

Do Nonpartisan Programmatic Policies Have Partisan Electoral Effects? Evidence from Two Large-Scale Experiments

Kosuke Imai, Harvard University

Gary King, Harvard University

Carlos Velasco Rivera, Institute for Advanced Study in Toulouse

A vast literature demonstrates that voters around the world who benefit from government discretionary spending increase their electoral support for the incumbent party. But, contrary to theories of political accountability, some suggest that voters also reward incumbent parties for “programmatic” spending, over which incumbents have no discretion, even when passed with support from all major parties. Why voters would attribute responsibility when none exists is unclear, as is why minority parties would support legislation that costs them votes. We study two prominent programmatic policies. For the first, we design and implement one of the largest randomized social experiments ever. For the second, we reanalyze studies that came to opposite conclusions, using a large-scale randomized experiment and a natural experiment. By improving statistical methods and correcting data errors, we show that evidence from all analyses of both policies is consistent: programmatic policies have no measurable effect on voter support for incumbents.

Political scientists in American and comparative politics have amassed considerable support for the theory that officeholders target *discretionary* government spending to gain votes, and voters reward them for doing so. As the magisterial literature review by Golden and Min (2013, 84) summarizes: “Studies overwhelmingly find that incumbent politicians are rewarded by voters for distributive allocations, and in particular for those that are clientelistic and from which recipients can be excluded.” Yet, perhaps paradoxically, some argue that voters react in the same way to *programmatic policies*, for which incumbents have little or no discretion in delivering benefits because citizens receive services on the basis of well-known, publicly stated rules (see Hicken 2011; Kitschelt and Wilkinson 2007; Stokes et al. 2013, for a definition). This is all the more puzzling in those situations when programmatic policies are passed with support from every major political party, including those who will be hurt electorally by their very action.

Although important efforts have been made to address the “programmatic incumbent support hypothesis” via qualitative argument (Cornelius 2004; Diaz-Cayeros, Estevez, and Magaloni 2009, 2016) and via the possibility of publication bias (Golden and Min 2013), neither the formal theoretical nor quantitative empirical literatures has offered a resolution. We resolve this puzzle by analyzing two large programmatic policies, each designed to reduce poverty and its effects—Seguro Popular de Salud (SPS) and Progresa. Studying SPS and Progresa together is also advantageous because the literature views the two as highly visible, highly impactful, programmatic policies passed and implemented in a manner close to a pure theoretical form (Camp 2013; Diaz-Cayeros et al. 2009, 2016).

We are fortunate in being able to study the electoral effects of these policies with two very large-scale randomized experiments and a natural experiment, which in total appears to be the best evidence ever brought to bear on this question.

Kosuke Imai is a professor of government and of statistics and Gary King is the Albert J. Weatherhead University Professor at Harvard University, Cambridge, MA 02138. Carlos Velasco Rivera is a visiting research fellow at the Institute for Advanced Study in Toulouse, 31080 Toulouse, France.

Data and supporting materials necessary to reproduce the numerical results in the article are available in the *JOP* Dataverse (<https://dataverse.harvard.edu/dataverse/jop>). An online appendix with supplementary material is available at <https://doi.org/10.1086/707059>.

The Journal of Politics, volume 82, number 2. Published online January 31, 2020. <https://doi.org/10.1086/707059>

© 2020 by the Southern Political Science Association. All rights reserved. 0022-3816/2020/8202-0022\$10.00

The first, which we designed and implemented originally for the purpose of evaluating SPS in Mexico, is one of the largest social experiments ever conducted and the largest randomized health policy experiment to date (King et al. 2007). The second, which includes the only randomized social experiment previously used to study the electoral effects of programmatic policies along with a novel natural experiment, left its authors' conclusions in disagreement (De La O 2013, 2015; Green 2006). Like the first experiment, it was also very large, originally designed as a policy evaluation and in Mexico.

We thus have two cases relatively close to the theoretical ideal for a programmatic policy. We are able to analyze each of these policies with unusually strong experimental frameworks. Our results resolve some disagreements in the literature and wind up not supporting the programmatic incumbent support hypothesis. We show that neither SPS nor Progresá has a causal effect on voter turnout or electoral support for the incumbent party.

The next two sections describe our two experiments. The section after those reinterprets the literature in light of our results. We also provide an appendix regarding some coding and analysis issues that led to prior disagreements in the literature and also an extensive online supplementary appendix with supporting technical information, measurement issues, statistical analyses, robustness checks, alternative measurement strategies, analyses of a formal theory, and other details. All data and replication information appear in Imai, King, and Velasco Rivera (2018).

EXPERIMENT 1: SEGURO POPULAR DE SALUD

Background

Although *seguro popular de salud* translates literally to “universal health insurance,” the Spanish word for “insurance” does not appear in the authorizing legislation, as it is a social welfare (income redistribution) program, not a self-sustaining insurance program. The program was designed to build or improve medical facilities and provide medical services, preventive care, pharmaceuticals, and financial health protection to the 50 million Mexicans with no regular access to health care (constituting about half the population of the country). SPS was aimed at those with low incomes; its main purpose is to reduce the devastating effects of catastrophic health expenditures, when, because of illness or injury, greater than 30% of a family's annual disposable income is spent on health care in one year. Before SPS, about 10% of the poor had catastrophic health expenditures each year. SPS was designed to eventually spend an additional 1 percentage point in gross domestic product; in 2005, expenditures totaled a substantial \$795.5 million. As it turned out, SPS was the most visible accomplishment of the Vicente Fox Quesada administration.

By all accounts, it was designed, passed, and implemented in a nonpartisan, programmatic fashion.¹

Because of Mexico's term limits, President Fox, along with Health Minister Julio Frenk Mora, decided they needed a way to convince whoever would succeed them to keep SPS in place. How one democratically elected government can “tie the hands” of the next democratically elected government is a classic question of normative political theory (Klarman 1997; Posner and Vermeule 2002; Sterk 2003; Thompson 2005), formal theory (Alesina and Tabellini 1990), and empirical political science (Franzese 2002). In this case, Fox and Frenk's approach was to commission an independent scientific evaluation by a team we led at Harvard's Institute for Quantitative Social Science (see <http://j.mp/ExpMex>). The idea was that if the evaluation favored SPS, it would at least have been more difficult for the next government to eliminate. And to make the evaluation worthwhile, we committed publicly and in print to say so if the evaluation was not favorable in any way.

The Mexican government signed legal contracts that gave us free rein to design and implement whatever evaluation, and spend whatever funds, we judged appropriate. The government gave us unfettered access to government officials, the ability to influence how SPS was implemented so we could more easily evaluate it, and convening power to speak with the numerous local officials across the country in charge of implementation. We retained the legal right to publish without prior review.

We developed a new experimental research design robust to the interventions by politicians who regularly—indeed usually—derail large public policy experiments, as they choose to be more attentive to the short-term desires of their constituents than any longer-term benefit of scientific evaluation (King et al. 2007). Details of the experiment include discussions of the background and design, published before any data analysis (King et al. 2007); novel statistical methods we developed for this design (Imai, King, and Nall 2009b); empirical results (King et al. 2009a); and publicly available replication data (Imai, King, and Nall 2009a; King et al. 2009b).

1. Julio Frenk, one of SPS's main architects said, “From the beginning I suggested that this [SPS] had to be a proposal from all parties, in other words, that it was not only a PAN, PRI, or PRD project, but a project for the country, for which everyone could claim credit. In a conversation with President Fox, after introducing the bill to Senate, I told him: ‘If we want the reform to move forward, it is extremely important that the project is handled as a shared project and that we give credit to everyone.’ The president agreed, and promised that credit would be given to everyone, which is what happened” (Ortiz 2006, 81, our translation). Whether we judge by public statements from minority parties or from the votes on legislation, there exists near consensus on this point.

We began by defining 12,284 *health clusters*, which are new, continuous geographic areas that tile all of Mexico’s 31 states in which fall one health clinic (or potential clinic that we could decide would be built) along with its catchment area (defined as less than a day’s travel time to the clinic, using locally available methods of transportation rather than “as the crow flies”). We recruited 13 of Mexico’s 31 states to participate, including 7,078 health clusters. We matched these health clusters in pairs based on background characteristics and then selected 74 pairs (based on the closeness of the match, likelihood of compliance with the experiment, and necessary political criteria). The experiment thus included 534,457 research subjects, in 118,569 households, within 1,380 localities, nested within 148 health clusters.

Within each of the 74 matched pairs, we randomly selected one health cluster to receive treatment and the other, as the control, to receive no change. Treatment included new (or upgraded) hospitals and other medical facilities, doctors, access to medicines and other medical interventions, advertising campaigns to encourage affiliation with the program, individual insurance, and funds to pay for it all. The treatment was applied August–September 2005, coincident with a baseline survey of 32,515 respondents randomly selected from 50 of the health cluster pairs; the outcome was measured the same way 10 months later, July–August 2006. We measured an extensive array of variables including individual opinions, attitudes, health status, and financial spending; household level variables, such as assets, wealth, demographics, and others; and physical health measures via three separate blood draws for each person.

The 2006 Mexican presidential election was held on July 2, coinciding with the start of the follow-up survey, which happens to be perfect timing for studying the effects of SPS on the election (see fig. 4 in the supplementary appendix). To merge federal election results reported by one governmental office, with census data reported by another, we define a new unit of analysis for this study within the confines of the experimental randomization. This is the *precinct cluster*, which we define as the largest possible geographic subset of a single health cluster in the SPS experiment for which we can accurately merge all relevant electoral and census information (for details, see the supplementary appendix, sec. 2). When insufficient information is available to identify a precinct cluster, we retained all the benefits of the original experimental design by removing the affected matched pair of health cluster (as suggested by Imai et al. 2009b; see also the discussion in Imai and Jiang 2018).² This left us in the end with 57 matched

2. This decision is not affected by which unit received the treatment within a given pair.

pairs (47 rural and 10 urban) out of the original 74 and all the benefits of matched pair randomization. Table 1 and figure 7 in the supplementary appendix give full details and descriptive statistics of the sample we use and to which our inferences apply. (As another robustness check, we repeated our analyses with all available precinct clusters even when no match was available, trading off more model dependence for less inefficiency, and found no substantive change in the results we present below. See tables 2 and 3 in the supplementary appendix.)

Especially useful for the current article is that our evaluation indicated that SPS had a massive effect on its intended outcome variable of financial assistance to the poor—a fact that is all the more impressive because government programs designed to help the poor in most countries typically have no measurable impact on the poor (Gwatkin, Wagstaff, and Yazbeck 2005). In only 10 months, SPS eliminated about a third of the catastrophic health expenditure problem among the poor in Mexico and about 60% among experimental compliers (those who would affiliate to SPS when in the treatment group and not when in the control group). Perhaps even more important for our purposes, the people of Mexico clearly thought the program would help them: fully 44% of those eligible in treatment areas enrolled in the program in the first month, which for each family involved taking a trip that could last as long as a full day to formally affiliate to SPS. Moreover, those who enrolled liked the program a great deal: 69% of those enrolled rated the quality of health services as good or very good, and 97% planned to enroll again in the follow-up period after the experiment. Some aspects of the program did not have the intended effects, as reported in King et al. (2009a), but for our purposes this large financial impact makes this an unusually strong test of the electoral effects of programmatic policies.

Results

We present two sets of outcome variables for the SPS experiment based on analyses of actual electoral results (using aggregation procedures described in the supplementary appendix, sec. 2) and on retrospective survey evaluations (which require no aggregation), respectively. The statistical methods used in both cases are fully nonparametric, enabling us to reap the benefits from a randomized experiment without making modeling assumptions (for a full description, see Imai et al. 2009b).

First, we estimate the total causal effect of SPS on electoral outcomes in figure 1. (This analysis and all other types of causal effects on electoral outcomes are based on the 57 precinct matched cluster pairs, resulting from mapping the 74 evaluation health cluster pairs into precincts.) A point estimate




Figure 1. Intention-to-treat point estimates and 95% confidence intervals for the total causal effect of SPS on voter turnout (dashed line) and the incumbent (PAN) vote share (solid line), overall (far left of left panel) and in three partitions of the data, including urban/rural (left), income quartile (middle), and asset quartile (right).

(dot) and 95% confidence interval (vertical line) appear for each causal effect, with a horizontal line at zero, indicating no effect. If the literature’s hypothesis is correct that nonpartisan programmatic policies increase support for incumbents, the points and confidence intervals would appear above the horizontal line. Instead, all the confidence intervals cross the zero, no effect, line, and thus none confirm the hypothesis.

The estimated total causal effect of SPS on the percentage voting for the incumbent party (National Action Party, PAN) is the first result, at the far left of figure 1 (solid vertical line); the estimated effect on federal voter turnout appears next to it (dashed vertical line). In both cases, the estimated total causal effect is not distinguishable from zero. Indeed, the confidence intervals are quite narrow for the estimated overall effects, indicating that the experiment is well powered for this hypothesis, and so we should have high confidence that the effect is negligible if not exactly zero.

To further search for possible support for the hypothesis, we partition our sample in three different ways to examine subgroup effects. We give the total causal effect within rural and within urban precinct clusters at the center and right, respectively, of the left panel of figure 1. All results are statistically indistinguishable from zero. The right two figure panels estimate total causal effects by quartile of the proportion of individuals in the first two income deciles living in a health cluster (center) and household assets (right). These may be especially relevant since households that report less income will pay less (or nothing if in the first two income deciles) for SPS services. These results can also be viewed as predicting compliance with receiving the experimental treatment, since less poor communities were expected to (and actually did) sign up for SPS services less often. This analysis follows Calvo and Murillo (2004), whose account suggests that the largest electoral effect of policies could be found among

individuals who benefit most from a given policy. Nevertheless, the estimated causal effect of SPS on voter turnout and on incumbent party vote within every one of these segments of the public is not distinguishable from zero.

Second, we avoid all issues involved in aggregating precincts to precinct clusters (discussed in the supplementary appendix, sec. 2) by estimating the total causal effect of SPS on individual survey evaluations. (This analysis and all causal estimates on individual evaluations are based on the 50 matched health clusters participating in the survey.) Recall that the beginning of our follow-up survey coincides with the election, and hence we are measuring voters’ opinion right after the election. We measure how well respondents thought the country was doing on economic, political, and social issues, compared to five years before (we provide descriptive statistics in fig. 8 of our supplementary appendix).³ This analysis also allows us to test the channels through which retrospective voting may work, one of the mechanisms hypothesized for why incumbents may benefit from programmatic policies (Pop-Eleches and Pop-Eleches 2012).

These results appear in figure 2 (in a format parallel to fig. 1). For the overall effect (far left) and for each of the three domain areas (economic as a solid line, political as dashed, and social as dotted), the effect is estimated at near zero with a confidence interval that overlaps zero. We find the same essentially zero effect of SPS, and failure of the programmatic policy hypothesis, for individual effects within different subgroups, including rural, urban (left panel), income quartile (center panel), and asset quartile (right panel).

3. The exact wording of the retrospective evaluation question in the survey was as follows: “Compared to five years ago, do you think Mexico is better, the same, or worse today economically/politically/socially?”




Figure 2. Intention-to-treat estimates and 95% confidence intervals of the total causal effect of SPS on economic (solid line), political (dashed line), and social (dotted line) retrospective evaluations of whether the country is doing better today than it was five years ago ($n = 32,515$ individuals in 50 matched health cluster pairs). Results are reported for all respondents and by urban/rural breakdown (left), income quartile (center), and asset quartile (right).

In our supplementary appendix, we also present a wide array of other analyses of the same data. For example, we reran the analysis from figures 1 and 2 with different coding rules and statistical techniques. These and many other analyses reveal no noticeable effect. To be specific, in the supplementary appendix, figure 9 analyzes PAN votes as a share of registered and eligible voters, as alternative measures of incumbent support. Figure 10 measures turnout with total votes cast as a share of eligible voters. Tables 2 and 3 report estimates for the impact of SPS on incumbent support and turnout when we break up precinct clusters pairs and analyze the data in a regression framework. Tables 4 and 5 do the same but for the sample of rural precinct clusters, for which we know the population measurement error is minimal. Figure 11 shows that the reallocation of opposition resources following the introduction of SPS does not explain the policy's null effect. Figure 12 reports intention-to-treat estimates by the share of the evaluation population in precinct clusters to address concerns of attenuation bias in rural areas. Finally, figure 13 reports estimates for the impact of SPS on individual retrospective economic, political, and social evaluations of the country when we control for the responses of the same people in a survey we conducted at baseline. Overall, the results are unambiguous: the highly successful SPS programmatic policy had little effect on turnout or the vote.

EXPERIMENT 2: PROGRESA

Background

We are fortunate to be in the possibly unprecedented position for political science of having a second large-scale randomized experiment, along with a natural experiment, to study the same substantive question. This analysis evaluates one of the largest poverty-alleviation programs in Mexico. The program consists of nutritional, educational, and health

components. The policy's key feature consists of cash transfers that eligible households receive on the condition that they attend regular health checkups and children enroll and attend school. Program benefits vary according to household composition. The average level of benefits is about US\$35 per month, which represents about 25% of income in poor rural households (Levy 2006, 23).

The origins of the program date back to the mid-1990s, when the country experienced one of the worst economic crises in its history. The government had previously relied on a myriad of food subsidies to alleviate poverty. However, government officials in the administration of then-president Zedillo concluded that subsidies benefited urban centers at the expense of the poorest rural areas in the country and were regressive and too costly to administer (Levy 2006, chap. 2). The motivation underlying the new administration's alternative approach was a recognition that in order to break the cycle of poverty, one had to recognize the relationship between the educational, health, and nutritional components of human capital.

The architects of the program were committed to the eradication of poverty in the country, and as a result they implemented a design that would increase its long-term viability. Progresa defined a target population and operated under clear, programmatic rules. In the eyes of the policy makers, this would ensure the political neutrality of the policy.⁴ Political scientists have also arrived at the same conclusion

4. Santiago Levy, one of Progresa's main architects, emphasized the nonpartisan goal of the program: "Congress's role in *Progresa-Oportunidades* has also contributed to its continuity in yet another way: it has established strong legal provisions against the 'political' use of the program. More particularly, it has sought to separate the program from the public image of the president and to provide information directly to beneficiaries about the

about the nonpartisan nature of Progresá.⁵ Analogous to the experience with SPS, the government hired the International Food Policy Research Institute as a trusted third party to evaluate Progresá and bolster the program's credibility (Levy 2006, 43).

The main evaluation of Progresá exploited the phasing of the policy across the country's localities. The sample consists of 506 rural localities in the country distributed across seven states, with 320 villages drawn from a population of localities eligible to receive the program by November 1997. Eligible localities were then assigned to the treatment group. The remaining 186 villages were randomly drawn from populations that would receive the policy in one of the later phases (November–December 1999 and March–April of 2000) as the control group. Although Progresá's design was "completely randomized," as distinct from SPS's more powerful "matched pair randomized" design, both have the advantages of large-scale experimental randomization. Then, the government carried out several surveys in each of these villages, first to determine household eligibility and then to measure outcomes over two years. Behrman and Todd (1999) find that treatment and control villages are fairly similar across a large battery of socioeconomic indicators.

The results of the evaluation based on this sample of villages show that Progresá increases school enrollment, improves the health and nutrition of children and adults, and increases household consumption (largely on food; Skoufias 2005, chap. 5). However, as first discussed in Green (2006), the evaluation potentially poses a challenge for estimating the effects on electoral variables because the 2000 Mexican presidential election was held on July 2. This means that for the purposes of studying the electoral effects of Progresá, the treatment group is defined as having received the pro-

gram for 31–32 months before the election, whereas the control group is defined as receiving the program for only 3–8 months (see the timeline in fig. 5 of our supplementary appendix). This is not as clean a test of the programmatic hypothesis as with our SPS experiment, since those who received the program more recently in the control group may be as or more grateful as those who received it earlier, although the distinction between the two groups remains unambiguous.

Fortunately, the way in which Progresá was rolled out across the country offers a different identifying assumption in the form of a natural experiment. Government officials relied on a poverty index to determine which communities (*localidades*) would be enrolled in the program.⁶ Authorities first enrolled in the program all localities reporting high and very high levels of poverty (and meeting other criteria, such as having access to a health center and educational facilities and having a threshold population level). Once authorities completed this phase, they proceeded to progressively incorporate localities in the poorest quintile, the second poorest quintile, and so on, among the set of localities reporting medium levels of poverty (Green 2006, 67). As figure 5 in the next section shows, this procedure generated two large exogenous discontinuities in the proportion of communities (and households) enrolled in the program.⁷ Following Green (2006), we exploit these discontinuities to estimate the impact of Progresá under a regression discontinuity design (RDD).

As with the SPS experiment, analyzing federal election outcomes in Mexico requires a procedure for merging or matching the boundaries of electoral precincts with often overlapping census geography, as the two are generated by different administrative offices that typically do not coordinate. The same issue exists in almost all analyses of electoral data around the world, but

nature of the benefits that they receive, their rights, and their obligations. . . . These factors, along with the program's positive results, have contributed to the program's transit through three shifts in the composition of the House of Representatives since 1998" (2006, 107–8).

5. Díaz-Cayeros et al. (2016) conducted interviews across communities in Oaxaca—a state with a long history of clientelism—about the experience of voters regarding the provision of Progresá. In their interviews, one voter noted that "one can be PANísta, PRIísta or PRDísta and still receive benefits from Oportunidades. . . . Before you had to be with the PRI to get anything from the government" (159). Another respondent noted that "the governor controls everything in Oaxaca. However, here you can be PANísta, be with the governor [back then from the PRI], and still get benefits from Oportunidades" (6). A final interviewee asserted that "although sometimes people who do not really need it get Oportunidades, it is less corrupt because benefits arrive regardless of which party you like" (6). On the basis of evidence like this, Díaz-Cayeros et al. conclude: "It is clear from our interviews in Oaxacan villages that the poor perceive important differences between the CCT and other social programs, and that they consider it an entitlement rather than a political handout that comes and goes with the election cycle" (159).

6. The index (*índice de marginación*) classified 105,749 localities in the country across the following five categories of poverty: very low, low, moderate, high, and very high. To distribute the localities across the five categories, officials used factor analysis to create a latent measure of poverty. Once this measure was obtained, authorities then implemented an optimal classification algorithm.

7. Green (2006, 74) reports that the exact cutoffs are located at the value of the index separating localities reporting low and moderate levels of poverty (threshold 1) and at the level separating localities in the three poorest quintiles from those in the two richest quintiles among the set of localities reporting moderate levels of poverty (threshold 2). However, as shown in supplementary app. fig. 51, we find that threshold 2 is slightly lower than the value separating the quintiles of interest. For our estimation, we set threshold 2 equal to -0.96 (instead of -0.932), which corresponds to the largest effect of the encouragement on Progresá enrollment. Figures 51 and 52 show, however, that our main estimates are robust to different values of threshold 2. Finally, we note that thresholds 1 and 2 are not deterministic because government authorities, in addition to the poverty index, took into account access to health and educational facilities for program enrollment.

the method of dealing with it is crucial. As we explain below, errors in this merging process for Progresá explain certain prior results in the literature.

Results

The literature analyzing Progresá is divided over whether the data support the programmatic incumbent support hypothesis. The first analysis in this literature, based on the natural experiment, finds no effect on either incumbent support or voter turnout—with estimated effects close to zero and small confidence intervals (Green 2006). In contrast, the second analysis, based on the randomized design, reports strong positive effects (De La O 2013, 2015). We show here that the reason for this discrepancy is not related to the differing identification strategy but rather faulty data-merging procedures (affecting about 70% of the observations) that happened to induce misleading results for the randomized experiment and not the natural experiment. Unconventional data analysis choices in the randomized experiment also contributed to the incorrect conclusions. Correcting either (or both) in the analysis of the randomized experiment generates results that mirror those from the natural experiment, which are also consistent with our analysis of SPS above—both indicating little or no effect of programmatic policies on voter turnout or incumbent support.

The appendix reveals these errors and shows how to correct them. The rest of this section replicates the original results reported in Green (2006) and De La O (2013, 2015). For the first analysis we exactly replicate the results using data that were incorrectly merged, with incumbent voting and turnout measured as a function of all people, and then with corrected data based on accurate geographic information system (GIS) coordinates along with more appropriate statistical techniques. We also add a new data source, with outcome variables based on voting as a function solely of those officially registered. This alternative coding has two advantages. First, no merging is necessary, and so no corrections are needed. And second, we offer a much stronger test of the hypothesis by excluding those who cannot vote and thus have zero causal effects (such as those underage, not citizens, not registered) from the denominator of the outcome variables. For the second analysis, we report results relying on the correctly merged data using GIS coordinates and show that the original conclusions of the study hold.⁸ (In our supplementary appendix we provide numerous other analyses,

8. We did not have access to the exact sample analyzed in Tina Green's unpublished dissertation (2006) and so are unable to report numerical estimates obtained in her analysis. However, this information would not change our conclusions or her's.

tests, and evidence, all of which yield the same conclusions as that offered here.)

We begin our analysis with figure 3, which replicates the regression estimate and 95% confidence intervals from Green (2006) and De La O (2013) for the total causal effect of Progresá on turnout (*left panel*) and incumbent vote (*right panel*). Results for De La O (2013) are based on the sample as originally coded with errors and including those who cannot vote in the denominator of the outcomes (*squares*). We also analyze official measures of the outcomes (turnout and incumbent support) with the incorrectly merged sample (*diamonds*) and with a corrected GIS sample, without merging problems (*circles*). Results for Green (2006) are based on official outcomes for the sample of precincts, each of which has only one village (*localidad*) obtained with the correct GIS procedure for merging (*triangles*). The horizontal line marking no effect appears at zero.

Reading the left panel in figure 3 from left to right, the different specifications we tried for De La O (2013) include the linear regression in the original article and book; a simple difference in means; a matching estimator;⁹ a regression controlling for log population; a regression with lag turnout on the same scale as the outcome; and a regression, under the original specification, after removing the two observations with the highest leverage. For Green (2006), we report the original pooled sharp RDD results.¹⁰ The panel on the right repeats all the analyses for incumbent (Institutional Revolutionary Party, PRI) vote share, including for Green (2006) the original estimates for the total effect of Progresá on PRI support in the Proportional Representation (PR) Senate election (*triangle*) and in the presidential election (*inverted triangle*).

The results in figure 3 exactly replicate results in De La O (2013, 2015), with positive point estimates for turnout and vote share for the incumbent party and a 95% confidence interval that excludes 0 for vote share but is insignificant for turnout.¹¹ Using the original variable (with errors uncorrected) reveals the same basic results, even using a simple difference-in-means estimator. However, once we use any of the four alternative approaches, each of which controls for

9. We did coarsened exact matching, adjusting the coarsening to deal with the presence of high leverage observations among the pretreatment covariate. The distributions of covariates before and after matching are reported in figs. 18, 19, 23, and 24 in the supplementary appendix, with full information in our replication data set.

10. Figures 39 and 40 in the supplementary appendix report additional results across the different thresholds for locality enrollment to Progresá, implementing different kernels (uniform and triangular) and employing different RD estimators (standard and bias correcting).

11. Figures 16 and 17 in the supplementary appendix display all the point estimates reported in this section but with 90% confidence intervals.




Figure 3. Intention-to-treat estimates of Progresa effect on turnout and incumbent party vote. *Left*, point estimates and 95% confidence intervals for the total causal effect of Progresa on turnout in the 2000 presidential election as originally, and incorrectly, measured in the De La O (2013) sample (squares), for official turnout among registered voters in the same sample (diamonds), and for official turnout among registered voters in the correct GIS sample (circles). The panel also replicates Green's (2006) total causal effect of Progresa on turnout in the sample of precincts with only one village under a sharp RDD design (triangle). *Right*, same analyses for incumbent (PRI) vote share and the effect of Progresa under sharp RDD on both PRI support in the 2000 PR Senate election (triangle), as in Green (2006), and in the presidential election (inverted triangle). Every estimate is indistinguishable from zero, except when using the flawed original measure from De La O (2013) and without controls (first two lines with squares representing point estimates in the right panel).

the large imbalance induced by the data errors, the positive effects vanish with no statistically significant evidence for the effect of Progresa on either turnout or vote share. Moreover, rerunning any of the six analyses, while dropping the original incorrectly coded variable and switching to official registration data (which has no possibility of data-merging errors and a higher probability of revealing an effect if present), reveals no evidence for the effect of this nonpartisan programmatic policy on either partisan outcome, regardless of how the data are analyzed. Moreover, with the clean registration data, the confidence intervals are much narrower, and all 12 include zero as a causal effect. Our reanalysis of Green (2006) strongly confirm the substantive conclusions reported in that study, namely, that Progresa did not have a substantial positive impact on either turnout or incumbent support.

Figure 4 repeats the same analyses, with the same robustness checks, for the instrumental variable analysis estimate of the causal effects in De La O (2013, 2015) and Green (2006).¹² The results here tell essentially the same story, with no statistically significant effect of nonpartisan programmatic policies on voter turnout or vote for the incumbent party. Although again, only the official turnout and vote figures (all

point estimates except squares) offer valid causal estimates, and these are not statistically different from zero.

Finally, to study the possibility that our estimated null electoral effect of Progresa is due to bias resulting from unaccounted heterogeneous treatment effects, we allow our estimated effect to vary by the pretreatment population in localities that participated in the evaluation (as a percentage of the population of all localities in a precinct). We first estimate the impact on the same outcomes as before, then conduct a placebo test, and finally implement a difference-in-difference design. Figure 29 in the supplementary appendix apparently suggests that Progresa increased PRI support in precincts with a higher percentage of experimental population. However, a placebo analysis shows that this result is not to be trusted: figure 31 in the supplementary appendix has almost the same apparently positive effect of Progresa on PRI support when the effect is known to be zero—the 1997 legislative elections, which occurred four months before localities were enrolled. (We do not use the 1994 presidential elections because of the 1996 redistricting that changed the boundaries of numerous precincts across the country and induced unacceptable levels of measurement error for pretreatment and outcome variables.) To address the imbalance between treatment and control precincts that the placebo test reveals, we implement a difference-in-difference model for the 2000 legislative elections. The findings, reported in figure 32 of the supplementary appendix, show that Progresa did not increase incumbent support. This analysis, along with numerous

12. The fuzzy RDD estimates reported in fig. 4 are based on a regression specification in which the treatment is a binary indicator for whether a locality was enrolled in Progresa as originally defined in Green (2006). Figures 43 and 44 in the supplementary appendix report results when the treatment is instead the proportion of families in a locality enrolled in Progresa.




Figure 4. Complier average treatment estimates of Progresa effect on turnout (*left*) and incumbent party vote (*right*). Directly parallel to figure 3, this figure replicates the instrumental variable estimation from De La O (2013) and the fuzzy RD design from Green (2006). Every estimate is indistinguishable from zero, except when using the wrong measure without controls (first two lines with squares representing point estimates in the right panel).

alternative specifications also reported in the supplementary appendix, leads to the same conclusion as with SPS: Progresa has little or no effect on either turnout or incumbent vote.

Robustness of our substantive conclusion

Two important shortcomings of the experimental evaluation of Progresa were first noted in the doctoral dissertation of Green (2006, n. 26) and mostly ignored thereafter. First, all villages (localities) in the evaluation study, including those in the control group, had received the treatment by the time of the 2000 election. Second, some treated localities share precincts with other localities that were not part of the evaluation, leading to a small number of households in the experiment who enrolled in Progresa. These two shortcomings of the randomized experiment may have contributed to the evidence that the program had little impact on election results.

Below, we address these two shortcomings of the randomized experiment by following Green (2006) and employing an alternative identification strategy based on an RDD. To do this, we exploit the arbitrary cutoffs government officials used to phase in localities to Progresa. These results strongly confirm those of Green (2006) and our substantive conclusion that Progresa had little impact on electoral results.

Although the original data from Green's (2006) dissertation were not available to us, we requested and received directly from the Mexican government information about the number of families incorporated to Progresa across all localities in the country for each of the expansion phases of the program. We thus focus on the 105,749 localities reporting a value of the poverty index used to enroll localities in the program. For these localities we create two versions of the treat-

ment: (1) an indicator variable (Progresa) that takes the value 1 if at least one family was enrolled in the program by the eleventh phase of program expansion (the last one before the year 2000 election) and (2) the proportion of families receiving Progresa within a given locality by the same program expansion phase.¹³ Relying on GIS, we then merge these localities with the set of precincts in the 2000 election. Following Green (2006), we examine precincts containing only one locality to avoid all issues of aggregation and merging.¹⁴ This process left us with a total of 7,865 precincts for our RDD analysis.¹⁵

Finally, the Mexican authorities relied on a poverty index to determine which localities were given priority to be enrolled in Progresa (Skoufias, Davis, and Behrman 1999, sec. 3).

13. Because we do not have data on the total number of families per locality, we use instead the total number of inhabited households as the denominator to compute the proportion of Progresa families across localities. This results in 231 localities reporting a value greater than 1. The likely reason for these values, according to the National Institute of Statistics and Geography's documentation, is that two or more families may share a household.

14. Chiapas and Mexico City are excluded from the sample. We exclude Chiapas because over 1,500 localities lacked geographic coordinates in the 1995 population count. We exclude Mexico City because it was not incorporated into Progresa by the program's eleventh phase of expansion.

15. The total number of precincts analyzed in Green (2006) is 3,379. The reason for the discrepancy between our sample and the sample analyzed by Green is the name-matching procedure the latter study used to merge localities with precincts. As we have shown, this procedure is unreliable, and in this particular case it may have led to an underestimate of the number of precincts with only one locality. This may have happened because a large number of precincts reporting two (or more) localities in the files from the electoral authority may in fact contain only one according to the way census authorities aggregate population at the local level.

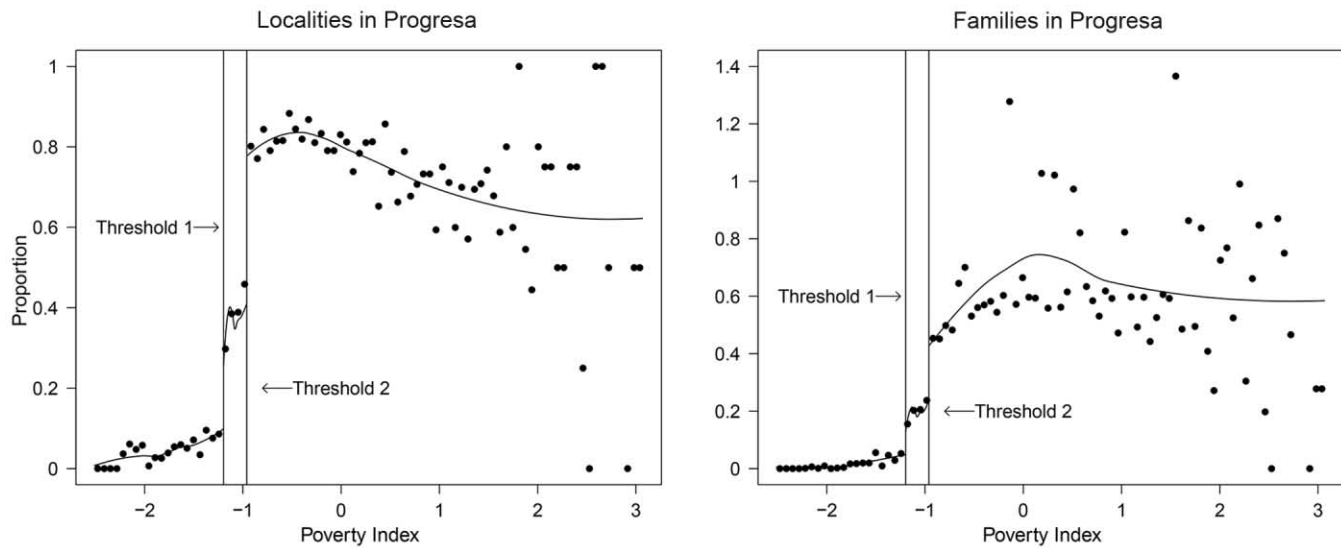


Figure 5. Proportion of localities phased into Progresa (left) and the proportion of Progresa-beneficiary families per locality (right) as a function of the census poverty index. Shows two large discontinuities in the proportion of localities incorporated to Progresa and in the proportion of Progresa-beneficiary families at the cutoffs the government used to phase in localities to the antipoverty program.

Authorities first enrolled in the program all localities reporting high levels of poverty (and fulfilling other criteria, such as having access to a health center, educational facilities, and a certain population level). Once authorities completed this phase, they proceeded to progressively incorporate localities in the poorest quintile, the second poorest quintile, and so on, among the set of localities reporting medium levels of poverty (Green 2006, 67).

Figure 5 shows that this procedure generated two large discontinuities in the proportion of communities and households enrolled in the program. At threshold 1, located at the value of the index separating localities reporting low and mod-

erate levels of poverty, we can see a substantial 28 percentage point increase in the proportion of localities incorporated to Progresa and 15 percentage points for households. Similarly, at threshold 2, we find a 40 percentage point increase the proportion of localities enrolled in Progresa and 23 percentage points for households. In section 3.3 of the supplementary appendix, we examine several pretreatment covariates including previous election results and show that there is no such discontinuity in these variables, giving strong support for the validity of the RDD analysis.

Figure 6 presents reduced-form intention-to-treat effects of Progresa on PRI vote share and turnout in the 2000 presidential




Figure 6. Intention-to-treat effects of Progresa on average official PRI vote share (left), PRI vote as a share of registered voters (center), and official turnout (right) in the 2000 presidential election as a function of the poverty index. There is no discernible discontinuity in the outcomes at either of the government cutoffs used to phase in localities to Progresa, indicating that the program had no effect on electoral outcomes.

election. The figure reveals the absence of any discontinuity at either of the two thresholds, indicating that Progresá had no discernible effect on turnout or PRI vote share. Indeed, when we use these data to compute the Sharp and Fuzzy RDD estimates, we find that the treatment effect of Progresá is almost exactly zero (last columns in the right panel of figs. 3 and 4, respectively). In section 3.3 of the supplementary appendix, we consider other election results examined by Green (2006)—PR Senate and simple plurality deputies elections. We find again that Progresá had little impact on any of these elections.

PRIOR RESEARCH

Our results, which clearly reject the programmatic incumbent support hypothesis, differ at least superficially from the largely positive conclusions reported in the literature (Labonne 2013; Larreguy, Marshall, and Trucco 2015; Manacorda, Miguel, and Vigorito 2011; Pop-Eleches and Pop-Eleches 2012; Zucco 2013). After some study, we conjecture that the key factor leading to these divergent results is the nature of the politics under which each policy was passed and the ways in which the policies were implemented (Layton and Smith 2015). Whereas SPS and Progresá are essentially pure forms of programmatic policies, most policies studied in prior research are partly programmatic and partly clientelistic. To be specific, programmatic policies in general, including those we study and all others in the literature, (a) impose objective rules that give incumbents no discretion over implementation. This condition implies that programmatic policies should (b) lead voters to expect to receive the same services regardless of the party in office. In addition, SPS and Progresá are distinctive in that they also (c) were passed with broad support in the legislature from all major parties and, perhaps as a consequence, (d) did not lead any single party to attempt to claim exclusive credit for its implementation. Our hypothesis is that conditions *c* and *d* play a major role in explaining the differing results.

It is true that studies of other programmatic policies have been almost exclusively observational. In contrast, randomized experiments enable researchers to estimate causal effects without risky modeling assumptions that are necessary in observational studies. Yet, experimental studies share with observational studies the problem of not being automatically representative outside the (political) context in which they are conducted (Imai, King, and Stuart 2008). This makes understanding the political context essential for any general understanding of the programmatic incumbent support hypothesis, experimental or observational. We thus now discuss the differences in the politics under which policies are adopted and implemented in relatively pure form and under

a mixed programmatic-clientelistic fashion. We first consider differences in the politics of adoption and implementation, then discuss broader theoretical explanations, and finally make suggestions about future research.

Politics of policy adoption

Clientelism has historically been the norm across developing countries because it yields large electoral benefits to incumbents (Golden and Min 2013). Incumbent chief executives generally prefer clientelistic policies because they are likely to yield large electoral benefits for themselves and their parties (Stokes 2005). However, adopting clientelistic policies is only possible with unified partisan control of government (De La O 2015). Under divided partisan control, the only policies that have a chance of passing are those supported by all parties (or at least all veto players). This explains why programmatic policies—SPS, Progresá, and every other major programmatic antipoverty program in the developed world—have only arisen under divided control. For example, in ensuring congressional approval for SPS and Progresá, policy makers took great pains to ensure no political actors could claim credit distinct from other parties and could not use the policies for their own electoral advantage (Levy 2006; Ortiz 2006).¹⁶ Voters were also well aware of these facts and, as a result, widely believed that they would receive program benefits regardless of the incumbent's partisan identity (Díaz-Cayeros et al. 2016).

Politics of implementation

As relatively pure programmatic policies, SPS and Progresá reduce the chances of partisan credit claiming via a clear policy design and resulting nonpartisan bureaucratic implementation. Progresá is known for having the strongest design and implementation protocol among the universe of conditional cash transfers in Latin America (De La O 2015), and SPS was quite similar in the period we examine. In contrast, none of the observational studies of the programmatic incumbent support hypothesis analyze pure programmatic policies (Larreguy et al. 2015; Manacorda et al. 2011; Pop-Eleches and Pop-Eleches 2012; Zucco 2008, 2013), and as such none exactly fit criteria *a–d* outlined above. These observational studies also analyze policies, political situations, and political contexts different from each other and our two experiments. Of

16. Recent work in lab experiments shows that individuals reward actors with proposal power for collective decisions (Duch, Przepiorka, and Stevenson 2015). A key condition for this finding is that individuals know the actor with such power. This is unlikely to hold in the Mexican case as the political bargain required for the adoption of Progresá and SPS prevented the party in power from advertising its role as the policy initiator.

course, every pair of studies that analyze a different policy or time period differs in an infinite number of ways, any one of which might be the linchpin that determines any differences in results. As such, future researchers have an important opportunity to help teach us more about how the programmatic incumbent support hypothesis, even though apparently inapplicable to pure programmatic policies, may work for more mixed programmatic-clientelistic government programs.

Consider that Plan de Atención Nacional a la Emergencia Social, the temporary emergency relief program analyzed in Manacorda et al. (2011), was a one-time effort. As such, and as discussed by the authors, voters may have associated the policy with the party responsible for its enactment and implementation. Similarly, in the case of Romania, Pop-Eleches and Pop-Eleches (2012) explain that the Euro 200 program had a well-known partisan intent. In Brazil, Hunter and Power (2007) and Zucco (2008) argue that one of the reasons for Bolsa Família's electoral success was the president's ability to claim credit for the program. In addition, Hall (2006) notes that the implementation of Bolsa Família was decentralized to municipalities, which led to charges of clientelism in the delivery of the program's benefits. Finally, Larreguy et al. (2015) study an interesting urban titling program crafted in 1973, before Mexico was a democracy. Under this program, incumbents organized events, claiming exclusive credit for the number of land titles granted during a given time period.

Similarly, in a field experiment in Uganda, Blattman, Eme-riau, and Fiala (2016) find that a program transferring cash to groups of unemployed youth increased support for and work on behalf of the opposition, among survey respondents. The policy was mostly funded by the World Bank, and a large proportion of survey respondents credited this institution with the creation of the program instead of the country's autocratic government. The authors hypothesize that this aspect of the program freed voters from clientelistic ties, thereby increasing their propensity to support and work for the opposition. Frey (2015) finds a similar result in Brazil, where local mayors have a limited ability to claim credit for Bolsa Família.

In the Philippines, Labonne (2013) finds, based in a randomized evaluation, that a conditional cash transfer program modeled after Brazil's Bolsa Família and Mexico's Progresa increased support for local incumbents only in competitive (i.e., nondynastic) municipalities with small federal budgetary transfers. The author does not find similar electoral benefits in competitive municipalities with large transfers. The explanation given for this difference is that in competitive municipalities with large transfers, mayors distribute funds among individuals not receiving the conditional cash transfer, making the voting behavior of program beneficiaries and non-beneficiaries indistinguishable from each other.

Theoretical explanations

The well-developed formal theory literature in this area is almost solely concerned with the electoral effects of discretionary, rather than programmatic, spending policies. For example, Dixit and Londregan (1996, 1132–33) devote the first two pages of their article to clarifying this point. Theoretical results suggest that targeting tactical or pork barrel spending occurs most often to benefit swing voters (Dixit and Londregan 1996) and marginal constituencies (Weingast, Shepsle, and Johnsen 1981), and this is indeed what the evidence shows (see Dahlberg and Johansson [2002] on the former and Keefer and Khemani [2009] and Primo and Snyder [2010] on the latter).

Studies have also found that such spending has substantial electoral payoffs (Evans 2006; Levitt and Snyder 1997). One clever exception, proposed in De La O (2015), holds that minority parties will rationally choose to support a programmatic policy that will hurt them electorally if opposing it imposes even larger costs. Unfortunately, as we show in section 4 of our supplementary appendix, the proposed formalization of this argument is consistent with both the programmatic incumbent support hypothesis and the opposite, and so the theory cannot be used to explain either. We also show there that the parameters of the theory are impossible to test from the Progresa (or SPS) experiment, and so either way the theory cannot explain the empirical results. However, the foundations of the model in De La O (2015) are of considerable value: when governments try to pass major public policies to cope with poverty in the presence of divided government they are much more likely to adhere to the programmatic policy ideal.

A variety of other perspectives might also help explain apparent divergent conclusions of different studies. For example, as we describe above, the Progresa experiment compared those who received the program recently and more distantly, unlike the SPS experiment, which had a clean control group that did not receive the program at all.¹⁷ Or, voters may simply reward incumbents for the implementation of programmatic policies as a result of reciprocity (Finan and Schechter 2012; Manacorda et al. 2011). Incumbents may also signal their commitment to the poor by adopting certain policies, making it rational for voters to reelect them (Diaz-Cayeros et al. 2009). Another possibility is that voters are retrospective in their voting behavior. Therefore, welfare improvements associated with government programs prompt them to reward incumbents (Pop-Eleches and Pop-Eleches

17. However, Green (2006), using a fuzzy RDD, also finds that Progresa does not have an impact on incumbent support or turnout.

2012). Finally, it may simply be that while incumbents pledge to implement a policy in a nonclientelistic manner, bureaucrats and local-level politicians fail to fulfill this promise (Hall 2006; Rocha-Menocal 2001).¹⁸

Future research

A valuable area for future research, then, would involve deriving theories that could shed light on when and under which political contexts and policy proposals the programmatic incumbent support hypothesis may or may not hold. Empirical research would also benefit from studies that seek to study the programmatic incumbent support hypothesis in policy and political contexts that are partially programmatic and partially clientelistic (Camp 2013; Diaz-Cayeros et al. 2009, 2016).

CONCLUDING REMARKS

When incumbent parties exercise discretion over government spending to benefit specific groups of voters, an impressive literature shows that these voters reward the incumbents with electoral support in subsequent elections. The reason for this powerful result would seem to be called into question by claims that nonpartisan programmatic policies—policies over which incumbents possess no discretion—have similar effects on voters making partisan decisions. We sought to contribute to an explanation of why this would happen and, also, why minority party politicians would, in the first place, support policies that would lead voters to oppose them in the next election.

We are in the unusual position for our discipline of being able to test an important hypothesis with two extremely large-scale, comparatively high-quality randomized experiments and a compelling natural experiment. We use these experiments to study two relatively pure forms programmatic policies and find no evidence in support of the programmatic incumbent support hypothesis. Our results may suggest why it is difficult for a country to break the cycle of perverse accountability (Stokes 2005): since only clientelistic policies yield electoral payoffs, incumbents are willing to adopt programmatic policies only in a context of high political competition.

The evidence we offer may raise other issues as well, such as the fact that parties in Latin America do not seem to have converged to political competition on a left-right dimension

as is common in Western democracies. Indeed, the cases of Progresá and SPS illustrate that certain policy-making domains, which may be of intense political contestation in other countries, have become depoliticized. The exact political context in which policies are adopted and implemented may have a major role in explaining differences of results in different settings. We look forward to future researchers shedding additional light on this important question in other national and policy contexts.

APPENDIX

FIXING DATA AND ANALYSIS ISSUES

We summarize here the coding errors created by faulty merging, evidence that these errors are sufficient to bias the causal effects, and the unconventional analysis techniques in De La O (2013, 2015). We also show how to fix these problems. More information is available in our supplementary appendix. Although we reach the unavoidable conclusion that these errors invalidate the claims in these works, De La O deserves credit for highlighting this important issue, thinking of the idea of repurposing a randomized experiment, gathering the necessary data, and making available a replication data set so that further discoveries became possible.

Coding Errors

In Mexico, as in many countries, electoral and population data are generated by government agencies that do not coordinate and so wind up with inconsistent geographies. For instance, while electoral authorities assign a population to a precinct, census officials may aggregate part of the same population with a neighboring village that happens to be outside that precinct. Neither office is necessarily correct or incorrect, as they have differing goals, but the result is difficult for analysts trying to merge incompatible data sources. This lack of interagency coordination is a common problem in the analysis of elections in many countries, but the issue here caused particular difficulties when De La O (2013, 5) “overlayed the 506 experimental villages to the smallest unit of outcome measure for which census, precinct, and electoral data roughly coincide: the *sección electoral* (precinct).” These precincts were “name matched” using the textual names of villages in different data files, from organizations that assigned different geographic locations and meanings to the same names. Unfortunately, the result turns out to be that 71.3% of the observations (villages) were incorrectly matched to areas with similar names but from places outside of designated treatment and control precincts (supplementary appendix, sec. 2.2). We confirmed these facts via formal Mexican Freedom of Information Act requests we filed and conversations with officials at the Dirección Ejecutiva de Organización

18. Blind retrospection may be yet another mechanism accounting for a positive effect of programmatic policies on incumbent support (Achen and Bartels 2004). We consider this possibility unlikely, as other studies have questioned the existence of such phenomena (see, e.g., Fowler and Hall 2016).

Electoral (DEOE) and Dirección Ejecutiva del Registro Federal de Electores at Mexico's National Electoral Institute (INE).¹⁹

Of the incorrect matches due to this procedure, two caused especially large biases. Correcting only these two units (or all of them) leads to the conclusion that Progresas's programmatic policies have little or no impact on turnout or voting for the incumbent. The large coding errors are apparent in the first column, for voter turnout, and the fourth, for the incumbent (PRI) vote, in the distributions portrayed in figure A1. Since it is impossible for more than 100% of people to vote, or to vote for any one party, every observation above the dashed line at 100 is mistaken. Moreover, these include some large errors that extend into the impossible region beyond 50% of the range of the original data and for turnout more than three times the range. Turnout obviously cannot be 375%, as the data analyzed in De La O (2013, 2015) indicate. In addition to these known errors, the large number of observations with nearly zero turnout (fig. A1, *left column, bottom*) in a presidential election with 65% overall turnout (IFE 2013) are of dubious validity.

Before turning to the consequences of these data errors, we analyze turnout and incumbent (PRI) vote among those officially registered (using precincts in De La O [2013] and an alternative GIS-determined sample). These alternative variables appear in figure A1 in columns 2 and 3 for turnout and columns 5 and 6 for incumbent vote. These data come from one source, with the same geographic boundaries and no possibility of coding errors. As a result, we can see in the figure that all the observations naturally fall within the possible region between 0% and 100%.

Our alternative turnout and incumbent vote share variables, based on officially registered voters, have no such measurement error bias problem since all those in the denominator would appear in the numerator if they choose to vote. If Progresas works in part by increasing levels of registration—as hypothesized in De La O (2013, 7–8)—then the interpretation of the turnout and vote variables would change, but no posttreatment or other bias or inefficiency would be introduced.²⁰ Since no evidence has been offered for this hypothe-

sis previously, we directly test it and present our results in tables 14–17 and 24–25 of our supplementary appendix. The results, with small confidence intervals around zero, demonstrate that the program had little or no impact on registration rates. These alternative codings for turnout and vote are considerably cleaner tests of the programmatic hypothesis, even if the original data had no coding errors.

Conditions for Bias

Thus far, we have revealed the existence of data errors, but that alone is insufficient to change any conclusion; the real question is whether these errors make a difference in the results. We find that this is indeed the case, especially for two huge outliers. The particular impact of these outliers is their correlation with the control variable “lagged population,” which we plot in figure A2. In the left panel, we plot the raw data, which reveals the two outliers (precincts 266 and 1502), each 90–100 times the size of the median and appearing in the treated group, with nothing comparable in the control group. De La O (2013, 6) was right to check the balance between the means of the treated and control groups, but checking solely based on the means led to missing the massive imbalance in the tails of the two distributions evident in this figure. The consequence of imbalance is model dependence (Ho et al. 2007). Since the most important advantage of a correctly analyzed randomized experiment is the absence of model dependence, introducing these coding errors eliminated this important benefit of randomization. Below, we illustrate and correct this problem. As we show in the next section, regardless of the reason for these outliers, when we correct for the imbalance they create, all evidence of significant effects of Progresas on turnout or incumbent support vanishes.

We go a step further and formally evaluate the bias these (incorrect) outlier observations have on the least squares analysis in De La O (2013). We do this by computing the “statistical leverage” of these observations, with and without the lagged population variable. Observations with larger values of leverage have more influence on the magnitude of coefficient estimates in linear regression.²¹ As the right panel

19. An important clue comes when De La O (2013, n. 17) tries to account for some of the data with turnout and vote percentages greater than 100 by noting that 10 precincts “contained special voting booths where by law out-of-precinct voters can cast a ballot.” Although such precincts do exist, none are in the author's data set—a fact that can be verified in data made publicly available by the Mexican government (<https://www.ine.mx/voto-y-elecciones/resultados-electorales/>). To be more specific, of the four types of polling stations in Mexico—*básicas*, *contiguas*, *extraordinarias*, and *especiales*—only *especiales* allow out-of-precinct voters. Yet, none of the polling stations included in the sample analyzed are of this type.

20. To see why no posttreatment bias is introduced, let $V(t)$ and $R(t)$ represent the potential number of those who turned out and registered under the treatment condition $t = 0, 1$, respectively. This notation allows for the possibility that the treatment can affect both turnout and registration. Our quantity of interest, i.e., turnout rate, is denoted as $Y(t) = V(t)/R(t)$, which is a well-defined potential outcome. Since we do not condition on the realized registration rate, no posttreatment bias is introduced.

21. The statistical leverage of observation i is defined as $x_i^T(X^T X)^{-1}x_i$, where X is an $n \times k$ matrix of pretreatment covariates with $k \times 1$ row x_i .

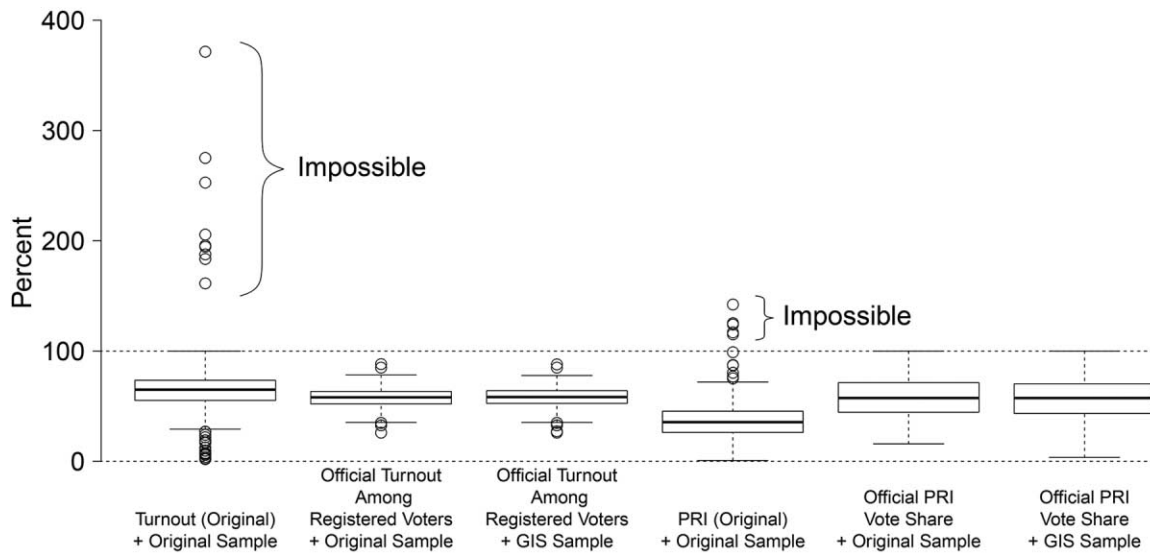


Figure A1. Univariate distribution of turnout and incumbent party vote in 2000. Compares the variables originally constructed in De La O (2013) via name matching (in cols. 1 and 4) with the official turnout among registered voters and PRI vote share in the name-matching sample (cols. 2 and 5) and in the GIS sample (cols. 3 and 6).

in figure A2 makes clear, the statistical results in De La O (2013) are driven largely by these two observations. Indeed, when controlling for lagged population, the two precincts have leverage that is 15–20 times larger than the median (see the vertical axis of the right panel). In contrast, when population is not included as a covariate, the leverage of these precincts is a modest 1.5 to 2.75 times larger than the median (see the horizontal axis). Since leverage is computed solely from the explanatory variables, this result applies identically to both outcome variables, turnout and incumbent vote. Particularly unfortunate is that both of these extremely high-

leverage observations happen to also have extremely (and unrealistically) low reported turnout rates (1.78 and 2.59 percentage points, respectively) and PRI vote shares (0.74 and 1.08 percentage points). Judging from how the data were constructed, the considerably larger vote and turnout results based on official registration, the much larger vote and turnout national figures, or by comparisons with similar precincts, it is likely that these extremely low turnout and vote percentages are incorrect.

We have thus demonstrated all the conditions for bias in the main results: errors in the data that matter, substantial

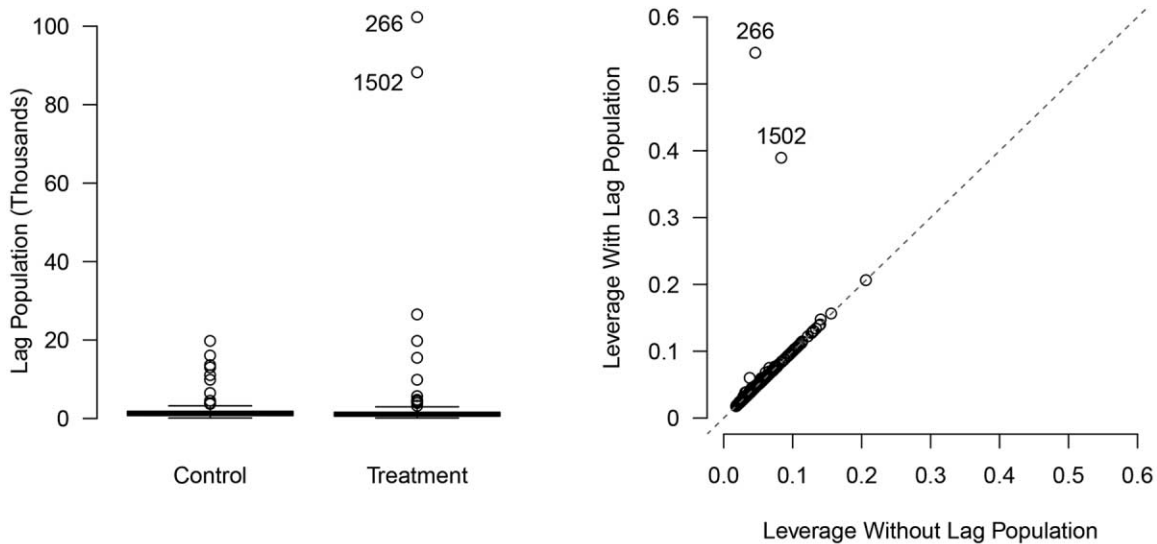


Figure A2. Population outliers and statistical leverage. Reveals the extreme degree to which two of the miscoded observations are outliers (left) with extremely high statistical leverage (right). Left, raw data; right, computed from the least squares analysis in De La O (2013).

imbalance in the treated and control distributions, huge outliers highly correlated with a covariate included in the regression, and extreme outcome values for the same outlier units.

Unconventional Model Specifications

Finally, we consider the unusual model specifications in De La O (2013, 2015). Across electoral studies in American and comparative politics, researchers almost always measure the vote for an incumbent party as the *incumbent vote share*—the number of ballots cast for the incumbent party divided by the total number of ballots cast. Researchers from most countries then typically model vote share by including a lagged value of vote share as a control variable.

In contrast, De La O (2013, 2015) measured the incumbent vote for a party as the number of voters for that party divided by the *voting age population*—including voters, nonvoters, noncitizens, those ineligible to vote for other reasons, and so on. The article then includes a lagged control variable that was not the incumbent vote share, and not the voting age population, but instead the number of people in the *total population*—including the voting age population as well as all those under 18. As a result of this unusual decision, the outcome variable in this analysis does not add to 100% across parties, or across parties plus nonvoters, and the outcome and control variables are not logically or necessarily related. Although the model is theoretically possible—more noncitizens or infants in an area could in principle lead to more incumbent voting—this idea was not suggested in De La O (2013, 2015); indeed, to our knowledge, no other published statistical model of the electoral politics in any country and election has ever made these assumptions or chosen this type of specification.

More importantly, because counts, unlike percentages, are unbounded, outliers on this scale greatly exacerbate the bias that results from influential outliers. This is why, as we show, either correcting the coding errors or switching to the dominant method of constructing vote variables eliminates any support for the effects of programmatic policies on partisan outcomes.

ACKNOWLEDGMENTS

We are grateful to Isadora Antoniano (formerly of the Mexican Electoral Institute), Miguel Rojano (Director of Electoral Cartography) and Luis Ruvalcaba (Deputy Director of Geographical Electoral Systems Development) for discussing with us the intricacies of electoral cartography in Mexico; to Tina Green and Ana De La O for data and replication information, and for help with follow-up questions; to Wangyal Shawa for his helpful GIS technical advice; to Rikhil Bhavnani, Graeme Blair, Chris Blattman, Ben Fifield, Miriam Golden, Ken Greene, Guy Grossman, Macartan Humphreys, John Londregan, Gabriel López-Moctezuma, Will Lowe, Gri-

goro Pop-Eleches, Jake Shapiro, and Cesar Zucco Jr. for helpful comments; and to Horacio Larreguy, Tara Slough, and John Marshall for very productive discussions.

REFERENCES

- Achen, Christopher H., and Larry M. Bartels. 2004. "Blind Retrospection Electoral Responses to Drought, Flu, and Shark Attacks." Working paper, Princeton University. <https://goo.gl/C4te4H>.
- Alesina, Alberto, and Guido Tabellini. 1990. "A Positive Theory of Fiscal Deficits and Government Debt." *Review of Economic Studies* 57 (3): 403–14.
- Behrman, Jere R., and Petra E. Todd. 1999. "Randomness in the Experimental Samples of Progesa: Education, Health, and Nutrition Program." Research report, International Food Policy Research Institute, Washington, DC. <http://goo.gl/YLWz8x>.
- Blattman, Chris, Mathilde Emeriau, and Nathan Fiala. 2016. "Do Employment Programs Buy Votes? A Field Experiment in Uganda." Working paper, Columbia University.
- Calvo, Ernesto, and Maria Victoria Murillo. 2004. "Who Delivers? Partisan Clients in the Argentine Electoral Market." *American Journal of Political Science* 48 (4): 742–57.
- Camp, Roderic A. 2013. "The 2012 Presidential Election and What It Reveals about Mexican Voters." *Journal of Latin American Studies* 45:451–81.
- Cornelius, Wayne A. 2004. "Mobilized Voting in the 2000 Elections: The Changing Efficacy of Vote Buying and Coercion in Mexican Electoral Politics." In Jorge I. Domínguez and Chappell Lawson, eds., *Mexico's Pivotal Democratic Election: Candidates, Voters, and the Presidential Election of 2000*. Stanford, CA: Stanford University Press, 47–65.
- Dahlberg, Matz, and Eva Johansson. 2002. "On the Vote-Purchasing Behavior of Incumbent Governments." *American Political Science Review* 96 (1): 27–40.
- De La O, Ana L. 2013. "Do Conditional Cash Transfers Affect Electoral Behavior? Evidence from a Randomized Experiment in Mexico." *American Journal of Political Science* 57 (1): 1–14.
- De La O, Ana L. 2015. *Crafting Policies to End Poverty in Latin America: The Quiet Transformation*. Cambridge: Cambridge University Press.
- Díaz-Cayeros, Alberto, Federico Estevez, and Beatriz Magaloni. 2009. "Welfare Benefits, Canvassing, and Campaign Handouts." In Jorge I. Domínguez, Lawson Chappell, and Alejandro Moreno, eds., *Consolidating Mexico's Democracy: The 2006 Presidential Campaign in Comparative Perspective*. Baltimore: Johns Hopkins University Press, 229–45.
- Díaz-Cayeros, Alberto, Federico Estevez, and Beatriz Magaloni. 2016. *The Political Logic of Poverty Relief: Electoral Strategies and Social Policy in Mexico*. Cambridge: Cambridge University Press.
- Dixit, Avinash, and John Londregan. 1996. "The Determinants of Success of Special Interests in Redistributive Politics." *Journal of Politics* 58 (4): 1132–55.
- Duch, Raymond, Wojtek Przepiorka, and Randolph Stevenson. 2015. "Responsibility Attribution for Collective Decision Makers." *American Journal of Political Science* 59 (2): 372–89.
- Evans, Thomas. 2006. "The Effects of Discretionary Federal Spending on Parliamentary Election Results." *Economic Inquiry* 44 (2): 234–48.
- Finan, Frederico, and Laura Schechter. 2012. "Vote-Buying and Reciprocity." *Econometrica* 80 (2): 863–81.
- Fowler, Anthony, and Andrew B. Hall. 2016. "Do Shark Attacks Influence Presidential Elections? Reassessing a Prominent Finding on Voter Competence." Working paper, University of Chicago.
- Franzese, Robert J. 2002. *Macroeconomic Policies of Developed Democracies*. New York: Cambridge University Press.
- Frey, Anderson. 2015. "Cash Transfers, Clientelism, and Political Enfranchisement: Evidence from Brazil." Working paper, University of British Columbia.

- Golden, Miriam, and Brian Min. 2013. "Distributive Politics around the World." *Annual Review of Political Science* 16:73–99.
- Green, Tinz. 2006. "Do Social Transfer Programs Affect Voter Behavior? Evidence from PROGRESA, 1997–2000." PhD diss., University of California, Berkeley.
- Gwatkin, Davidson R., Adam Wagstaff, and Abdo S. Yazbeck, eds. 2005. *Reaching the Poor with Health, Nutrition, and Population Services: What Works, What Doesn't, and Why*. Washington, DC: World Bank.
- Hall, Anthony. 2006. "From Fome Zero to Bolsa Família: Social Policies and Poverty Alleviation under Lula." *Journal of Latin American Studies* 38:689–709.
- Hicken, Allen. 2011. "Clientelism." *Annual Review of Political Science* 14:289–310.
- Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart. 2007. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15:199–236.
- Hunter, Wendy, and Timothy J. Power. 2007. "Rewarding Lula: Executive Power, Social Policy, and the Brazilian Elections of 2006." *Latin American Politics and Society* 49 (1): 1–30.
- IFE (Instituto Federal Electoral). 2013. "Estudio Censal de la Participación Ciudadana en las Elecciones Federales de 2012." Report, Dirección Ejecutiva de Capacitación Electoral y Educación Cívica.
- Imai, Kosuke, and Zhichao Jiang. 2018. "A Sensitivity Analysis for Missing Outcomes Due to Truncation-by-Death under the Matched-Pairs Design." *Statistics in Medicine* 37 (20): 2907–22.
- Imai, Kosuke, Gary King, and Clayton Nall. 2009a. "The Essential Role of Pair Matching in Cluster-Randomized Experiments, with Application to the Mexican Universal Health Insurance Evaluation." Replication data, Harvard Dataverse. <https://doi.org/10.7910/DVN/9RJGWB>.
- Imai, Kosuke, Gary King, and Clayton Nall. 2009b. "The Essential Role of Pair Matching in Cluster-Randomized Experiments, with Application to the Mexican Universal Health Insurance Evaluation." *Statistical Science* 24 (1): 29–53.
- Imai, Kosuke, Gary King, and Elizabeth A. Stuart. 2008. "Misunderstandings among Experimentalists and Observationalists about Causal Inference." *Journal of the Royal Statistical Society A* 171 (2): 481–502.
- Imai, Kosuke, Gary King, and Carlos Velasco Rivera. 2018. "Do Nonpartisan Programmatic Policies Generate Partisan Electoral Effects? Evidence from Two Large Scale Experiments." Replication data, Harvard Dataverse. <https://doi.org/10.7910/DVN/70SNIS>.
- Keefer, Philip, and Stuti Khemani. 2009. "When Do Legislators Pass on Pork? The Role of Political Parties in Determining Legislator Effort." *American Political Science Review* 103 (1): 99–112.
- King, Gary, Emmanuela Gakidou, Kosuke Imai, Jason Lakin, Clayton Nall, Ryan T. Moore, Nirmala Ravishankar, Manett Vargas, Martha María Téllez-Rojo, Juan Eugenio Hernández Ávila, Mauricio Hernández Ávila, and Héctor Hernández Llamas. 2009a. "Public Policy for the Poor? A Randomised Assessment of the Mexican Universal Health Insurance Programme." *Lancet* 373:1447–54.
- King, Gary, Emmanuela Gakidou, Kosuke Imai, Jason Lakin, Clayton Nall, Ryan T. Moore, Nirmala Ravishankar, Manett Vargas, Martha María Téllez-Rojo, Juan Eugenio Hernández Ávila, Mauricio Hernández Ávila, and Héctor Hernández Llamas. 2009b. "Public Policy for the Poor? A Randomized Ten-Month Evaluation of the Mexican Universal Health Insurance Program." Replication data, Harvard Dataverse. <https://doi.org/10.7910/DVN/P6NCOM>.
- King, Gary, Emmanuela Gakidou, Nirmala Ravishankar, Ryan T. Moore, Jason Lakin, Manett Vargas, Martha María Téllez-Rojo, Juan Eugenio Hernández Ávila, Mauricio Hernández Ávila, and Héctor Hernández Llamas. 2007. "A 'Politically Robust' Experimental Design for Public Policy Evaluation, with Application to the Mexican Universal Health Insurance Program." *Journal of Policy Analysis and Management* 26 (3): 479–506.
- Kitschelt, Herbert, and Steven Wilkinson. 2007. "Citizen-Politician Linkages: An Introduction." In Herbert Kitschelt and Steven Wilkinson, eds., *Patrons, Clients, and Policies: Patterns of Democratic Accountability and Political Competition*. Cambridge: Cambridge University Press.
- Klarman, Michael J. 1997. "Majoritarian Judicial Review: The Entrenchment Problem." *Georgetown Law Journal* 85:491–554.
- Labonne, Julien. 2013. "The Local Electoral Impacts of Conditional Cash Transfers." *Journal of Development Economics* 104:73–88.
- Larreguy, Horacio, John Marshall, and Laura Trucco. 2015. "Breaking Clientelism or Rewarding Incumbents? Evidence from an Urban Titling Program in Mexico." Working paper, Harvard University. <http://goo.gl/P4YXky>.
- Layton, Matthew L., and Amy Erica Smith. 2015. "Incorporating Marginal Citizens and Voters: The Conditional Electoral Effects of Targeted Social Assistance in Latin American." *Comparative Political Studies* 48 (7): 854–81.
- Levitt, Steven, and James M. Snyder. 1997. "The Impact of Federal Spending on House Election Outcomes." *Journal of Political Economy* 105 (1): 30–53.
- Levy, Santiago. 2006. *Progress against Poverty: Sustaining Mexico's Progres-Oportunidades Program*. Washington, DC: Brookings Institution.
- Manacorda, Marco, Edward Miguel, and Andrea Vigorito. 2011. "Government Transfers and Political Support." *American Economic Journal* 3 (3): 1–28.
- Ortiz, Mauricio. 2006. *El Seguro Popular: Una Cronica de la Democracia Mexicana*. Mexico City: Fondo de Cultura Economica.
- Pop-Eleches, Cristian, and Grigore Pop-Eleches. 2012. "Targeted Government Spending and Political Preferences." *Quarterly Journal of Political Science* 7 (3): 285–320.
- Posner, Eric A., and Adrian Vermeule. 2002. "Legislative Entrenchment: A Reappraisal." *Yale Law Journal* 111 (7): 1665–705.
- Primo, David M., and James M. Snyder. 2010. "Party Strength, the Personal Vote, and Government Spending." *American Journal of Political Science* 54 (2): 354–70.
- Rocha-Menocal, Alina. 2001. "Do Old Habits Die Hard? A Statistical Exploration of the Politicisation of Progres, Mexico's Latest Federal Poverty-Alleviation Programme, under the Zedillo Administration." *Journal of Latin American Studies* 33:513–38.
- Skoufias, Emmanuel. 2005. *PROGRESA and Its Impacts on the Welfare of Rural Households in Mexico*. Washington, DC: International Food Policy Research Institute.
- Skoufias, Emmanuel, Benjamin Davis, and Jere R. Behrman. 1999. "An Evaluation of the Selection of Beneficiary Households in the Education, Health, and Nutrition Program (PROGRESA) of Mexico." Research report, International Food Policy Research Institute, Washington, DC.
- Sterk, Stewart E. 2003. "Retrenchment on Entrenchment." *George Washington Law Review* 71 (2): 231–54.
- Stokes, Susan C. 2005. "Perverse Accountability: A Formal Model of Machine Politics with Evidence from Argentina." *American Political Science Review* 99 (3): 315–25.
- Stokes, Susan C., Thad Dunning, Marcelo Nazareno, and Valeria Brusco. 2013. *Brokers, Voters, and Clientelism: The Puzzle of Distributive Politics*. Cambridge: Cambridge University Press.
- Thompson, Dennis F. 2005. "Democracy in Time: Popular Sovereignty and Temporal Representation." *Constellations* 12 (2): 245–61.
- Weingast, Barry R., Kenneth Shepsle, and Christopher Johnsen. 1981. "The Political Economy of Benefits and Costs: A Neoclassical Approach." *Journal of Political Economy* 89 (4): 642–64.
- Zucco, Cesar. 2008. "The President's 'New' Constituency: Lula and the Pragmatic Vote in Brazil's 2006 Presidential Elections." *Journal of Latin American Studies* 40 (1): 29–49.
- Zucco, Cesar. 2013. "When Payouts Pay Off: Conditional Cash Transfers and Voting Behavior in Brazil, 2002–10." *American Journal of Political Science* 57 (4): 810–22.