeiro et al. (2013). They study the effect of the 1959 Cuban revolution on Cuba's subsequent per-capita income, finding that it is substantially depressed relative to their control. They do not consider any other possible economic effects of the revolution.²

Our paper is organized as follows. Section 2 provides some background on the Venezuelan economy. Section 3 explains how the synthetic control is created, and how potential control countries are chosen. Sections 4, 5, 6 and 7 present our results on per-capita income, health, poverty, and inequality, respectively. Section 8 concludes.

2 Recent Venezuelan history

Venezuela is a highly urbanized, ethnically diverse South American country of roughly 30 million people. Despite gaining independence in 1821, the country did not enjoy a peaceful transfer of power from one party to another until 1969. Venezuela is a founding member of OPEC, and since the 1973 oil crisis, its economic fortunes have risen and fallen with oil prices. From 1958 to 1998 Venezuela was democratic with its politics dominated by two political parties. However, poor economic performance in the 1990's created popular unrest.

²We thank Alberto Abadie for bringing this working paper to our attention when giving us comments on an earlier draft of our paper.

Hugo Chavez was a military officer who launched an unsuccessful coup in 1992. He was pardoned in 1994 by then president Rafael Caldera. In 1998 Chavez handily won the presidency and launched his “Bolivarian Revolution” in 1999 with the formation of a constitutional assembly that produced a new constitution, which placed a lot of emphasis on social progress and human rights. It also converted the legislature from bi-cameral to unicameral and greatly increased the powers of the executive branch. It increased the presidential term from 5 to 6 years and allowed for the president to hold the office for 2 consecutive terms. In 2000 a mega-election was held for the presidency, the new national assembly and other offices. Chavez won and his party won 101 of 165 seats in the Assembly, which then granted Chavez the right to rule by decree. In 2004, Chavez and the national assembly increased the number of Supreme Court justices from 20 to 32, allegedly “packing” the court with his supporters. Chavez won re-election in 2006, and in 2007 was granted the right to run again in 2012 in a referendum.³

In sum, whatever its virtues, Chavez’s revolution concentrated power in the executive branch and weakened the checks and balances that existed in the previous system.

The Chavez regime was known for its stated emphasis on improving access to food, education and medical care for the poor, its large increases in social spending, its ties to Cuba, its policy of selective nationalizations and price controls, and its antipathy to the U.S. government. The fact that on the one hand, Chavez is viewed by many as a hero who improved the lives of average Venezuelans, and, on the other hand, is considered by many others as one who severely damaged the Venezuelan economy speaks to the difficulty of establishing a clear counterfactual, which we undertake in what follows.

3 Creating the synthetic control

3.1 The method

We want to estimate the difference between the observed levels of our outcome variables and where that variable would have been if Chavez had never taken office. To do this, we need a control group that gives us an idea where the outcome variable for the country in question would have been. Since an exact control does not exist (we have no post-1999 observations on Venezuela where Chavez was not president), we create a control group by synthesizing the

³A full analysis of the events and policies of the Chavez regime is beyond the scope of this paper, but beyond what is listed in the text, Chavez took greater control over the national petroleum company, nationalized a string of other companies, imposed price controls on many items and limited access to foreign exchange to many individuals and businesses. He also experienced a failed coup in 2002 and a failed recall in 2004.

performance of countries similar to Venezuela.⁴ The synthetic Venezuela is simply a weighted average of the other countries; for example, per-capita income in synthetic Venezuela is a weighted average of per-capita income in the control countries. Since none of the control countries have Chavez (or someone with similar policy interventions) as their leader, the synthetic control should be independent of the direct effects of his election. To the extent that the synthetic control accurately captures the other influences on the Venezuelan economy, it represents precisely the counterfactual we desire.⁵

The two most important concerns here are the choice of countries (which we address in the next subsection) and the assignment of weights across countries. One solution would be to simply choose weights that minimize the difference between the pre-Chavez performance of the synthetic control and the pre-Chavez performance of Venezuela. We impose the restriction that the weights on each country in the control be non-negative and sum to one. In practice, many countries are assigned a weight of zero, with only a handful of countries receiving positive weight for each outcome variable.

While the above approach might simply put high weight on countries that happen to perform similarly to Venezuela during our pre-Chavez sample, we wish to select countries that actually have similar economic processes to Venezuela. That is, we seek countries with similar relationships between outcome variables such as per-capita income and their determinants. To achieve this, we use time-invariant indicator variables and minimize the difference between the weighted average of these variables for the synthetic control and the values of these variables for Venezuela. Countries with more similar indicator variables to Venezuela will receive higher weight in creating the control. We also put more value on similarity in explanatory variables that have more influence over the outcome variable.⁶

⁴For a more detailed technical discussion of the synthetic control method, see Abadie et al. (2010).

⁵We can think of synthetic Venezuela as representing the outcomes the country would have experienced if it had elected a weighted average of the leaders of the control countries. To the extent this alternate electoral outcome would have required different preferences by Venezuelan voters, we are identifying the treatment effect of both Chavez and the voter preferences that led to his election.

⁶Mathematically, let X_1 represent the vector of all indicator variables for Venezuela and Y_1 represent the vector of pre-Chavez outcome variables. Similarly, let X_0 be the matrix whose columns are vectors of indicator variables for the potential control countries and Y_0 is the matrix whose columns are vectors of their pre-Chavez outcome variables. The column vector of weights is designated W , with the weight assigned to country j equal to the j -th individual element of W . We choose weights to minimize the distance function $D = \sqrt{(X_1 - X_0W)' V (X_1 - X_0W)}$, where V is a positive-definite diagonal matrix. The procedure for creating the synthetic control follows three steps: (1) select any positive-definite diagonal matrix V (we use the identity matrix) and solve for weights as a function of this matrix, $W(V)$; (2) given this function $W(V)$, choose V so that it minimizes the mean squared prediction error over the pre-Chavez period $MSE = (Y_1 - Y_0W(V))' (Y_1 - Y_0W(V))$; (3) repeat until MSE is minimized, then find the final distance-minimizing weight matrix W^* .

3.2 The control countries

Abadie and Gardeazabal (2003) make the point that the synthetic control method described above depends on choosing countries that plausibly have similar economic processes to the country with the intervention. Because of this, we need to use data for countries with important similarities to Venezuela.

In order to capture similarities of geography and heavy oil exports, we use almost all Latin American countries and OPEC member countries with available data.⁷ We also include Canada and Norway, which are major oil exporting countries but are not OPEC members.

We are studying 5 outcome variables: Per capita income, life expectancy, infant mortality, poverty, and inequality. We thus create a distinct synthetic control for each of these variables. A full list of countries included in each synthetic control group, along with the weights assigned, is included in Table 1. In total, we consider 20 potential control countries, although data is not available from all countries for all synthetic controls.

3.3 The synthetic control for per-capita income

For the case of per-capita income we have 19 potential control countries. For indicator variables we use income levels from four pre-Chavez years (1970, 1977, 1989, 1998). We use the average population growth rate 1970-1998 from the Penn World Table, openness to trade in 1970 and 1990 also from the Penn World Table, and two educational attainment variables in 1970 and 1990 from the Barro and Lee (2010) database.

When we apply the algorithm to create a synthetic control, we end up with a control that is composed of 1% Brazil, 15% Canada, 43% Iran, 32% Mexico, and 9% Peru. Table 2 shows the pre-Chavez values of the indicator variables for actual Venezuela and for the synthetic Venezuela. Average income in the two series differs by only \$14 and the values of the control variables between the control and Venezuela all match fairly closely. When we use the control to predict Venezuela's per capita income from 1970 – 1998, the root mean squared percentage error (RM-SPE) is 0.068, or less than 7 percent.

⁷We do not consider Bolivia, Ecuador, or Nicaragua as potential controls as they experienced a political transition and constitutional change similar to Venezuela's during our sample period.

3.4 Why do we need this fancy control?

It is important to note that the synthetic control mimics pre-Chavez Venezuela much better than simpler types of controls. In Figure 1 we display the evolution of average per-capita income for Latin American (LA) control countries and our OPEC control countries. As can be seen, neither are appropriate controls. While the OPEC average series seems to move in a similar fashion, it is very far away from the Venezuela experience and would require a high degree of extrapolation to use as a control.⁸ Furthermore, Table 2 shows that the average values in Venezuela of the covariates we used to construct the synthetic control are not well matched by either the LA or OPEC average values making any extrapolation fairly dangerous.

Even judiciously picking a particular country as a control is not as good as the synthetic control. Figure 2 compares Venezuela to the single closest LA country (Argentina) and the single closest OPEC country (Iran). While these fit better than the regional averages, their RMSPEs are .1467 and .2718 respectively, which are far worse than the RMSPE for our synthetic control (0.068). Finally, Figure 3 displays our weighted average synthetic control and its superior ability to match per capita incomes in pre-Chavez Venezuela.

4 Results on per-capita incomes

4.1 Main result

Now we are ready to perform our experiment. We are taking the synthetic control as our representation of “business as usual” in Venezuela and we seek to see whether outcomes under Chavez deviate from this counterfactual benchmark. We have confidence in our control because it matches closely to actual outcome for the 30 pre-Chavez years and the indicator variables match fairly well between the control and Venezuela.

Figure 3 plots actual and synthetic Venezuela both over the pre-Chavez period that we used to construct the control and the post Chavez period where the control serves as our counterfactual. As can be seen, synthetic Venezuela tracks actual Venezuela upward through the high oil prices of the 1970s, downward during the financial turmoil of the late 1970s and early 1980s, and through the income increase in the early 1990s. However, after 1998 (the year of Chavez’s successful presidential campaign) synthetic and actual Venezuela sharply diverge. By 2003 Venezue-

⁸The synthetic control technique does not allow any extrapolation outside of the support of our control country data. This is one practical reason to include higher income countries like Norway and Canada—to ensure that the Venezuelan experience lies inside the support of the control data.

lan per-capita income is more than \$4,000 below that of synthetic Venezuela, and at the end of our sample in 2009, actual income is more than \$4,300 less than the synthetic control.⁹ Figure 4 shows the same information, only highlighting the divergence of actual Venezuela from where it would have been (as projected by synthetic Venezuela). It appears that Chavez's leadership and policies were quite bad for the overall level wealth in Venezuela. The RMSPE post Chavez is 0.28 or almost 4.5 times as large as the pre-Chavez error.

4.2 Placebo test

The above findings are provocative, but need some additional context. To truly claim that Venezuelan performance is an outlier, we need some indication of out-of-sample precision in the absence of similar interventions. If overall predictability somehow fell after 1999, or if our synthetic control was actually a poor predictor of Venezuela pre-Chavez compared to what could be achieved for other countries, then our results would not be informative. We address these issues with a series of placebo tests. Specifically, we repeat the above exercise for each of the potential control countries we considered above.

That is to say, for each of the 19 countries, we develop a synthetic control using 1970 – 1998 data and compare the control to the actual country in both the pre- and post-intervention periods. We find that the average 1970 – 1998 RMSPE is 0.132 and the average 1999 – 2009 RMSPE is 0.151. In other words, the control countries as a group do not exhibit anything like the large deterioration in fit that Venezuela does. This finding lends support to the idea that the effect we see in Venezuela is attributable to the leadership and policies of Hugo Chavez.

Figure 5 presents a graph of the placebo test results for those countries whose pre 1999 RMSPE is within 150% of Venezuela's. The results show that the post-Chavez divergence in Venezuela is visibly larger than any of the divergences in the other countries. We see a strong Chavez effect in Venezuela and no Chavez effect outside of Venezuela.

4.3 Move Chavez test

As an additional robustness test, we pretend that Chavez came to power in 1988 and repeat our experiment. That is, using the same potential control countries and indicator variables, but this time only from 1970 – 1988, we create a synthetic control for Venezuela and compare its

⁹We have experimented with using a larger number of control variables and not enforcing covariate matching between actual and synthetic Venezuela and we obtain very similar results. These other exercises are available from the authors on request.

performance with that of actual Venezuela before and after 1988. This test allows us to see if the result we found above is due to an inability to predict Venezuela out-of-sample for any time period, not just the Chavez era.

The composition of the new synthetic controls are presented in Table 3. The weights are nearly identical, with a slight shift from Peru and Mexico to Brazil. The RMSPE for 1970 – 1988 is 0.070 while the 1989 – 1999 RMSPE is 0.085. We see no real deterioration of the control’s ability to predict Venezuelan per-capita GDP in the out of sample interval here where there was no actual intervention. Figure 6 shows very little change in our ability to track actual Venezuela post-intervention when we falsely assume that the intervention occurred in 1988.¹⁰

All in all, there is a strong case to be made that in terms of real income per capita, The Chavez administration’s leadership and policies seriously hurt economic performance in Venezuela relative to what likely would have occurred with a “business as usual” set of policies (as given by the predictions of the synthetic control).

5 Results on health

Of course, GDP per capita is not the only dimension along which we should evaluate the Venezuelan economy. The Chavez government may not have considered average incomes a high priority, instead preferring to focus on improving health outcomes, reducing poverty or reducing inequality. We address each of these issues below, starting with health.

To identify the effect of Chavez’s presidency on health outcomes, we consider annual observations from two variables: infant mortality and life expectancy, 1970-2012.¹¹ We have 17 control countries for our two measures.¹² The compositions of the synthetic control for both measures of health are presented in Table 1. Figure 7 shows that the synthetic controls do a good job of capturing the pre-Chavez performance for both measures, and that post-Chavez there is almost no discernible effect of Chavez’s presidency on infant mortality. For life-expectancy we do see a small degree of divergence, as life expectancy slowly drops roughly one year below the synthetic control, a gap which persists through the end of the sample. Figure 8 shows the placebo graphs with all 17 control countries.¹³

¹⁰We performed this same check with our inequality and infant mortality data, yielding similar results. The composition of these synthetic controls are also displayed in Table 3.

¹¹We also considered the mortality rate for children under age 5 for the same years, which yielded similar results.

¹²Our indicator variables are investment and government shares in 1970 and 1990; the three Polity IV scores from 1970 and 1990; the standard deviation of inflation; access to improved sanitation in 1990; the 1970 level per capita GDP; 1990 maternal mortality rate; oil wealth per capita in 1995; and the 1970 level of each annual series.

¹³For life expectancy, only 6 of the 17 countries have a pre 1999 RMSPE within 150% of Venezuela’s; for infant

6 Results on poverty

We next consider the effect of Chavez's election on the poverty rate, measured by the percentage of the population living below five dollars per day. Comparable cross-country poverty data is difficult to find, and would likely be impossible if we required annual observations. Fortunately, the synthetic control method allows us to use observations that cover large enough periods that we have at least one observation per period for Venezuela and each of 13 control countries.¹⁴ We are able to find poverty rates for three periods prior to 1998, and annual observations from 1998 to 2006, giving us four pre-Chavez periods and eight post-Chavez.¹⁵

Figure 9 shows the poverty measure for Venezuela and for the synthetic control, which is composed of 20% Argentina and 80% Colombia. With so few pre-election data points and a limited set of control countries, it is unsurprising that the model does not fit as well as the average income model. Nevertheless, the synthetic control still captures the general upward trend in the poverty rate over the 1980s and 1990s. After the election, poverty experiences upward jumps in 2002 and 2003, but after 2003 trended down, falling significantly below the election-year levels by 2006. This pattern is roughly matched by the synthetic control, with the control in 2006 lying only slightly above the actual poverty rate in Venezuela.

To see how different actual and synthetic Venezuela really are, we create placebos just as we did for average income. We perform the same analysis for all 14 control countries and plot the deviations from their synthetic controls in Figure 10, with the difference between synthetic and actual Venezuela shown in bold. This figure illustrates that the divergence between actual and synthetic Venezuela is well within the range seen from countries which do not experience the same change in leadership.¹⁶ We find no evidence that Chavez had a significant impact on the poverty rate in Venezuela.

mortality, only 2 countries do.

¹⁴We use the same indicator variables for poverty that we did for health.

¹⁵For countries with more than one observation in each pre-1998 period, we average the available observations. We extrapolate the first-period (1980–1987) poverty rate for El Salvador and Paraguay, and interpolate the second-period (1988–1992) rate for Peru and the third-period (1993–1997) rate for Guatemala. We also interpolate a number of annual observations for the post-1998 observations. In most cases, the change in observed rates before and after the missing observation is only a few percentage points. Because of this we believe that linear interpolation is appropriate for the analysis.

¹⁶The RMSPE for the 1999 – 2006 period is actually less than the RMSPE for 1998 and earlier, indicating that we have a closer out-of-sample fit than pre-intervention. This is likely due to the small number of pre-intervention observations.

7 Results on inequality

Having considered average incomes, health outcomes, and poverty rates, we next use the synthetic control method to study the effect of Hugo Chavez as the leader of Venezuela on inequality. We want to acknowledge up front that there is no perfect panel dataset of Gini coefficients for applied researchers to use. We took two approaches to obtaining one. First we used the WIDER dataset and created our own panel by using all Ginis with quality at least as good as Venezuela's, averaging any multiple observations in a single year, and interpolating to fill in any missing years. That is to say, we made our own sausage.

Our other approach was to use the SWIID database of Solt (2009), which uses its own methods to make a panel from the WIDER and other data (store bought sausage). Since both approaches produced almost identical results, we only present here results using the SWIID database.

We have 14 potential control countries, with a synthetic control for Venezuela that is composed of 3% Colombia, 4% Norway, and 92% Uruguay. Figure 11 shows that, pre-Chavez, our synthetic control generally captures the trend of inequality in Venezuela. Post Chavez, the control continues to track actual Venezuela until 2003, when Venezuelan inequality falls rapidly and steadily compared to the level predicted by the synthetic control, although the gap levels off toward the end of the sample.

Figure 12 shows our now familiar placebo tests; as with income, we show the countries with pre-1999 RMSPEs within 150% of Venezuela's. We see that relative to other countries' synthetic controls, our control for Venezuela is fairly accurate pre-Chavez. The figure shows that, although inequality after 2003 has certainly fallen, for the entire post-Chavez period we do not have strong evidence of lower inequality in Venezuela relative to what might have been without Chavez.

8 Conclusion

Identifying the effect of national leaders on economic outcomes is an important, but tricky, business. While in experimental sciences comparison to the right control group is relatively straightforward, clean experiments at the level of national economies are few and far between. Using the synthetic control method of Abadie and Gardeazabal (2003) and Abadie et al. (2010), we create a synthetic Venezuela, indicating how the country is likely to have performed if Hugo Chavez had not been elected to the Venezuelan Presidency in 1998.

We find that although average incomes rose somewhat during his time as president, they

lagged far behind where they might have been if Chavez had not taken office. During his leadership, life expectancy grew while mortality and poverty rates fell. Yet we find no evidence that these gains are any different than they would have been with another national leader. If anything life expectancy rose slower than what the control predicted. On the positive side of the ledger, it may be the case that inequality fell faster than it would have without Chavez in office after 2003; but we see no evidence that his term in office as a whole reduced inequality more than it would have without him. While these results are not produced in a complete simultaneous model, they suggest that the observed reductions in inequality and poverty should not be attributed to Chavez's leadership, and that Chavez's legacy may have more to do with the harm caused to per capita incomes.

References

- Abadie, A., Diamond, A., Hainmueller, J., 2010. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association* 105 (490), 493–505.
- Abadie, A., Gardeazabal, J., 2003. The economic costs of conflict: A case study of the Basque Country. *American Economic Review* 93 (1), 113–132.
- Barro, R., Lee, J.-W., April 2010. A new data set of educational attainment in the world, 1950–2010, NBER Working Paper No. 15902.
- Besley, T., Montalvo, J., Reynal-Querol, M., 2011. Do educated leaders matter? *The Economic Journal* 121 (554), 205–227.
- Billmeier, A., Nannicini, T., 2013. Assessing economic liberalization episodes: A synthetic control approach. *The Review of Economics and Statistics* 95 (3), 983–1001.
- Cavallo, E., Galiani, S., Noy, I., Pantano, J., 2013. Catastrophic natural disasters and economic growth. *Review of Economics and Statistics* 95 (5), 1549–1561.
- Easterly, W., Pennings, S., January 2014. How much do leaders explain growth? an exercise in growth accounting, Working Paper.
- Fowler, A., 2013. Electoral and policy consequences of voter turnout: Evidence from compulsory voting in Australia. *Quarterly Journal of Political Science* 8, 159–182.

- Garcia Ribeiro, F., Stein, G., Kang, T., May 2013. The cuban experiment: Measuring the role of the 1959 revolution on economic performance using synthetic control, Working Paper.
- Gautier, P., Siegmann, A., Van Vuuren, A., 2009. Terrorism and attitudes toward minorities: The effect of the Theo van Gogh murder on house prices in Amsterdam. *Journal of Urban Economics* 65 (2), 113–126.
- Hinrichs, P., 2012. The effect of affirmative action bans on college enrollment, educational attainment, and the demographic composition of universities. *The Review of Economics and Statistics* 94 (3), 712–722.
- Jones, B., Olken, B., Aug. 2005. Do leaders matter? National leadership and growth since World War II. *The Quarterly Journal of Economics* 120 (3), 835–864.
- Montalvo, J., 2011. Voting after the bombings: A natural experiment on the effect of terrorist attacks on democratic elections. *The Review of Economics and Statistics* 93 (4), 1146–1154.
- Solt, F., 2009. Standardizing the world income inequality database. *Social Science Quarterly* 90 (2), 231–242.

Table 1: Estimated Synthetic Control Weights for each Outcome Variable

	Outcome Variables				
	Income	Poverty	Infant Mort	Life Exp	Inequality
Algeria	0.00	-	0.00	0.00	-
Argentina	0.00	20.38	4.98	83.79	0.00
Brazil	0.95	0.00	0.00	0.00	0.00
Canada	15.02	-	0.00	0.00	0.00
Chile	0.00	0.00	0.00	0.00	0.00
Colombia	0.00	79.62	0.00	0.00	3.47
Costa Rica	0.00	0.00	0.00	0.00	0.00
El Salvador	0.00	0.00	0.00	0.00	0.00
Guatemala	0.00	0.00	12.28	0.00	-
Honduras	0.00	0.00	0.00	0.00	-
Indonesia	0.00	-	-	-	0.00
Iran	43.06	-	0.00	0.00	-
Iraq	0.00	-	-	-	-
Mexico	31.92	0.00	0.00	0.68	0.00
Nigeria	-	-	0.00	1.97	0.00
Norway	0.00	-	27.06	0.00	4.29
Panama	0.00	0.00	55.68	7.06	0.00
Paraguay	0.00	0.00	-	-	-
Peru	9.06	0.00	0.00	6.50	0.00
Uruguay	0.00	0.00	0.00	0.00	92.24

Note: Columns show the weight assigned to each country in the synthetic controls for Venezuela. Each column includes a synthetic control for a different outcome variable. A dash (-) indicates that the country is not available in the dataset for the given comparison. Values are in percentage points. Rounding errors may prevent columns from summing to 100.

Table 2: Average Income Predictor Means

Variables	Actual Ven	Synthetic Ven	LA Average	OPEC Average
avg. GDP per capita	9,071	9,085	5,573	4,376
Pop. Growth Rate	2.67	2.40	2.10	2.66
Openness 1970	46.64	48.50	40.18	76.57
Openness 1990	46.65	47.92	43.87	93.86
Total ed. 15+, 1970	3.74	3.81	3.98	1.93
Primary ed. 15+, 1970	2.91	2.66	3.11	1.47
Total ed. 15+, 1990	5.11	6.39	6.17	4.50
Primary ed. 15+, 1990	3.76	3.88	4.29	2.80

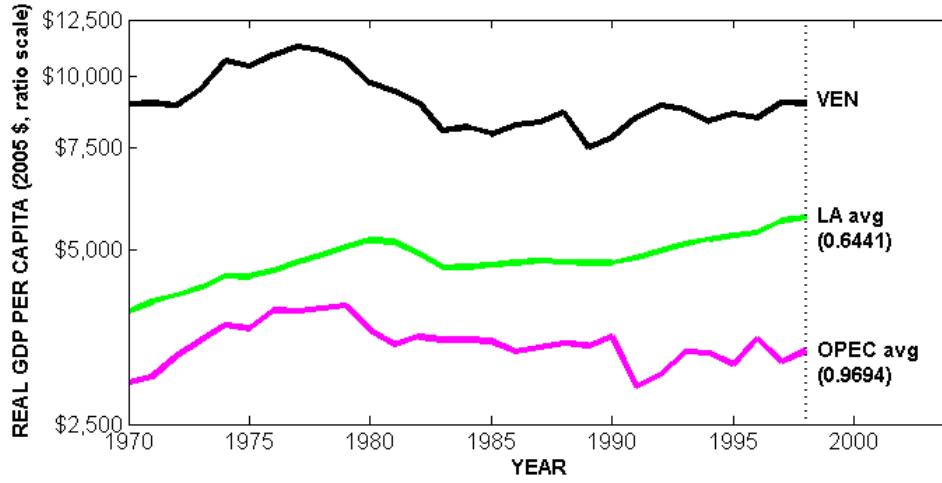
Note: Table shows the values of indicator variables and the average pre-Chavez outcome variable for actual and synthetic Venezuela. Average GDP is an average over the pre-Chavez period, 1970-1998. GDP per capita is measured in 2005 dollars. Population growth is in percentage points. Government share is a fraction of GDP, in percentage points. Polity and Ex. Const. follow the polity scores scale. Durability and the education variables are measured in years.

Table 3: Countries and Weights for Move Chavez Tests

	Income	Infant Mort	Inequality
Algeria	0.00	0.00	-
Argentina	0.00	0.00	0.00
Brazil	9.88	0.00	0.00
Canada	16.38	0.00	0.00
Chile	0.00	0.00	0.00
Colombia	0.00	0.00	46.29
Costa Rica	0.00	0.00	0.00
El Salvador	0.00	0.00	0.00
Guatemala	0.00	23.82	-
Honduras	0.00	0.00	-
Indonesia	0.00	-	0.00
Iran	42.86	0.00	-
Iraq	0.00	-	-
Mexico	27.72	0.00	0.00
Nigeria	-	0.07	0.00
Norway	0.00	48.02	26.77
Panama	0.00	20.09	26.94
Paraguay	0.00	-	-
Peru	3.16	0.00	0.00
Uruguay	0.00	8.00	0.00

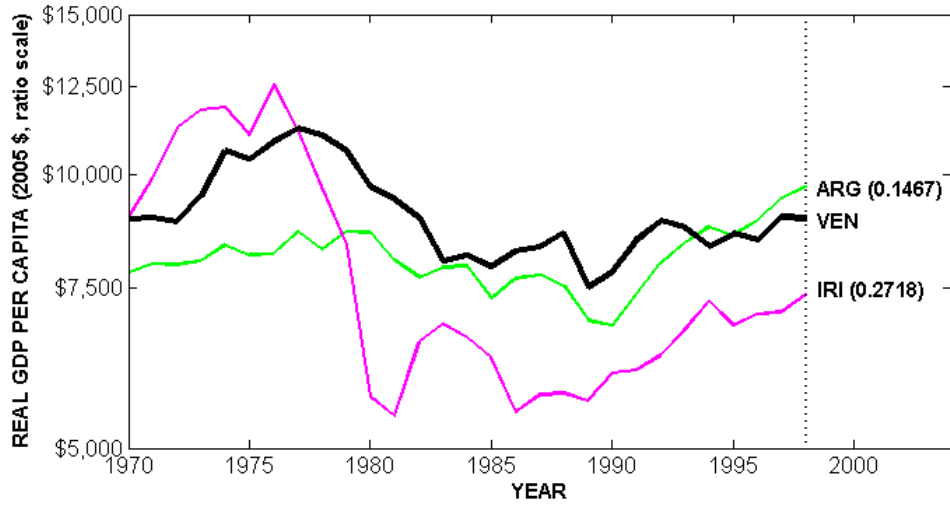
Note: Table shows the weight assigned to each country in the synthetic control group for Venezuela. A dash (-) indicates that the country is not available in the dataset for the given comparison. Values are in percentage points.

Figure 1: Potential controls using regional averages



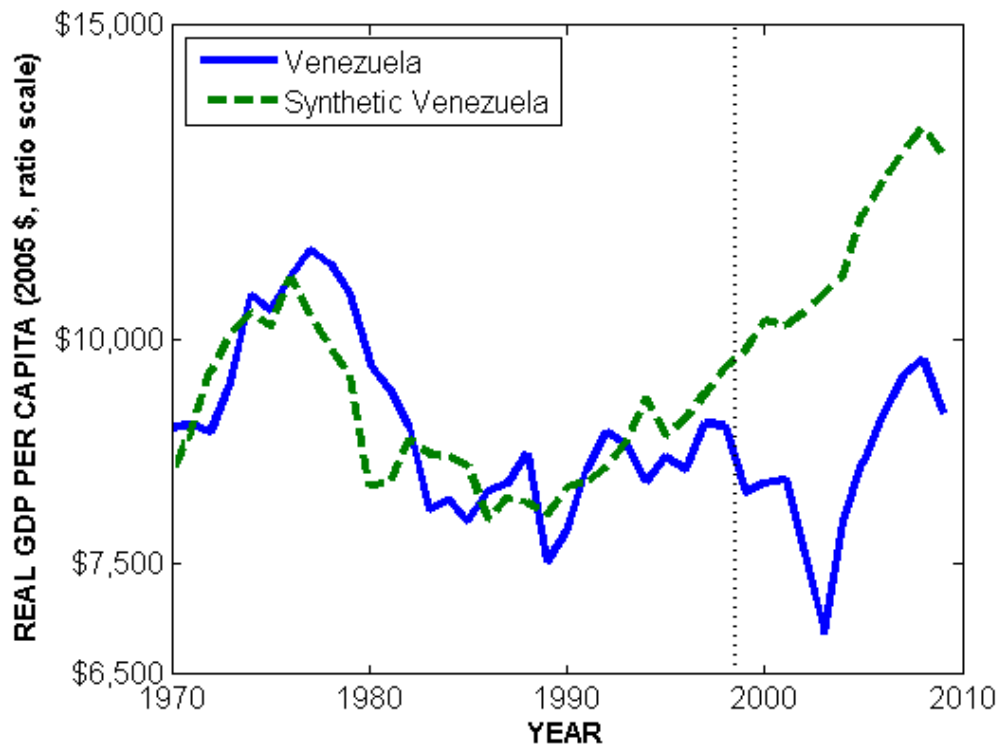
Note: Lines represent pre-1999 real GDP per capita for Venezuela (VEN), a simple average of Latin American countries (LA avg), and a simple average of OPEC countries (OPEC avg). The root mean square prediction error (RMSPE) of the two averages are in parentheses.

Figure 2: Potential controls using single countries



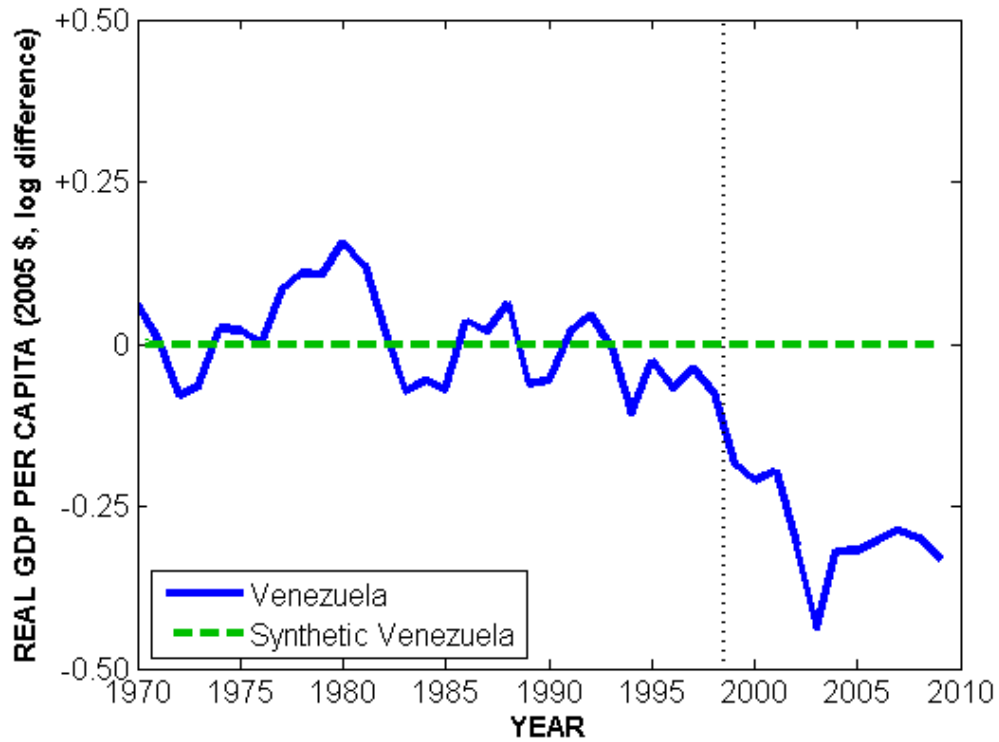
Note: Lines represent pre-1999 real GDP per capita for Venezuela (VEN), Argentina (ARG), and Iran (IRI). The root mean square prediction error (RMSPE) of the two control countries are in parentheses.

Figure 3: Per-capita income



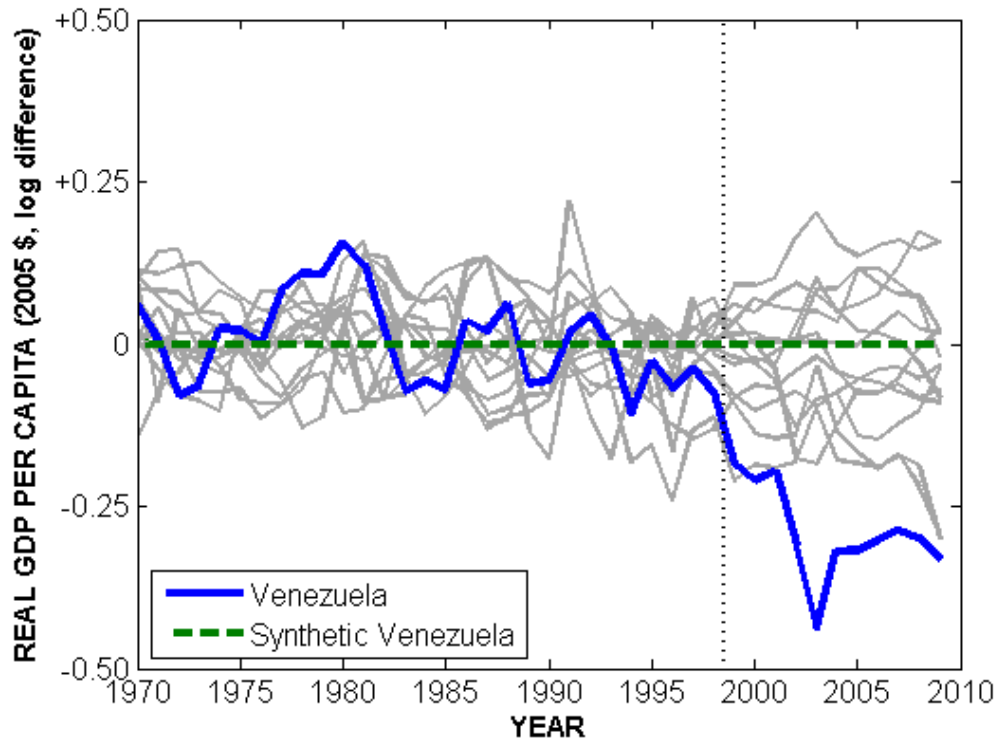
Note: The solid line represents observed (log) income per-capita in Venezuela, 1970-2009; the dashed line represents the synthetic control. The vertical dotted line indicates the end of the pre-Chavez years (1970-1998).

Figure 4: Per-capita income



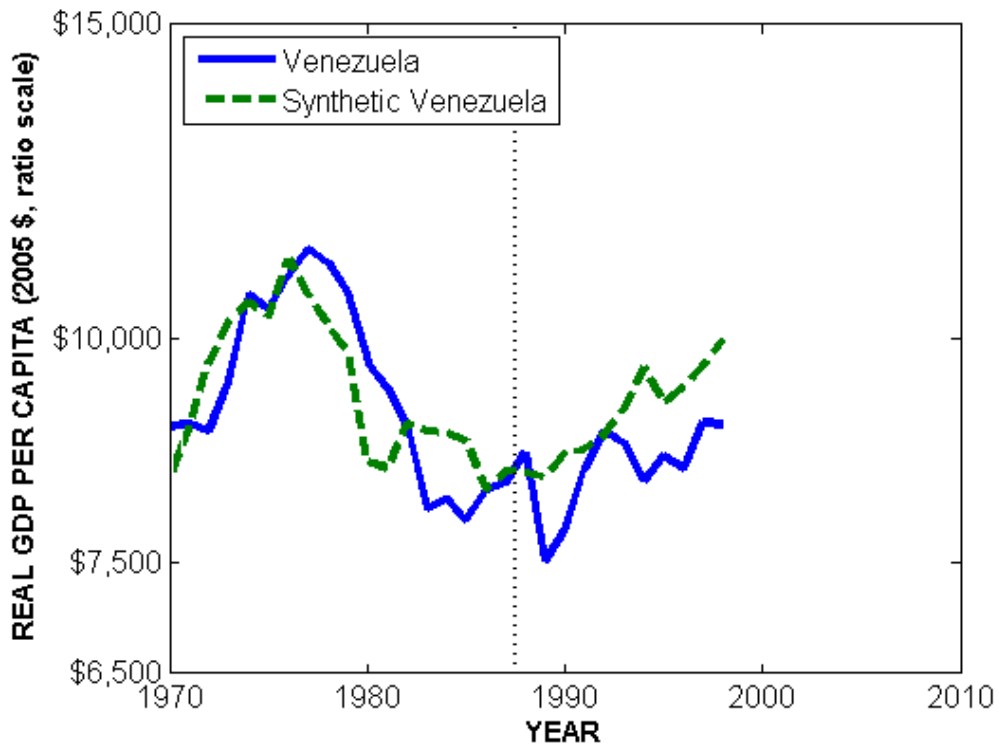
Note: The solid line represents the difference between observed (log) income per-capita in Venezuela, 1970-2009, and the synthetic control; the synthetic control (dashed line) is normalized to zero. The vertical dotted line indicates the end of the pre-Chavez years (1970-1998).

Figure 5: Per-capita income Placebo Tests, restricted countries



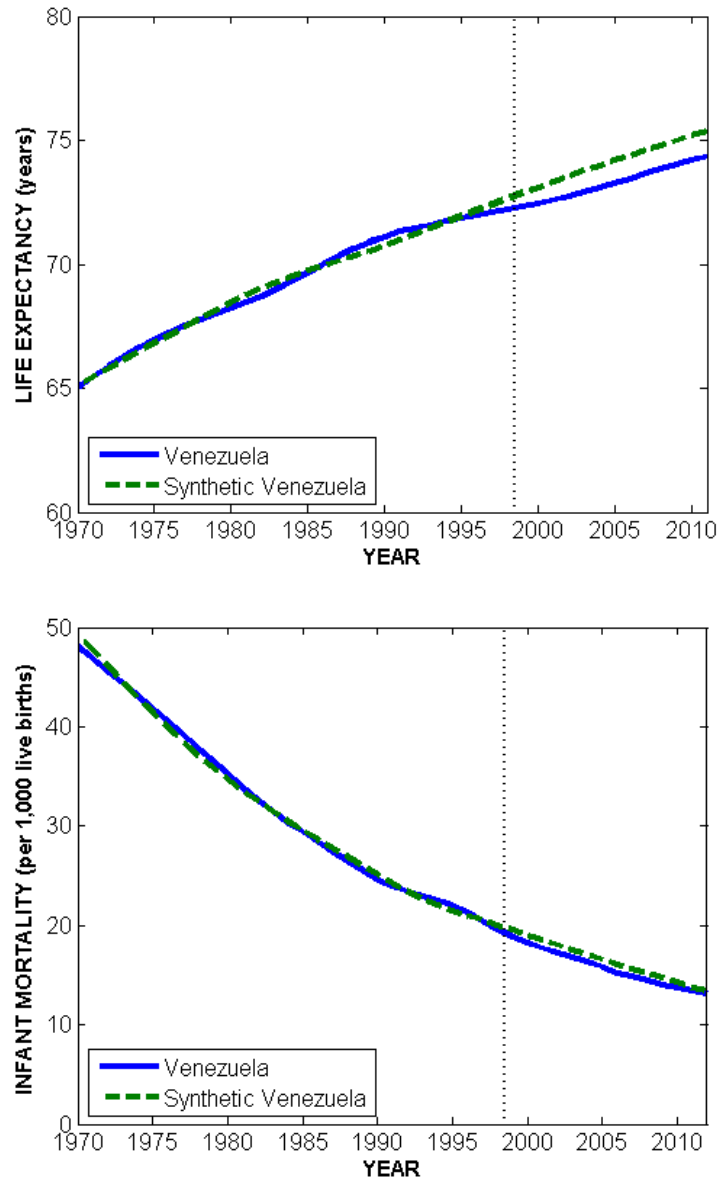
Note: The bold line represents the difference between observed (log) income per-capita in Venezuela, 1970-2009, and the synthetic control; the synthetic control (dashed line) is normalized to zero. Gray lines represent placebo tests: deviations from synthetic control for the other countries in the dataset. This graph only include countries with pre-intervention root mean squared prediction error (RMSPE) less than 0.1025 (1.5 times that of Venezuela). It drops Chile, Indonesia, Iran, Iraq, Norway, and Panama.

Figure 6: Income Move Chavez Test



Note: The graph reproduces the synthetic control for income with the timing of the intervention (election of Chavez) artificially set to 1988 and run through 1998 in order to have the same number of out-of-sample observations while avoiding overlap with the actual Chavez years.

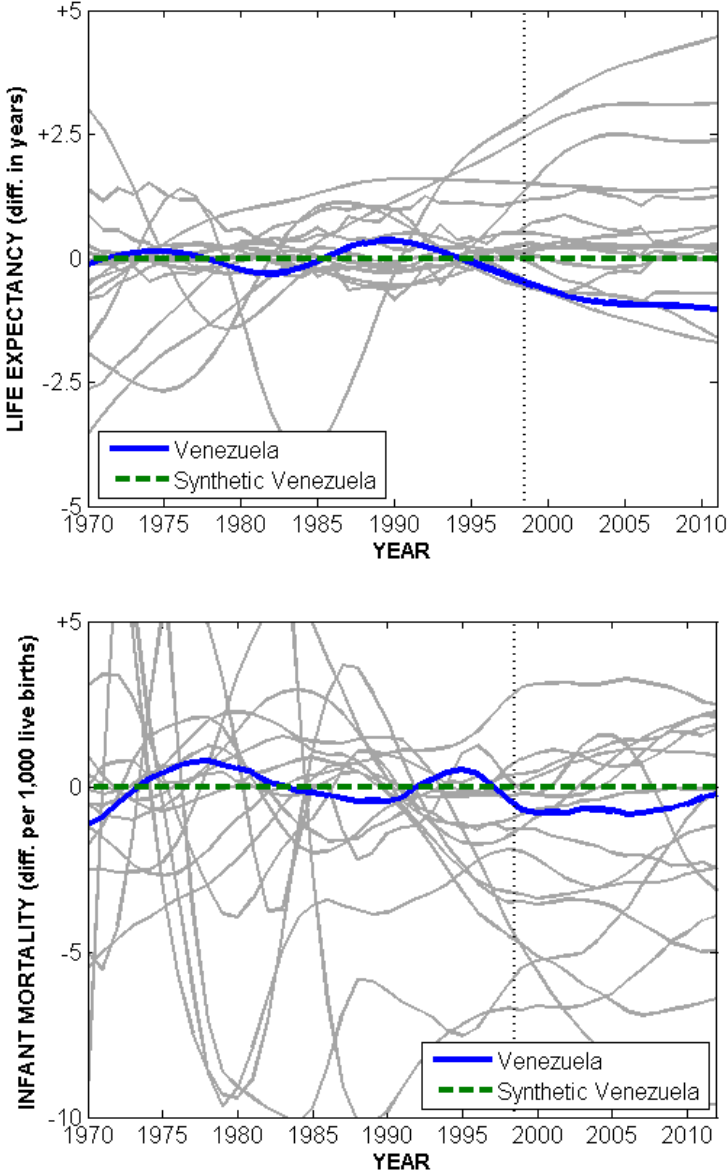
Figure 7: Health Outcomes



Top Graph Note: The solid line represents observed life expectancy in Venezuela, 1970-2012; the dashed line represents the synthetic control.

Bottom Graph Note: The solid line represents observed infant mortality rate in Venezuela, 1970-2012; the dashed line represents the synthetic control.

Figure 8: Health Outcomes Placebo Tests

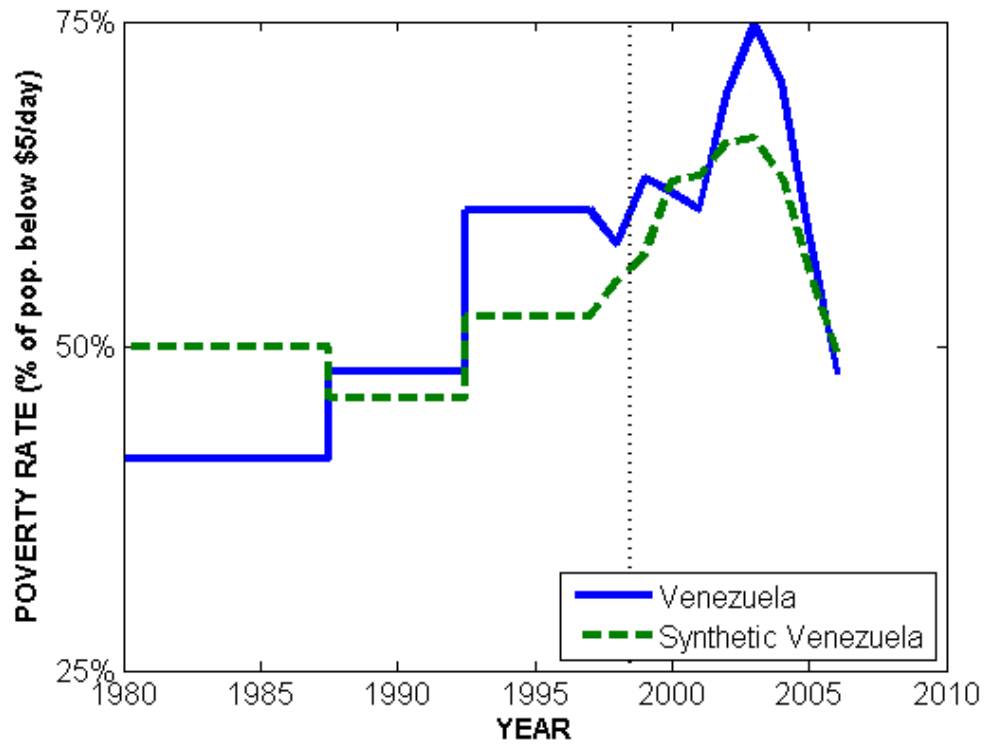


Top Graph Note: The bold line represents observed life expectancy in Venezuela, 1970-2012; the dashed line represents the synthetic control.

Bottom Graph Note: The bold line represents observed infant mortality rate in Venezuela, 1970-2012; the dashed line represents the synthetic control.

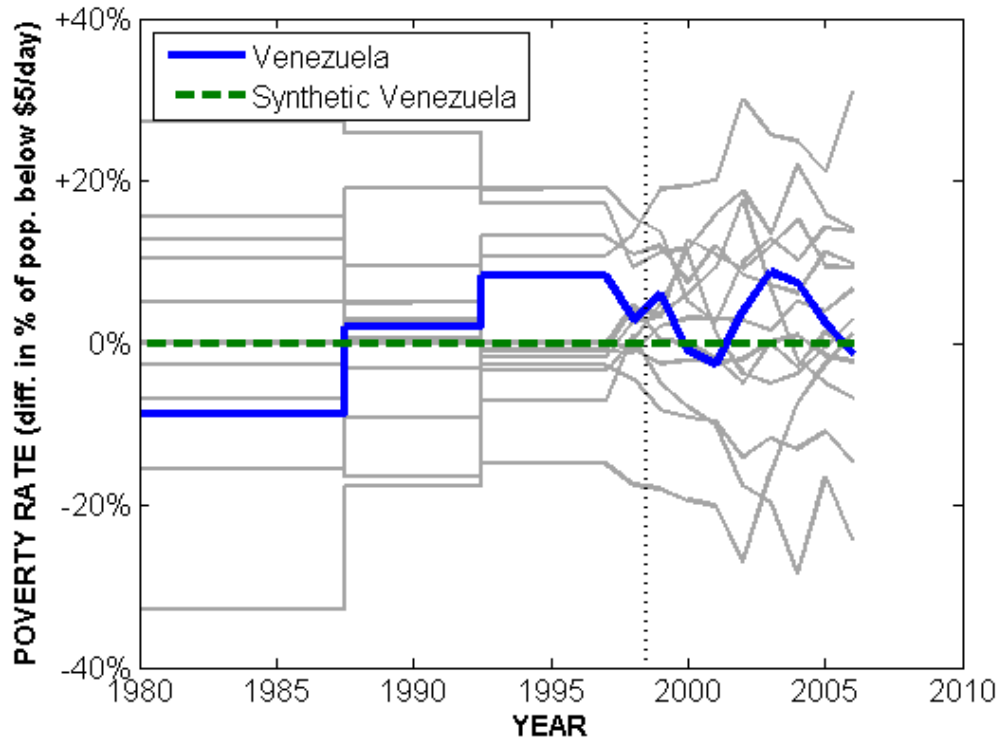
Both graphs show placebo tests (gray lines) for all other countries in the dataset, including those with poor pre-1999 fit.

Figure 9: Poverty Rate



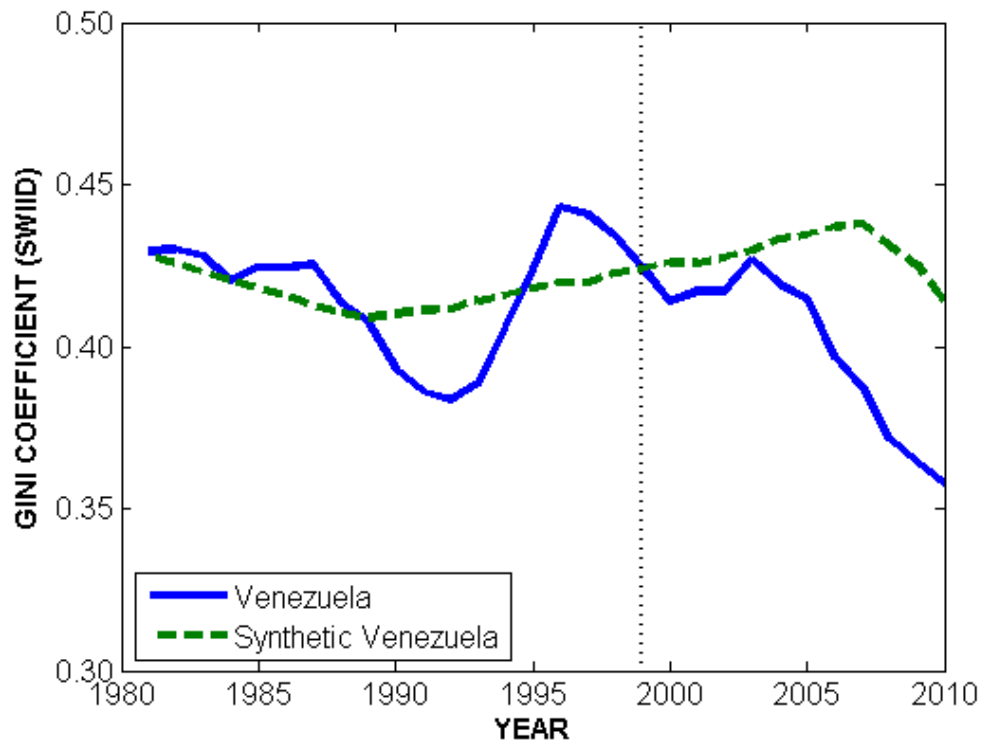
Note: The solid line represents observed poverty rate in Venezuela, 1980-2006; the dashed line represents the synthetic control. The vertical dotted line indicates the end of the pre-Chavez years (1980-1998).

Figure 10: Poverty Rate Placebo Tests



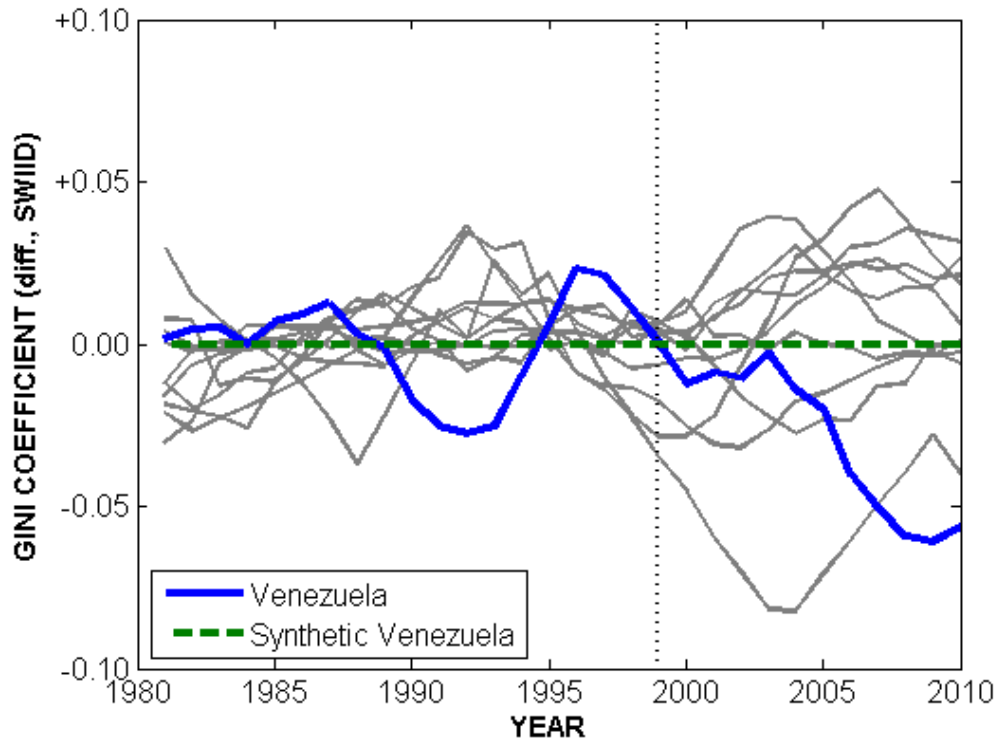
Note: The bold line represents the difference between the observed poverty rate in Venezuela, 1980-2006, and the synthetic control; the synthetic control (dashed line) is normalized to zero. Gray lines represent placebo tests: deviations from synthetic control for all other countries in the dataset. The graph shows all countries, including those with poor pre-1999 fit.

Figure 11: Inequality (SWIID Gini coefficient)



Note: The solid line represents observed Gini coefficient (SWIID) in Venezuela, 1981-2010; the dashed line represents the synthetic control.

Figure 12: Inequality Placebo Tests



Note: The bold line represents the difference between the observed Gini coefficient (SWIID) in Venezuela, 1981-2010, and the synthetic control; the synthetic control (dashed line) is normalized to zero. Gray lines represent placebo tests: deviations from synthetic control for the other countries in the dataset. This graph only include countries with pre-intervention root mean squared prediction error (RMSPE) less than 2.2212 (1.5 times that of Venezuela). It drops Costa Rica, Indonesia, Norway, and Peru.