

# The medium-run effects of a foreign election intervention: Haiti's presidential elections, 2010–2015

Craig Palsson

Huntsman School of Business, Utah State University, Logan, Utah, USA

## Correspondence

Craig Palsson, Huntsman School of Business, Utah State University, Logan, UT 84322, USA.

Email: [craig.palsson@usu.edu](mailto:craig.palsson@usu.edu)

## Abstract

The US and other foreign actors often intervene in elections in countries transitioning to democracy. I examine the effects of such interventions on voter behavior. In 2010, the United States intervened in Haiti's presidential election, advancing Michel Martelly over Jude Célestin. I look at the relationship between the intervention and voter behavior in the next election using Célestin's 2010 vote share as a measure of the intervention's intensity in a modified difference-in-differences analysis. Areas with greater Célestin support in 2010 had lower turnout in 2015. The relationship is robust to many sensitivity tests that account for possible confounding effects from fraud.

## KEYWORDS

disenfranchisement, elections, Haiti, turnout

## JEL CLASSIFICATION

D72, F5, K16

## 1 | INTRODUCTION

Democracy is an important institution for economic outcomes—both macroeconomic (Acemoglu et al., 2019) and microeconomic (Kudamatsu, 2012; Naidu, 2012). Since democracy complements a market-based economy, a major component of the US's foreign policy has been democracy assistance (Rose, 2000). From 2015 to 2018, expenditures on democracy assistance programs by United States Agency for International Development and the US State Department averaged \$2.2 billion per year (U.S. Government Accountability Office, 2020). These programs support elections, but sometimes they include direct interventions. In the short-run, such interventions tend to succeed in promoting the US's preferred candidate (Levin, 2016). But interventions might replace the median voter's preferences with the US's, and we do not know whether this has lasting consequences for democracy in the targeted country.

To investigate the medium-run consequences of foreign interventions in elections, I look at Haiti in the 2010 and 2015 presidential elections. In 2010, in the aftermath of a devastating earthquake, Haiti's election faced serious challenges after the first round. For the first time in Haiti's history, no candidate gained at least 50% of votes in the first round, so the presidential election advanced to a second-round runoff. But there was conflict on which candidates should advance. Three candidates were in contention: Mirlande Manigat, the undisputed first-round leader; Jude Célestin, the incumbent party's candidate and second-place finisher; and Michel Martelly, who took third place by less than one percentage point. The election rules dictated that the top two finishers, Manigat and Célestin, would advance,

**Abbreviations:** BRIDES, Bureau de Recherche en Informatique et en Développement Economique et Social; CEP, Provisional Electoral Council; HIRS, Haitian Internal Revenue Service; OAS, Organization of American States; USAID, United States Agency for International Development.

but doubts about the first-round's legitimacy and a strong opposition campaign brought the country to a political impasse over whether Célestin should proceed. In an unprecedented move, the United States and the Organization of American States (OAS) pressured the Haitian government to remove Célestin from the election and advance Martelly, a solution the Haitian government resisted but ultimately accepted. I examine how the intervention affected turnout in the 2015 presidential election.

To explore the intervention's consequences, I exploit that the intervention disproportionately affected Célestin's supporters in 2010. The rationale behind the OAS recommendation was that a random sample of tally sheets uncovered some sheets that the OAS deemed "irregular." After discarding these sheets, Martelly's vote share exceeded Célestin's. Although discarding the sheets eliminated votes for every candidate, it implicitly validated Martelly's and Manigat's votes while disregarding the pivotal Célestin voters. Thus, I assume that the intervention affected Célestin voters more than others: the intervention disregarded a greater share of voters in sections where 40% of the vote went to Célestin than where he won 20%. With this definition of treatment, the empirical strategy uses a modified difference-in-differences specification, comparing turnout in sections across the two elections.

The data come from section-level vote tallies in both the 2010 and 2015 elections, and I provide new evidence on the data's legitimacy. Some might worry about whether the vote tallies mean anything because there were allegations of fraud. Although suspicious tally sheets were discarded, there might be concern that the filtered sheets still contain fraud. To address these concerns, I match the election results to pre-election surveys and find the results are strongly correlated with the survey's predictions. Moreover, visual tests for vote fraud described in Klimek et al. (2012) report no significant signs of fraud. Although the results do not say there was no fraud or that fraud did not influence the election, they do imply that such irregularities can be treated as noise.

The modified difference-in-differences results reveal that while turnout decreased across elections, it decreased more in sections that supported Célestin. Relative to a section where he won 20% of the 2010 vote, a section where Célestin won 70% of the vote dropped its 2015 turnout by 3.5 percentage points. The result is driven entirely by a decrease in votes; there is no difference in voter registration. To support the results, I run a falsification test using Manigat's 2010 vote share as the treatment. The same regression specification shows no correlation between Manigat's 2010 support and 2015 turnout. If the effect is causal, then the results suggest that the intervention did not discourage Manigat's pivotal voters, since she remained in the race, but that removing Célestin adversely affected his pivotal voters.

Although the data exhibit no strong evidence of fraud, there is still a concern that the noise might induce a spurious correlation. To address the concerns about fraud, I test the robustness of the results with three approaches. First, I omit problematic districts identified by the OAS and find the results persist. Second, I run a sensitivity check to see what patterns of fraud would create spurious results. Since fraud is unobservable, and since there are not even clear indicators on which sections would be most affected by fraud, the test has to explore different assumptions about the patterns of fraud. The patterns are represented by two parameters: (1) what share of sections experienced fraud and (2) what percentage of votes in those sections were fraudulent. I run two versions of the test. First, I assume the intensity of fraud is uncorrelated with Célestin's vote share. Second, to bias the results against Célestin, I assume that his top-performing sections are the ones where fraud occurred: if the assumption is that 10% of sections experienced fraud and that 20% of those votes were fraudulent, then the simulation removes 20% of Célestin's votes in sections in his top 10% of support. The test's results show that, for the results to be spurious, fraud would have to be much more pervasive than evidence supports. Finally, I test the sensitivity of the results to assumptions about voter behavior on sheets discarded for irregularities.

The results warn about the effects of foreign powers intervening in elections. There is some evidence that foreign aid might improve the integrity of elections (Ubert & Jackson, 2020), but this paper shows that interventions that go beyond financial support could hurt electoral integrity. Previous work has shown that in the concurrent election, voters in a targeted country tend to support the candidate promoted by the intervention (Levin, 2016). In hypothetical interventions, we know that voter support for the interventions depends on whether it fits their political preferences (Corstange & Marinov, 2012). But we do not understand how foreign interventions affect subsequent elections. The extent to which the election is hurt could depend on the strength of the electoral institutions. For instance, work by Allcott and Gentzkow (2017) suggests that Russia's interference in the US 2016 presidential election likely had an inconsequential effect on the results, and the historic-high turnout in the 2018 midterm elections suggests that any discouragement effect would have been small too. But since many countries that receive democracy assistance are only in the initial stages of a political transition, the intervention's effects could be different. When weighing the costs and benefits of a democracy assistance intervention, the United States should recognize that short-run benefits could be offset by long-run costs.

These results also contribute to our understanding of the economics of voting. The canonical voting model from Downs (1957) posits that voters go to the polls based on the cost of voting, the benefits of the policy, and the expectation that the voter will swing the decision; later models include utility from the act of voting (Riker & Ordeshook, 1968). With the Downsian model in mind, there are two ways to think about the results. First, the intervention discarded votes from pivotal voters, so it may reduce a voter's subjective beliefs that he will swing the decision. The importance of beliefs about being a pivotal voter have been debated—lab results show a positive correlation between beliefs and voting behavior (Duffy & Tavits, 2008) and that voters will abstain from voting if they are uninformed and believe they are pivotal (Battaglini et al., 2010), but data from small-scale elections suggest that winning margins are too large for the pivotal voter model to explain voting behavior (Coate et al., 2008). While I cannot conclude that the pivotal-voter model explains Haitian behavior, one interpretation is that the intervention changed the pivotal voter from Haitian citizens to foreign actors, diminishing the incentive to vote. A second possible mechanism is that the intervention decreased the expected utility from voting. Since utility from voting and performing civic duties strongly influence voter behavior (Ali & Lin, 2013; Dellavigna et al., 2017; Gerber et al., 2008), foreign actors need to be aware that an intervention could significantly affect how voters enjoy elections.

Finally, these findings are a significant deviation from similar empirical work on disenfranchisement from fighting fraud. The intervention disenfranchised many Haitian voters by disregarding their votes due to apparent irregularities. This parallels many debates in the United States about policies that prevent election tampering but that might also disenfranchise voters. For instance, a rising political issue in the United States is the fight over voter identification laws. While some believe that such laws will fight fraud, others believe requiring identification will disenfranchise groups who are less likely to have government issued IDs. Yet empirical work has shown that such laws have no significant effects on fraud or on turnout (Cantoni & Pons, 2019; Hoekstra & Koppa, 2021).<sup>1</sup> On the other hand, voters in Florida were more likely to vote if their right to vote had been threatened by an aborted state-wide purge (Biggers & Smith, 2020). But these studies look at voter behavior in the United States, a country with established democratic institutions. Countries with weak or new democratic institutions, on the other hand, may respond differently to disenfranchisement (Collier & Vicente, 2012).

## 2 | BACKGROUND ON HAITIAN POLITICS, 1990–2016

Haiti has a long history of fraught politics, and foreign intervention is not unfamiliar. Haiti's modern political system began with the 1987 Constitution, which eliminated the office of President for Life and opened the executive position to popular election. The constitution created a permanent, independent council to run elections, but since the beginning the council has been neither permanent nor independent. Instead, the council that runs elections is the Provisional Electoral Council (CEP) whose authority lasts for one election. A regular concern is that the president appoints the members of the CEP, which has led to accusations in the past (Miles & Feeney, 2001) and in 2010 (Johnston & Weisbrot, 2011) that the council unjustly favors the incumbent party. This will be discussed below.

The cumulative effect of Haiti's history has shaped politics over the past 3 decades.

### 2.1 | Elections and the international community, 1990–2006

Ever since democracy started in Haiti in 1990, the international community has been involved to some degree. The smallest and least controversial role it has played has been to monitor the elections. Monitoring began in the first election in 1990, when Haiti's provisional government invited several hundred international observers to monitor the election, which Jean-Bertrand Aristide won with 70% of the vote (von Hippel, 1995). The first major intervention was in 1995. After a coup had ousted Aristide in 1991, less than 8 months after taking office, the international community helped return him to the presidency in 1994 on the condition that he step down in 1995 and hold an election (Mobekk, 2001). The nature of the international community's involvement changed in the election of 2006 with two major interventions. First, the international community forbade Aristide's party (Fanmi Lavalas) from running a candidate (Dupuy, 2006). Second, the international community influenced the vote count. The election committee, when counting the votes, discovered a significant number of ballots were blank. It was unclear whether the blanks were mistakes, protest votes, or intentional fraud, but according to election rules, they had to be counted. But there were enough blanks that including them in the total would dilute the leading candidate's vote share enough to force a second

round. The international community intervened, saying that while the blanks had to be counted, the rules did not specify how they had to be counted, and awarded them proportionally to each candidate, thus avoiding a runoff (Dupuy, 2006).

With the coups and electoral interventions through the 1990s and 2000s, Haitian voters became disillusioned with the election process, which led to large variation in turnout. One source of disillusionment was that people did not see results from the political process. The international community was so focused on elections that they often ignored the abusive administrations they created, which eroded trust in the international community (von Hippel, 1995). Many Haitians lost faith in the election process, believing that it never led to change, and they came to refer to their system as *pepe* (secondhand) democracy. With this disillusionment came large swings in voter turnout. In 1990, 50% of registered voters turned out, but in 1995, when Aristide was forced to step down despite his interrupted first term, turnout for the election was around 28% (Nohlen, 2005, p. 392). The return of Aristide brought a return of participation, though the exact numbers are disputed: the CEP and other election observers reported turnout at 60% (Miles & Feeney, 2001).<sup>2</sup> In 2006, turnout hit 63% (Dupuy, 2006, pp. 168–169).

While the international community has meddled with Haitian elections, we do not know whether its influence has caused Haitians to change their voting behavior. One difficulty with attributing any adverse effects is that the interventions were often small. For instance, although one of the conditions for restoring Aristide was that he allow the election to proceed as scheduled in 1995 and to not run as a candidate, the condition was enforcing constitutional restraints. Haitian presidents cannot serve two consecutive terms. The controversy was whether Aristide's interrupted tenure counted as a full term. Similarly, the 2006 interventions did not seem to alter the outcomes of the election. While the international community prevented Aristide's party from officially running a candidate, the party's 1995 candidate and winner René Préval created a new party and won in 2006. And although Préval's victory was helped by the intervention that pushed his first-round vote share above 50%, he was the clear favorite leading up to the election, with the runner-up only receiving 12% (Dupuy, 2006). Thus, while the international community intervened, its actions were arguably inconsequential. But that changed significantly with the 2010 election.

## 2.2 | The 2010 presidential election

In November 2010, 10 months after an earthquake that killed 200,000 people, Haiti held a presidential election. Turnout was low at 23%, which was to be expected because of the tumultuous post-earthquake environment. But that was not the only problem with the election; the OAS, whom the Government of Haiti had invited to monitor the election, reported, “the day of elections was marred by disorganization, dysfunction, various types of irregularities, ballot stuffing and incidents of intimidation, vandalism of polling stations and violence” (Organization of American States, 2011b). While acknowledging the problems, the OAS expressed confidence in the election: “Based on its observations in the eleven electoral departments, the Joint Mission does not believe that these irregularities, serious as they were, necessarily invalidated the process” (Organization of American States, 2010). Despite this supporting statement, counting the votes was contentious. The CEP failed to receive 1365 tally sheets (about 12.2% of the vote). Then, before reporting the results, the CEP discarded 312 tally sheets (about 7.6% of the votes) that appeared “irregular” (Johnston & Weisbrot, 2011).

With a first round that had 18 candidates, no candidate won a sufficient share of the votes to become the outright winner. Haitian elections consist of a first round of voting with all candidates, but if no candidate wins more than 50% of votes then a second round runoff is held with the top two candidates. After counting the remaining votes, the leading candidate, with 31.4% of the vote, was Mirlande Manigat, a university administrator and former First Lady. In second place, with 22.5%, was Jude Célestin, the candidate handpicked by then-president René Préval. Finally, in the third place, with 21.8% of the vote (less than 1 percentage point behind second), was Michel Martelly, a music star on his first political foray. For the first time in Haiti's history, the presidential election would require a runoff.

Despite the CEP's removal of “irregular” tally sheets, some accused Célestin of using Préval's administration to manipulate the election, and the election results were protested. The OAS observers expressed little faith in the claims, “More subversive of the process was the toxic atmosphere created by the allegations of ‘massive fraud’. The [election monitors] observed instances where even before the voting started, any inconvenience or small problem led to the immediate cry of fraud” (Organization of American States, 2011b). Indeed, the claims were so tenuous that while Manigat and Martelly initially led the charge for canceling the vote, they withdrew their complaints once they heard they were in contention for the runoff (Katz, 2013, p. 255). A week after the election, the Government of Haiti requested

that the OAS review the tally sheets and verify the results (Organization of American States, 2011a, p. 24). After its investigation, the OAS recommended that the CEP discard another 234 tally sheets consisting of nearly 51,000 votes (4.7% of the counted votes). Discarding these sheets would put Célestin in third place and move Martelly into the second round against Manigat.

But the government of Haiti initially refused to implement the recommendation. The refusal created weeks of a political stalemate. Finally, the US Secretary of State Hillary Clinton visited Haiti on January 30, 2011 and told President Préval the OAS recommendations would be implemented (Katz, 2013, p. 270). The evening of Clinton's appearance, the government announced the final round of the election would be between a former First Lady (Manigat) and a celebrity with no political experience (Martelly). The second round was held on March 20, 2011. Despite removing Célestin from the election, there were similar problems in the second round, and 15.3% of tally sheets were discarded for irregular counts (Organization of American States, 2011a, p. 28). On May 15, Michel Martelly became president.

The Martelly administration did not improve confidence in democratic institutions. During Martelly's 5-year term, the country failed to hold an election at any level of government, and by 2015, the majority of legislators' terms had expired, leaving the president to rule by decree. Despite difficulties, presidential elections were held in October 2015. Since the constitution prohibits presidents from serving consecutive terms, Martelly could not run again. Instead, he picked Jovenel Moïse, a banana farmer with no political experience, as the candidate for his party. Although 54 candidates ran, the principal contender was again Jude Célestin.

The 2015 election had similar problems as 2010. After the first round, there were many reports of fraud and misconduct. A commonly cited problem was the election monitors.<sup>3</sup> The monitors were representatives from each political party and were stationed at each polling location. While the logic for the monitors was that they would prevent each other from cheating, instead it was widely reported that smaller parties were selling their monitors to larger parties. The problems with the 2015 election were so large that no one would allow the elections to advance to the second round, and in February 2016 President Martelly resigned and a provisional president took his place.

Uncertainty about the next election persisted through 2016. The provisional president's mandate ended in June, without a clear plan on when the election would be held. Then in October, Hurricane Matthew disrupted the country, killing many. Finally, after a year of uncertainty, another first-round presidential election was held in November 2016. Many of the candidates from the 2015 election had dropped, halving from 54 to 27, and Jovenel Moïse, Martelly's handpicked successor, won the election outright. Although accusations of fraud were minimal, there was a significant concern that participation was low, with only 18% voter turnout.

Did the events of 2010 affect the 2015/2016 elections? Although a causal connection has yet to be established, there is evidence that confidence in the democratic process had faltered. Between the 2015 and 2016 round of voting, Kolbe and Muggah (2016) surveyed over 2000 households about their feelings toward the election. The overwhelming favorite candidate among these households was Célestin, receiving support from 42% of respondents, and Moïse was the least preferred candidate at 4%. But only 16% of households said they had voted in the 2015 election, and only 3% said they would vote in the rerun of the first round in 2016. While the report does not report cross-tabulations of candidate preference and participation, it is clear that Célestin voters were not participating. The most popular response for non-participation was concern over fraud or lacking confidence that the vote would be counted (26% gave it as their main reason) and the second most popular response was that there was no point in voting (19%). The voters did not rationalize their responses, but when asked about hypothetical solutions to the 2015/2016 electoral crisis, some of the least popular options involved interventions by foreign states or organizations. Thus, while none of this evidence is strong enough for a causal connection between the intervention in 2010 and the low turnout in 2015/2016, there is certainly enough circumstantial evidence to suggest further investigation of whether Célestin's supporters lost faith in the democratic process.

### 3 | DATA ON HAITIAN ELECTIONS

Election data in Haiti are scarce, but in this case, we have enough data to at least investigate the intervention's medium-run effect. The data come from election results in the first round of the Haitian presidential elections in 2010 and 2015 and the repeat first-round election in 2016. Ahead of the presidential elections, there were concerns about preventing fraud and creating a legitimate election. To assuage concerns, the CEP posted vote tallies for each polling station in the country. The 2010 vote tallies, without the 312 sheets discarded by the CEP, were collected by Johnston and Weisbrot (2011), and I collected the 2015 and 2016 data directly from the CEP tallies.



Unfortunately, there are no publicly available disaggregated data for the second round of 2010 (the runoff between Manigat and Martelly) or for any election before 2010. On the one hand, the absence of data prevents me from exploring the validity of the parallel trends assumption. On the other hand, investigating how foreign interventions affect political participation almost by definition occurs in countries with scarce data. Governments that regularly track election data and make the disaggregated results publicly available have stronger electoral institutions and therefore are not a target for foreign intervention.

The vote tallies allow us to measure turnout in both elections and Célestin's vote share in 2010. Calculating turnout is challenging because we do not have the exact number of registered voters in each section. But we can proxy for the number of registered voters using the number of booths in each section because booths were assigned to sections according to the number of registered voters: one booth for every 550 registered voters. In 2010, the average section had about 16 booths, and by 2015 the average section had 19 booths, which reflects that voter registration increased between the two elections. Since booths are determined by the number of registered voters, for a section with  $N$  booths, we know there are at least  $(N - 1) \times 550$  registered voters. The complication is in the last booth, which could represent anywhere between 1 and 550 voters. Because the data do not allow me to identify the marginal booth, I assume all booths have 550 registered voters. This assumption's robustness is explored in Table A3 of Supporting Information S2.

Summary statistics for turnout and Célestin's vote share are reported in Table 1.<sup>4</sup> There are 614 sections with turnout data in 2010: 609 also have turnout data in both 2015 and 2016, while 2 only have 2015 and 3 only have 2016 data. In 2010, Célestin's average vote share was 26%; average turnout was also 26%. While the 2010 turnout should have been depressed by the post-earthquake recovery environment, the average turnout across sections remained constant in 2015 at 26%. But average turnout in 2016 dropped to 19%, widely attributed to voter fatigue after a year of uncertainty about elections. Since the empirical strategy measures disenfranchisement through Célestin's vote share, it is helpful to split the data by Célestin's vote share, comparing the top quartile of his support with the bottom. In each election, Célestin's top quartile of 2010 sections had higher (statistically significant) turnout than the bottom quartile. Across elections, turnout changed differentially by quartile of Célestin support. From 2010 to 2015, we see that sections in the bottom quartile increased their turnout (23% to 25%) but that the top quartile's turnout declined (29% to 27%). In 2016, turnout had declined relative to 2010, but in the bottom quartile it had dropped only 5 percentage points, whereas in the top quartile it declined by 9. Thus, there is evidence that voter behavior responded differently across elections, though the mechanism behind the difference is yet unclear.

We can look at which sections supported Célestin using a 2009 amenities census. In 2009, the Ministry of Agriculture, Natural Resources and Rural Development surveyed sections to assess the amenities available around the country. Such amenities include the types of school available (elementary, secondary, and technical), the presence of state services (post office and court), access to information technology (presence of Internet cafe and cell phone coverage), and other amenities (sanitation facilities, gas station, and recreation center). Table 1 shows that of the 12 amenities measured, sections in Célestin's top quartile of support were less likely to have each one than the sections in the bottom quartile, with 8 of the 12 differences being statistically significant at the 5% level or stronger. This strongly suggests that Célestin's support came from poorer voters. This is unsurprising since Célestin was running under the Unity party, whose founder was then-president René Préval, who owed his 2006 victory to large turnout in poor areas (Dupuy, 2006, p. 169).

Since there is a clear difference in the characteristics of sections that supported Célestin in 2010, we need to pause and address how this will affect the analysis. Ideally, the analysis would account for the initial differences and control for changes over the 5 years. But this is Haiti, a country where data are scarce. While Haiti has a national statistics office, before the earthquake it did not regularly collect data, and it certainly did not collect much data in the chaotic period following the earthquake. Even accounting for population changes is difficult because the population projections were based on assumptions of constant population growth (see Section A.4 in Supporting Information S2). Thus, the analysis cannot account for changes in section characteristics over time. Nevertheless, Haiti was not a dynamic economy. Of course it experienced a significant shock with the 2010 earthquake, but this survey measured amenities in 2009, so the differences in characteristics predated the earthquake. After the earthquake, Haiti's economy stagnated, with GDP per capita growing an average of 0.5% per year (in contrast with the Dominican Republic, its island neighbor, which during the same year grew 4.2% per year). Thus, the analysis assumes that the initial differences between sections are absorbed in section fixed effects and that there was no significant divergence in these conditions between the elections. The reliance on this assumption is another reason I am cautious about claiming the analysis is recovering a causal effect.

TABLE 1 Summary statistics on election turnout, Célestin's vote share, and section characteristics

	All sections	Bottom quartile	Top quartile	Difference
A. Turnout				
Célestin 2010 vote share	0.26 [0.17]	- -	- -	- -
Turnout 2010 ( $N = 614$ )	0.26 [0.10]	0.23 [0.10]	0.29 [0.11]	0.06*** [0.012]
Turnout 2015 ( $N = 611$ )	0.26 [0.07]	0.25 [0.07]	0.27 [0.07]	0.02*** [0.0081]
Turnout 2016 ( $N = 612$ )	0.19 [0.08]	0.18 [0.08]	0.20 [0.08]	0.02** [0.0087]
B. Section characteristics (2009)				
School—Elementary	0.95 [0.21]	0.97 [0.18]	0.94 [0.25]	-0.030 [0.028]
School—Secondary	0.51 [0.50]	0.62 [0.49]	0.34 [0.47]	-0.28*** [0.062]
School—Technical	0.15 [0.36]	0.30 [0.46]	0.06 [0.25]	-0.24*** [0.047]
Post office	0.03 [0.16]	0.05 [0.22]	0.01 [0.09]	-0.044** [0.022]
Court	0.10 [0.29]	0.13 [0.34]	0.07 [0.26]	-0.057 [0.039]
Internet cafe	0.15 [0.36]	0.32 [0.47]	0.08 [0.27]	-0.24*** [0.049]
Cell coverage—Total	0.19 [0.39]	0.23 [0.42]	0.15 [0.36]	-0.081 [0.051]
Cell coverage—Partial	0.64 [0.48]	0.66 [0.47]	0.58 [0.50]	-0.088 [0.063]
Pharmacy	0.24 [0.43]	0.39 [0.49]	0.12 [0.33]	-0.27*** [0.053]
Sanitation	0.45 [0.50]	0.57 [0.50]	0.28 [0.45]	-0.29*** [0.061]
Gas station	0.08 [0.27]	0.20 [0.40]	0.02 [0.15]	-0.17*** [0.039]
Recreation	0.13 [0.34]	0.22 [0.41]	0.07 [0.26]	-0.14*** [0.044]
$N$	483	116	125	-

Note: The columns with “quartile” in the heading refer to the quartiles of Célestin's 2010 vote share. The bottom quartile is sections where Célestin won less than 12.5% of the vote, and the top quartile includes sections where he won more than 35.5%. In the Difference column, standard errors are in brackets, for all other columns standard deviations are in brackets.

\*\*\* $p < .01$ ; \*\* $p < .05$ .

There was one section characteristic, however, that changed during this period, and probably influenced the election: the cholera rate. After the earthquake, cholera was introduced to the country by Nepalese UN workers (Katz, 2013, Chap. 11). Because of Haiti's poor sanitation infrastructure, the cholera spread through the country, but it of course had differential effects. Controlling for cholera is important because it evolved into a contentious political issue and could spark an electoral response. Thus, all regressions control for the district's cholera rate (note: districts are made of several sections).

#### 4 | EXAMINING THE CONNECTION BETWEEN THE INTERVENTION AND TURNOUT

Although it was a political travesty, several aspects of the 2010 presidential election provide a convenient environment for testing whether foreign intervention hurt democratic institutions. According to Shulman and Bloom (2012), foreign interventions in elections are most likely to trigger pushback when they are salient, partisan, and directed by the state rather than a non-state actor. Haiti in 2010 met all of these conditions: the intervention was ultimately effected by the US Secretary of State (state actor) in an openly acknowledged process (salient) that excluded a particular candidate (partisan).

In addition to meeting the theoretical conditions, there are other reasons why this election is appropriate for this question. First, even though Haitian elections have a history of foreign influence, this intervention was more explicit than previous elections. A good contrast for the 2010 intervention was the previous election's intervention. In 2006, the international community assured a quick victory for Préval by exploiting a legal ambiguity about counting blank votes. The intervention was subtle and could easily go unnoticed. In 2010–2011, however, the international community had delayed the second round for months calling for Célestin's removal, which finally happened the same day as a high-profile visit from the US Secretary of State. Second, the intervention unambiguously changed who became president. Again, 2006 is a good contrast. While assuring Préval's victory required a slight push, he was the clear favorite and no one was surprised by the outcome. In 2011, however, removing Célestin allowed Martelly to advance against Manigat. If Manigat had won, one could always debate whether the intervention affected the outcome. But Martelly won, which would not have been possible without the intervention.

A second convenience is that the intervention discarded legitimate votes. While the intervention was rationalized as a way to remove fraud from the election, the implementation was crude. If a tally sheet was deemed "irregular," every vote on the tally sheet was discarded. Rosnick (2011) gives a simple example of the harm this could cause: suppose the true votes on a tally sheet are 120 for Célestin and 50 for Manigat, but that the sheet reads 180 for Célestin and 50 for Manigat. The extra 60 votes are fraudulent. While a true correction would eliminate those 60 votes and properly represent the section's preferences, discarding the entire tally sheet trades an error of +60 votes for Célestin for a net error of +70 votes for Manigat, which is a 130 vote swing in that section. Rosnick then uses a simulation of the discarded tally sheets to show that Célestin's legitimate votes likely gave him an indisputable lead over Martelly. The intervention not only directly disenfranchised voters by discarding their votes, it indirectly disenfranchised all voters who had voted for Célestin.

Another convenience of this setting is that Célestin ran in both elections. Having Célestin in both elections is nice because it keeps the personality constant. Personality drives Haitian politics because there are no stable political coalitions that attract consistent support. In each election, dozens of candidates ran for president across just as many parties, and between elections Célestin even switched political parties. Having Célestin, the target of 2010s intervention, run in both elections means that we can control for the personality and be more confident that the effects are driven by the intervention.

To investigate the connection between the intervention and voter behavior, I examine turnout across the two elections using Célestin's vote share to measure intensity of treatment. I use a modified difference-in-differences strategy similar to Jones et al. (2017) and estimate the following regression:

$$y_{st} = \beta_1 \text{Disenfranchisement}_s \times \text{PostInterference}_t + \delta_t + \delta_s + \beta_2 X_{st} + \varepsilon_{st}. \quad (1)$$

The dependent variable  $y_{st}$  includes election outcomes such as turnout or registered voters in section  $s$  in election  $t \in \{2010, 2015, 2016\}$ .  $\text{Disenfranchisement}_s$  measures the degree of disenfranchisement in section  $s$ , which I proxy with



Célestin's 2010 vote share. The dummy variable  $\text{PostInterference}_i$  indicates whether it is the 2015 or 2016 election. The regression includes election and section fixed effects ( $\delta_t$  and  $\delta_s$ , respectively) and controls for the cholera rate ( $X_{st}$ ).

Predicting the sign of  $\beta_1$  is difficult because the relationship between disenfranchisement and turnout is theoretically ambiguous. According to reactive theory, disenfranchisement may cause voters to cherish their voting rights, increasing turnout in the next election (Biggers & Smith, 2020). For instance, when the Voting Rights Act of 1965 reversed the disenfranchisement of Blacks in the South, Black voters registered to vote and turned out in high numbers (Cascio & Washington, 2014). On the other hand, the intervention could discourage Célestin's supporters, fracturing their fragile faith in the democratic system. Such voters may question the purpose of participating in the election, especially when inefficient polling stations and dangerous election-day conditions significantly increase the cost of voting. In this case, we would observe a decrease in turnout.

For  $\beta_1$  to be a causal effect, we must assume that Célestin's 2010 vote share did not affect changes in turnout within a section across elections except through the response to the intervention. One way to frame the identifying assumption is to say that in a world where Célestin was not dismissed from the 2010 election we do not expect turnout to change differently within a section that voted 40% in favor of Célestin than in a section where only 10% of voters supported him. Of course, a section that voted 10% for Célestin must have given more of its votes to another candidate, so another way to frame the identifying assumption is to say we do not expect turnout to change differently within a section that voted 40% in favor of Célestin than in a section where 40% voted in favor of Manigat or Martelly.

The causal assumption's validity rests on how well the analysis controls for any other confounding factors. Unfortunately, it is hard to eliminate every confounding effect. For example, the analysis assumes parallel trends, but because detailed voting data is not available before 2010, I cannot provide evidence in favor of this assumption. The analysis also assumes that the voters who were disenfranchised in 2010–2011 were still in their sections in 2015–2016. If voters move to other sections but retain the consequences of their response to the intervention, then the effects will spill over into other sections. Unfortunately, we do not know whether migration is a problem or not. Haiti's last census was in 2003 and there are no regular surveys of migration. Nevertheless, the effect on the analysis should be minor. As long as the migration destination is orthogonal to the Célestin's 2010 vote share in the origin, this means there is error in measuring the treatment status, which biases the estimates of  $\beta_1$  toward zero.

But another threat to the analysis is fraud. Even though the OAS expressed confidence in the election, there is still concern that there is fraud present in the data. Most importantly, if fraud in 2010 was correlated with Célestin's vote share—an easy case to make since he was removed on allegations of fraud—and if that fraud was eliminated in 2015—also reasonable since his party was no longer in power—then the analysis would produce a negative coefficient.

Since accusations of fraud tainted the 2010 election, it is worth exploring the evidence for fraud, its potential effect on the analysis, and how to handle it.

## 4.1 | The extent of fraud in the 2010 election

Before even analyzing the data, we already have some protections against fraud affecting the analysis. First, the CEP and OAS reportedly eliminated the most egregious examples of fraud. The remaining sheets, which are the only ones used in this analysis, passed scrutiny. Thus, while we cannot guarantee the remaining ballots are free of fraud, the fraud was noisy enough to avoid detection. Of course, one concern is that the CEP, which was formed by the incumbent party, could have systematically excluded ballots in a way that favored the incumbent candidate. We will investigate this possibility in Section 6.3.

Along these lines, a significant threat to the analysis is whether fraud is correlated with treatment. Because Célestin was eventually removed from the election, it seems that he was the main perpetrator of fraud. But while Célestin was under the most scrutiny, and while some cheating on election day was observed, there was no conclusive evidence that he perpetuated the cheating nor evidence that the other candidates did not cheat. The OAS reports never attribute fraud to any party nor exempt any parties from fraud. In fact, as mentioned above, more “irregular” sheets were discarded in the second round, when it was only Manigat and Martelly, than were discarded in the first. Thus, while fraud likely exists, it is possible that all candidates received fraudulent votes.

Even though there are reasons to believe fraud was resolved or that it is uncorrelated with treatment, there are still two important issues to consider. First, many ballots were missing or not included in the vote tallies—how should such cases be handled? Second, if fraud remains, do the election results have any meaning?

Resolving the concern about missing ballots is straightforward. The CEP reports how many ballots were missing or discarded in each election, and both elections had problems—in both 2010 and 2015, 30% of sections had at least one missing ballot. Although in both elections the percentage was equal, it was uncommon for the same section to have missing ballots in both years, as seen in Figure A1 of Supporting Information S2. If a ballot is missing, we can still estimate turnout and Célestin's 2010 vote share from the other ballots in the section. But missing ballots decrease our confidence in the measure's accuracy. To address the concern, regressions are weighted by the fraction of ballots recovered—we weight an observation more if we have 100% of its ballots than if we have only 50%. For robustness, unweighted results are reported in Table A2 of Supporting Information S2.

The concern about whether the fraud destroys any meaning behind the results is more difficult. Unfortunately, it is impossible to identify election fraud relying on just election results; however, it is possible to analyze such results and see if the patterns match those predicted by models of election fraud. In what follows, I present evidence from pre-election surveys and election patterns to argue that nothing in the data suggests fraud confounds the analysis. But let me be precise about how to interpret the analysis: I am making an argument about fraud's effect on the analysis, not its effect on democracy. Fraud was clearly an issue in the elections, and such concerns should be taken seriously because they affect the election's integrity. My argument is that to the extent which it did occur, fraud did not significantly skew the data and can be interpreted as noise, obscuring the precision but not the estimates. This is consistent with the ballots passing the CEP's inspection, as described above.

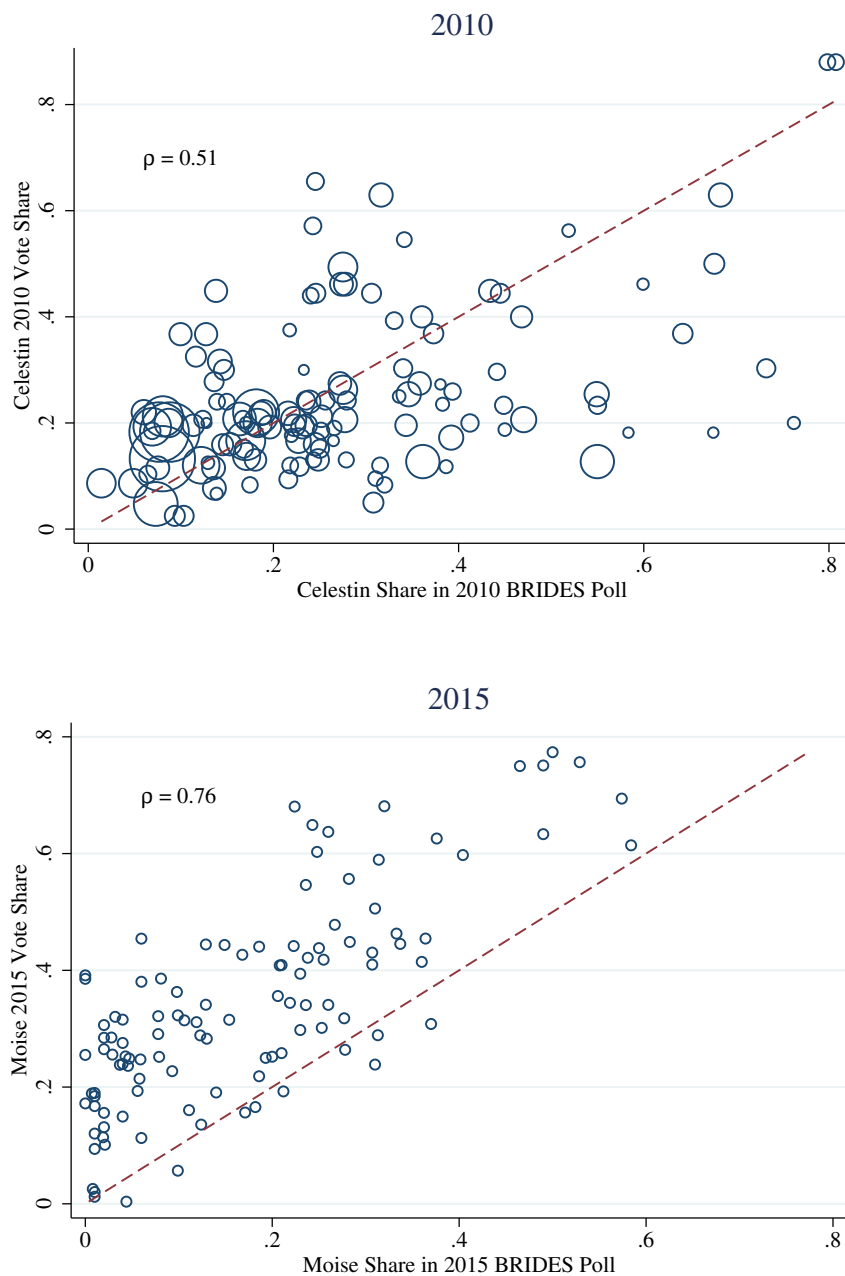
Evidence from pre-election surveys suggests voting results reflected the district's preferences. In both 2010 and 2015, pre-election polling was conducted by the Bureau de Recherche en Informatique et en Développement Economique et Social (BRIDES) at the district-level (which is one step of geographical aggregation above sections, the unit of analysis in the regressions). The BRIDES survey is the only source of data outside of election results that measures local political preferences, but it is not perfect. Although the BRIDES survey allegedly followed sound survey practices, it is clear from the demographic data presented in Table A1 of Supporting Information S2 that the respondents do not represent the country's population (though, in defense of BRIDES, it might be a fair representation of the people who voted). Furthermore, the sample size is small: in 2010, the average number of respondents per district was 57, and in 2015, it was about 100, leaving significant room for sampling error. Nevertheless, if we want to validate the election results with an outside source, the BRIDES survey is the only benchmark available.

Despite the survey's weaknesses, the district-level results correspond favorably to the election results. Figure 1 shows scatter plots of the election results against the survey results for both 2010 and 2015, and there is a clear relationship between them. The correlation is stronger in 2015 ( $\rho = 0.76$ ) than in 2010 ( $\rho = 0.51$ ), but this is expected with the difference in sample sizes. Although there is room to question the independent validity of the BRIDES survey and the election results, their strong correlation suggests that the election results reflect the district's preferences.

While that test used external data to validate patterns in the results, another test looks for fraud indicators just using patterns in the election data. Klimek et al. (2012) suggest looking at the election fingerprint, which is a plot of the suspicious party's vote share against the voter turnout at that location.<sup>5</sup> Empirically, elections without fraud tend to have data points that cluster around a given turnout and vote percentage, but elections with fraud see a smear from that cluster up to the top right corner. The intuition is simple: if someone illegally adds 100 votes for a candidate, that will increase the candidate's vote share, but it will also increase the measured turnout because it looks like 100 more voters participated. Thus, sections with both anomalously high turnout and high support for the winner are likely locations of ballot stuffing. Indeed, in Haiti, one criterion the OAS used for examining tally sheets for irregular voting patterns was whether the sheet had over 50% turnout (Organization of American States, 2011a, p. 98). While this is a visual test without test statistics, it provides a simple, intuitive check for ballot stuffing.

Figure 2 plots the fingerprint for four presidential elections: Haiti's 2010 and 2015, Spain's 2008, and Russia's 2012. In Russia's 2012 election, practically no booths fell into the lower triangle of low turnout and low support, but there is a distinct smear to the northeast corner of the scatter plot. Klimek et al. (2012) conclude that the smear is compelling evidence of fraud in Russia's election as ballot stuffers pushed up both Putin's share and the observed turnout. Spain, on the other hand, is shown as an example of a fingerprint that displays no indication of fraud: turnout is high, but there are few booths in the upper triangle of both high turnout and high support for the winner. For Haiti, the two fingerprints appear much more like Spain than Russia, with no booths reporting high turnout. Even conditional on low turnout numbers, very few booths report high support for Célestin.

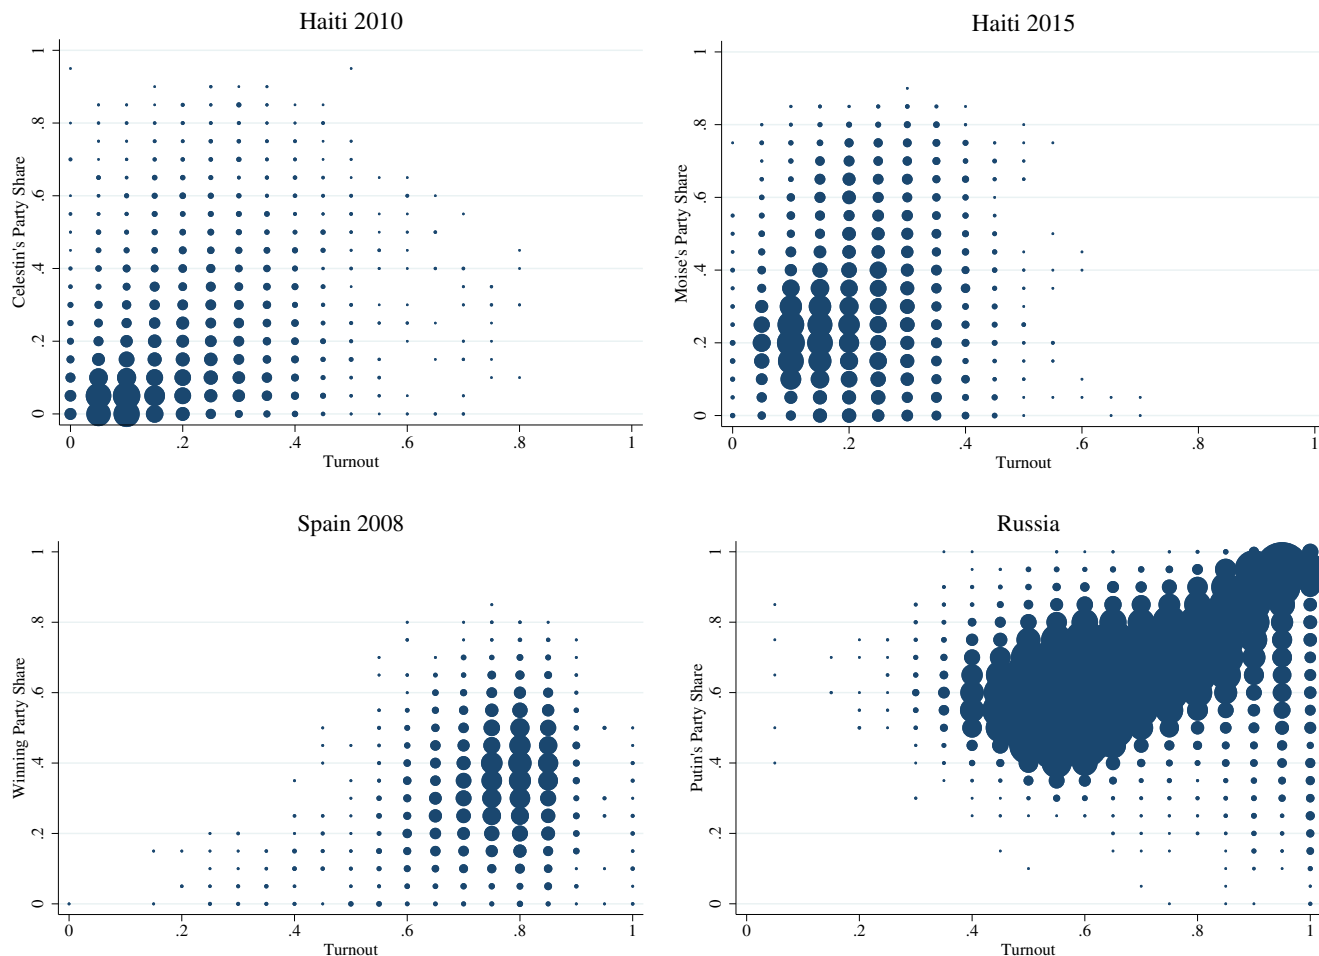
The data do not raise any red flags for fraud and are suitable for analysis. The findings are consistent with the CEP filtering suspicious ballots and presenting only the results it trusted. As mentioned above, the analysis cannot



**FIGURE 1** Relationship between pre-election polls and 2010/2015 election results by district. Data come from pre-election surveys conducted by Bureau de Recherche en Informatique et en Developpement Economique et Social (BRIDES) in October of the election year. Each dot represents a district. The 2010 plot weights districts according to sample size, but in 2015 the sample size was the same for every district. The top-left corner of each plot reports the correlation coefficient ( $\rho$ )

unconditionally reject the presence of fraud or determine the effect of fraud on the election results. Furthermore, the analysis cannot reject interventions with marginal impacts on vote counts, such as vote buying or voter intimidation. But the data appear to accurately reflect district preferences, and fraud might add some noise to those measures. In a regression analysis, any fraud is treated as measurement error.

Although the analysis in this section has increased our confidence in proceeding with the analysis, there are still concerns that election tampering may bias the main results. Thus, I address these concerns with three empirical exercises. The first and second rely on observations from the OAS election monitoring report. First, I assume that the OAS report accurately identified that the problems with fraud were concentrated in Port-au-Prince and that the provinces' problems were not with fraudulent voting. Under this assumption, the easy test for the



**FIGURE 2** Election fingerprints for presidential elections in Spain (2008), Russia (2012), and Haiti (2010, 2015). Election fingerprints let you see if any sections had unusually high turnout and unusually high support for the winner. Russia is an example where booths with high turnout also often had a high share of votes going to Putin, evidence that ballot stuffing influenced the election. Dots represent cells of turnout and vote share, and dots are weighted by the number of sections in each cell. Data for Spain and Russia come from Klimek et al. (2012)

results' robustness is to remove observations from Port-au-Prince and the surrounding areas.<sup>6</sup> Second, I use the OAS observations to design a sensitivity test. The report claims the problems were concentrated in a few areas. The question is, how many sections experienced problems? And for such sections, how many of Célestin's votes were fraudulent? These questions establish a clear foundation for a sensitivity test that examines how robust the results are to perturbations in the data. Third, I examine how sensitive the results are to assumptions about the voting behavior on the sheets discarded by the CEP and OAS.

## 5 | RESULTS

Table 2 reports the results from estimating Equation (1). I use two forms of the dependent variable: raw turnout and the log-transformed turnout. The benefit of the log transformation is that we can gain insight into the mechanism by decomposing the result into the effect on registration and on votes, both included in the same table. To account for serial correlation across elections, standard errors are clustered at the section level, and because turnout on missing ballots have to be imputed, the regression gives more weight to sections with fewer missing 2010 ballots (unweighted regressions, reported in Table A2 of Supporting Information S2, give similar results).

The results show a large effect on turnout. The results are negative and statistically significant at the 1% level, suggesting a discouragement effect. A 10% increase in support for Célestin in 2010 is associated with a 3% drop in

TABLE 2 The 2010 intervention and turnout in Haiti's 2015 and 2016 presidential elections

	Turnout	ln(Turnout)	ln(Registered)	ln(Votes)
Célestin Share in 2010 × Post Interference	−0.069*** [0.018]	−0.31*** [0.078]	−0.073 [0.057]	−0.38*** [0.099]
2015 Election	−0.0099 [0.0084]	0.013 [0.035]	0.16*** [0.028]	0.18*** [0.046]
2016 Election	−0.074*** [0.0087]	−0.31*** [0.037]	−0.056* [0.029]	−0.37*** [0.049]
$\sinh^{-1}(\text{Cholera Rate}) \times 2015 \text{ Election}$	0.0075*** [0.0015]	0.026*** [0.0061]	0.0038 [0.0050]	0.030*** [0.0080]
R-squared	.305	.339	.245	.452

Note: All regressions have data on 614 sections. For 609 sections, there are 3 years of data, and for the remaining 5 there are only 2 years. Standard errors are clustered at the section level. Because the votes on missing ballots have to be imputed, regressions are weighted by the fraction of the section's ballots recovered.

\* $p < .10$ ; \*\*\* $p < .01$ .

turnout in 2015. Relative to a section where Célestin won 20% of the 2010 vote, a section where Célestin won 70% of the vote dropped its 2015/2016 turnout by 15 log points.

Where is the effect coming from: registration or votes? Decomposing the effect is important because it clarifies the mechanism of retribution versus discouragement. A reactive effect could increase voter registration and total votes, but if registration increased more than votes then  $\hat{\beta}_1$  would be less than zero. On the other hand, a discouragement effect would depress both registration and votes. Table 2 looks at registration and votes separately. Célestin's 2010 districts had significantly fewer votes in 2015 and 2016, and while there was not a significant difference in voter registration, the point estimate suggests it was also negative. There seems to be no evidence for a reactive effect, but negative effects on votes and registration is consistent with a discouragement effect.

The analysis in Table 2 pools the 2015 and 2016 elections, which may hide important heterogeneity between the two elections. There is not a strong reason to use one or the other since both exhibited anomalies: the 2015 election was fraught with accusations of fraud, but the 2016 election had exceptionally low turnout. Additionally, the 2016 election was preceded by Hurricane Matthew, which devastated the southern peninsula and possibly deterred turnout. In Table 3, I explore the heterogeneity by disaggregating the samples, running the analysis separately for the 2015 and 2016 elections. The results are similar across both samples. The consistency is not surprising: voters who were discouraged in 2015 had no reason to believe 2016 would be any better.

## 5.1 | The Manigat falsification

To validate and explore the results, I perform a falsification test using Manigat as the treatment. Since Manigat was not dismissed from the election, her supporters should show neither retribution nor discouragement. Thus, I estimate Equation (1) using Manigat's 2010 vote share in place of Célestin's.

Table 4 shows the results. Manigat's share in 2010 is unrelated to turnout, registration, and total votes. As predicted, there is no evidence for discouragement or retribution from Manigat's voters. Like with the main results, I report the disaggregated analysis in Table A6 of Supporting Information S2 and find no heterogeneity across elections.

## 5.2 | Turnout and discarded tally sheets

While the analysis has so far focused on the ultimate decision to remove Célestin from the second round, the 2010 election might have discouraged voters because of the votes that were discarded. When the CEP counted the votes, it implemented a blunt solution: any tally sheet that looked suspicious was discarded entirely. That means the CEP discarded legitimate votes. Victims of this filter may feel specifically targeted and, therefore, abstain from voting in the



TABLE 3 The 2010 intervention and turnout in Haiti's 2015 and 2016 presidential elections, disaggregated

	Turnout	ln(Turnout)	ln(Registered)	ln(Votes)
2015 Election				
Célestin Share in 2010 × 2015 Election	−0.063*** [0.019]	−0.30*** [0.077]	−0.067 [0.066]	−0.37*** [0.11]
2015 Election	0.002 [0.0089]	0.049 [0.035]	0.15*** [0.032]	0.20*** [0.048]
$\sinh^{-1}(\text{Cholera Rate}) \times 2015 \text{ Election}$	0.0045*** [0.0016]	0.018*** [0.0059]	0.0058 [0.0056]	0.023*** [0.0083]
R-squared	.032	.059	.281	.236
2016 Election				
Célestin Share in 2010 × 2016 Election	−0.075*** [0.021]	−0.32*** [0.099]	−0.081 [0.069]	−0.40*** [0.12]
2016 Election	−0.085*** [0.010]	−0.35*** [0.047]	−0.045 [0.033]	−0.39*** [0.057]
$\sinh^{-1}(\text{Cholera Rate}) \times 2016 \text{ Election}$	0.010*** [0.0018]	0.035*** [0.0081]	0.0019 [0.0058]	0.037*** [0.0098]
R-squared	.335	.330	.034	.322

Note: All regressions have data on 614 sections. For 609 sections, there are 3 years of data, and for the remaining 5 there are only 2 years. Standard errors are clustered at the section level. Because the votes on missing ballots have to be imputed, regressions are weighted by the fraction of the section's ballots recovered.

\*\*\* $p < .01$ .

TABLE 4 Using Manigat's 2010 vote share as a falsification test

	Turnout	ln(Turnout)	ln(Registered)	ln(Votes)
Manigat Share in 2010 × Post Interference	0.0065 [0.016]	−0.058 [0.065]	−0.031 [0.047]	−0.089 [0.080]
2015 Election	−0.028*** [0.0088]	−0.043 [0.034]	0.16*** [0.028]	0.11** [0.047]
2016 Election	−0.092*** [0.0087]	−0.37*** [0.035]	−0.064** [0.028]	−0.43*** [0.046]
$\sinh^{-1}(\text{Cholera Rate}) \times 2015 \text{ Election}$	0.0071*** [0.0016]	0.025*** [0.0061]	0.0037 [0.0050]	0.029*** [0.0081]
R-squared	.345	.38	.201	.445

Note: All regressions have data on 614 sections. For 609 sections, there are 3 years of data, and for the remaining 5 there are only 2 years. Standard errors are clustered at the section level. Because the votes on missing ballots have to be imputed, regressions give more weight to sections with fewer missing ballots.

\*\*\* $p < .01$ ; \*\* $p < .05$ .

future. Since discarded tally sheets came from sections that favored Célestin (Johnston & Weisbrot, 2011), there might be a confounding effect between the CEP filter and the foreign intervention.

To address this confounding factor, I introduce a control for discarded tally sheets. I try two measures: first, an indicator for whether the section had any tally sheets discarded (13% had at least one discarded), and a continuous measure of the fraction of tally sheets discarded (the average section was had 3% of sheets discarded). Table 5 shows how the control affects the results. In the first column, I replicate the main results from Table 2 for comparison. Columns 2 and 3 use the measures of discarded ballots without controlling for Célestin's 2010 vote share. When the

TABLE 5 Controlling for tally sheets discarded by CEP

	(1)	(2)	(3)	(4)	(5)
Panel A. Turnout					
Célestin Share in 2010 × Post Interference	−0.069*** [0.018]	-	-	−0.057*** [0.018]	−0.054*** [0.018]
Any Discarded × Post Interference	-	−0.053*** [0.017]	-	−0.048*** [0.017]	-
Frac. Discarded × Post Interference	-	-	−0.21** [0.084]	-	−0.18** [0.086]
R-squared	.305	.314	.312	.32	.317
Panel B. ln(Turnout)					
Célestin Share in 2010 × Post Interference	−0.31*** [0.078]	-	-	−0.29*** [0.079]	−0.30*** [0.081]
Any Discarded × Post Interference	-	−0.09 [0.058]	-	−0.067 [0.059]	-
Frac. Discarded × Post Interference	-	-	−0.28 [0.30]	-	−0.15 [0.31]
R-squared	.339	.332	.331	.34	.339

Note: All regressions have data on 614 sections. For 609 sections, there are 3 years of data, and for the remaining 5 there are only 2 years. Standard errors are clustered at the section level. Because the votes on missing ballots have to be imputed, regressions are weighted by the fraction of the section's ballots recovered.

Abbreviation: CEP, Provisional Electoral Council.

\*\*\* $p < .01$ ; \*\* $p < .05$ .

dependent variable is raw turnout, there is a negative, statistically significant relationship between discarded tally sheets and turnout, but the significance disappears when using the log transformation of turnout. Columns 4 and 5 include both the Célestin share and the discarded sheets, and under all specifications the Célestin coefficient remains large, negative, and statistically significant. The evidence in Table 5 suggests that discarding the tally sheets may have discouraged voters, but the effect was separate from the effect of the intervention.

## 6 | ROBUSTNESS OF RESULTS

The analysis so far has suggested that the intervention in Haiti's 2010 presidential election is associated with a decline in turnout in the 2015/2016 elections. But a barrier to fully accepting these results is the concern about fraud. As shown above, the concerns about fraud have been minimized because the CEP discarded the most concerning tally sheets and the remaining data display no indicators of fraudulent behavior. Yet since we know there was fraud, we should still be cautious about whether the results are robust to accounting for fraud.

The motivation for the robustness tests comes from a report written by the OAS. In 2009, before the earthquake even hit, the government of Haiti had invited the OAS to monitor the election. Such requests are common, with the OAS observing more than 187 elections from 1960 to 2010 (Organization of American States, 2011a, p. 9). The first members of the observation team arrived on August 3, 2010, almost three months before the first round of the election was held, and stayed until May 18, 2011, two months after the second round (p. 8). Thus, OAS had the most informed outsider opinion on fraudulent behavior.

After describing the disorder in the election, the OAS report states, “the problems seemed to be concentrated in a few regions, including *Port-au-Prince*” [emphasis added] (Organization of American States, 2011b). The report continues, “For the most part the voting in the provinces appeared to have been unfolding relatively smoothly though similar irregularities had been observed (late opening of polling stations, voters having difficulty in finding their polling

stations or finding their names on the partial electoral list inside the polling station)” [emphasis added]. Two key insights from the report are, first, election problems were concentrated in Port-au-Prince and the surrounding areas, with the provinces operating relatively well; and, second, election problems were isolated to a few regions. The two insights motivate the first two robustness checks below. The first is to drop the areas with the biggest concern for fraud. The second is to test the sensitivity of the results by perturbing the data.

The third test explores the sensitivity of the results to questions about the CEP’s motives. So far, we have assumed that the CEP was unbiased in discarding sheets, eliminating irregular sheets regardless of the party it favored. But since the CEP was appointed by Préval, we might be worried the discarded sheets were meant to benefit his candidate, Célestin. Thus, the third test checks the sensitivity of the results to assumptions about voting behavior on the discarded sheets.

## 6.1 | Robustness to dropping Port-au-Prince and surrounding areas

The first robustness check excludes the most problematic areas. The OAS report identified Port-au-Prince and the surrounding areas as the areas with the most election-day problems. Thus, I run the same specification but drop the all sections in the following districts: Port-au-Prince, Carrefour, Cite Soleil, Delmas, Gressier, Kenscoff, Petionville, and Tabarre. Table 6, shows the results are nearly identical and are robust to dropping those districts.

Beyond the concern about fraud, this robustness check also addresses concerns about how the earthquake might have affected turnout across elections. These districts experienced the most damage from the earthquake. Since the 2010 election happened after the earthquake, the section fixed effects should absorb any direct effects from the earthquake. There might have been, however, some dynamic effects during the recovery that the fixed effects miss. But dropping these districts does not affect the main coefficient, so it appears that any recovery effects were independent.

## 6.2 | Sensitivity check

The second test is a sensitivity check to see how robust the results are to perturbations in the data. The goal is to see how much fraud would have to have existed in the 2010 election for the decrease in turnout to come from eliminating fraud rather than from voters discouraged from participating. The main advantage of this test is to be as transparent as possible about the conditions needed to believe the results.

The sensitivity test begins with a thought exercise. Suppose that the only fraud left in the filtered data (i.e. after the CEP removed irregular tally sheets) is one section where 50% of Célestin’s 2010 votes were fraudulent.

TABLE 6 Testing the robustness of the results to dropping Port-au-Prince and surrounding areas

	Turnout	ln(Turnout)	ln(Registered)	ln(Votes)
Célestin Share in 2010 × Post Interference	−0.060*** [0.018]	−0.23*** [0.075]	−0.063 [0.059]	−0.29*** [0.095]
2015 Election	−0.012 [0.0084]	−0.01 [0.034]	0.16*** [0.028]	0.15*** [0.045]
2016 Election	−0.076*** [0.0087]	−0.33*** [0.036]	−0.065** [0.029]	−0.39*** [0.048]
$\sinh^{-1}(\text{Cholera Rate}) \times 2015 \text{ Election}$	0.0072*** [0.0016]	0.025*** [0.0060]	0.0056 [0.0050]	0.030*** [0.0079]
R-squared	.354	.399	.211	.465

Note: All regressions have 3 years of data on 581 sections. The following districts are omitted: Port-au-Prince, Carrefour, Cite Soleil, Delmas, Gressier, Kenscoff, Petionville, and Tabarre. Standard errors are clustered at the section level. Because the votes on missing ballots have to be imputed, regressions give more weight to sections with fewer missing ballots.

\*\*\* $p < .01$ ; \*\* $p < .05$ .

Would eliminating that fraud be enough to overturn the results? The results so far suggest that even if one section had egregious fraud, the overall pattern that turnout decreased among Célestin voters would survive. But what if 50% of Célestin's 2010 votes in every section were fraudulent? Eliminating that much fraud in the next election would be good for democracy but it could also look like turnout had decreased. So the question is: under what patterns of fraud do the results still hold?

The sensitivity test is designed to answer this question. It has two parameters: (1) the share of sections in 2010 with fraud and (2) the share of Célestin's votes that were fraudulent. The test then runs rounds for every combination of the two parameters between 0 and 1, at intervals of 5 percentage points. For instance, the test has one round where 10% of sections experienced fraud, and in those sections 5% of Célestin's votes were fraudulent. In each round, the test has four steps. First, the sections with fraud are selected. Since we do not know which sections had fraud, I perform the test twice with different assumptions about the ballot-stuffing strategy. One strategy is to stuff ballots in areas where Célestin is weak. Under this strategy, the vote share in these sections would look closer to sections with strong support, and therefore there would be no correlation between fraud and vote share. Thus, the first test assumes no correlation between fraud and Célestin's 2010 share. If the round says 10% of sections experienced fraud, then 10% of sections are selected at random. But another strategy is to target sections that already support the candidate, which would induce a positive correlation between fraud and vote share. Thus, the second test assumes a perfect correlation between fraud and Célestin's vote share. If the round assumes 10% of sections experienced fraud, the assumption is that it is the sections at the top 10% of Célestin's performance.

The next three steps of the test are the same regardless of how the first step is performed. Second, in the selected sections, Célestin's votes are reduced by the amount of fraud in that section (for this example, 5% of his votes are removed in each section). Third, 2010 turnout is recalculated accounting for the elimination of fraud. Fourth, Equation (1) is reestimated. For each combination of parameters, the test is performed 200 times, and the result from the fourth step is saved. The primary outcome of interest is to see whether there is a statistically significant result after removing votes from the 2010 election.


Note that one of the strengths to this approach is that it is robust to *when* the ballot stuffing occurred. The fraud could happen on election day, per the OAS evidence, but it can also happen out of the public eye when the votes are counted. Indeed, the OAS noted its concerns about errors introduced during the counting (Organization of American States, 2011a, p. 24). If the ballot stuffing came from distortion by election officials, then the OAS observations would not identify the troublesome sections. Since the test assigns fraud based on the final vote counts, the timing of the ballot stuffing is irrelevant.

The results are summarized in Figure 3. The axes of Figure 3 are the two parameters: the x-axis is the share of sections with fraud and the y-axis is the share of Célestin's votes that were fraudulent. The dark gray region indicates the combinations of parameters under which turnout decreased across the two elections and the result is statistically significant at the 5% level (determined by running the simulation 200 times). The light gray region shows the parameters where the results are negative but no longer statistically significant.

The sensitivity test shows that even under large allegations of fraud, the main result holds. In the uncorrelated case, even if we believed there were sections where every Célestin vote was fraudulent, the results would still hold as long as that fraud was contained to less than 10% of sections. Or suppose that 20% of sections had ballot stuffing. Then up to 80% of Célestin's votes in those sections could be fraudulent and the intervention would still have a negative, statistically significant effect on turnout. In the second case, where Célestin's top performers were the culprits, then 20% of sections could have more than 10% of their votes fraudulent without eliminating the result.

While the sensitivity test does not eliminate the concerns of fraud, it puts very clear boundaries on what patterns of fraud one would have to argue to disbelieve the results. Note that much of the evidence has limited the chance of significant fraud in the filtered data. The OAS claimed problems were limited to a few areas. The filtered data match pre-election predictions, and the results do not raise flags of abnormal patterns. So claiming 20% of sections had significant fraud would require new evidence to support that claim.

One caveat with these results is that they are performed just with the data that made it through the CEP filter. On the one hand, that might bolster our confidence that the fraud has already been removed. In that case, the assumptions about fraud need to be amended to be "X% are fraud after the CEP already removed a significant amount of fraud." On the other hand, there are concerns that the CEP might have introduced another bias in their filtering of the tally sheets. The next section addresses that concern.



**FIGURE 3** Simulating the robustness of the results to the degree of fraud. Results from the simulation described in the text. The dark region represents the combination of parameters where the decrease in turnout across elections is negative and significant at the 5% level. The lighter region includes combinations that point estimates that are not statistically significant but still negative

### 6.3 | Discarded sheets

Although the above sensitivity check looks at whether the filtered results contain fraud, there is another source of fraud we might be worried about: the discarded tally sheets. During the vote count, the CEP discarded irregular tally sheets, but it did not explain why the sheets were deemed irregular. Since the CEP was appointed by Préval (Célestin's party leader), one concern is that the CEP discarded sheets to ensure Célestin advanced. Removing tally sheets with this systematic bias could distort the results since this would inflate Célestin's 2010 vote share.

While it is impossible to know the vote counts on the discarded sheets, we can see if the results are sensitive to a range of assumptions about the ignored votes. There are two parameters of interest: how many votes were in the



TABLE 7 Testing the robustness of the result to assumptions about discarded tally sheets

	0.05	0.1	Turnout same as section	0.9	1
Célestin vote share					
0	−0.050*** [0.017]	−0.055*** [0.017]	−0.063*** [0.017]	−0.0063 [0.020]	0.005 [0.021]
0.1	−0.049*** [0.017]	−0.055*** [0.017]	−0.067*** [0.017]	−0.040** [0.018]	−0.034* [0.019]
Same as section	−0.045** [0.017]	−0.049*** [0.017]	−0.069*** [0.018]	−0.12*** [0.021]	−0.13*** [0.023]
0.9	−0.044** [0.017]	−0.051*** [0.017]	−0.086*** [0.019]	−0.21*** [0.023]	−0.24*** [0.025]
1.0	−0.044** [0.018]	−0.050*** [0.017]	−0.087*** [0.019]	−0.22*** [0.023]	−0.24*** [0.024]

Note: Each cell is a separate regression. Each regression estimates the same specification as Table 2, but turnout and Célestin's vote share are recalculated to account for assumptions about the tally sheets discarded by the CEP. Each column represents a different assumption about turnout, and each row represents a different assumption about Célestin's vote share. Each cell reports the coefficient on Célestin's 2010 Vote Share  $\times$  Post Interference.

Abbreviation: CEP, Provisional Electoral Council.

\*\*\* $p < .01$ ; \*\* $p < .05$ ; \* $p < .10$ .

discarded sheets, and how much did the discarded sheets favor Célestin. Since we know how many sheets were missing for each section, we can impute values for the discarded sheets, recalculate the section's turnout and Célestin vote share, and reestimate the regression.

The results from the sensitivity test are shown in Table 7. The columns show different assumptions about turnout for the discarded sheets, and the rows show different assumptions about Célestin's vote share. For example, the first row assumes that the discarded sheets had no votes for Célestin, and the first column assumes turnout on the discarded sheets was only 5%. Under the combination of these two assumptions, the coefficient is still negative and statistically significant. This is true for every combination of assumptions except for the case where there was high turnout (above 90%) and no support for Célestin. But an extreme case like that would only occur from another candidate stuffing the ballots, and those ballots should rightfully be discarded.

The most reasonable assumption is that the missing sheets had the same turnout as other in the section but did not favor Célestin. Discarding those sheets would create large net gains in Célestin's vote total because it would eliminate relatively few Célestin votes at the benefit of removing relatively more votes for the opposition. The middle column of Table 7 shows what the results would look like under this strategy. Even if every discarded sheet had no Célestin votes, as long as the turnout was similar to the rest of the section's sheets, we would find almost the same effect.

With the first-round difference between Célestin and Martelly being less than 7000 votes, there was certainly an opportunity to swing the election through strategically eliminating sheets that did not favor Célestin. But this strategy would have had little effect on the changes in voting behavior between the two elections.

## 7 | CONCLUSION

This paper has sought to understand the medium-run effects of a foreign election intervention by examining voter behavior in Haiti from 2010 to 2016. While voter participation was lower in 2015 and 2016 than in 2010, I find that it declined faster in sections where there was stronger support for Célestin, the candidate removed from the 2010 election by foreign intervention. To conclude that this association is causal, we have to assume that voter behavior would have followed similar trends across sections without the intervention. Unfortunately, there are many limitations to arguing that we can assume there is a causal effect. Data limitations prevent us from examining the parallel trends assumption. And concerns about the election's integrity leave questions about confounding factors. While I try to address concerns about fraud with sensitivity tests, it is impossible to rule out every confounding

effect of fraud or election tampering (such as vote buying). Nevertheless, any investigation of the effect of foreign election intervention on voter behavior by definition will occur in an environment with these challenges. Hopefully future research will address these concerns.

There is an even bigger concern about the proper counterfactual analysis and the tension of truly understanding the medium- and long-run effects of such interventions. It is impossible to know where Haiti would be today without the 2010 intervention, but the consequences of the 2010 and 2015 elections were felt through 2021. During the earthquake recovery, the Martelly and Moïse administrations accumulated about \$2 billion in debt to Venezuela. Intended to be used to rebuild the country, the funds were instead squandered and embezzled. The scandal was exposed in November 2017 through a Senate probe with 650 pages of substantiating evidence, and among the many implicated in the scandal was President Jovenel Moïse. By August 2018, the scandal had gained enough political momentum to become a social media phenomenon, inspiring widespread protests.<sup>7</sup> The protests had become violent in February 2019 then continued through the year, disrupting and impeding most of Port-au-Prince. In January 2020, Moïse was left to rule the country by decree because parliament dissolved after failure to hold elections in 2019. And in July 2021, political uncertainty escalated with the assassination of Moïse.

Haiti's 2010 election is a reminder that the United States spends billions of dollars on democracy assistance programs. While such programs hope to promote democratic institutions, we do not understand their long-run impact. This paper has shown that when the United States intervened in Haiti's 2010 presidential election, the voters who supported the removed candidate participated less in the next election. Counterintuitively, an intervention that was ostensibly about maintaining election integrity might have instead reduced faith in elections. These findings are consistent with survey evidence that revealed the most common reasons for not voting in 2015/2016 were concerns that votes would not be counted or a belief that "there is no point in voting" (Kolbe & Muggah, 2016). It is likely that the negative effect was amplified by Haiti's unfortunate history with democracy. But whether a stronger history of democratic institutions makes a difference is left to future work.

This paper provides new insights on the medium-run effects of interventions on democracy, but there remains a question of the long-run effects. Perhaps a good comparison is the long-run effects of disenfranchisement on Black voting participation in the United States. Before many of the legal restrictions that disenfranchised them, Blacks in the South were deterred from voting by lynchings (Jones et al., 2017). Today, US counties with higher historical rates of lynchings have lower Black voter registration (Williams, 2021). The persistent effects of lynchings seems to stem from a lack of intergenerational behavior transmission: parents who were discouraged from voting cannot model democratic participation for their children. If the foreign interventions discourage political participation, and if intergenerational behavior transmission is important for sustaining democratic institutions, then such interventions might have much higher long-run costs than assumed.

## ACKNOWLEDGMENTS

I thank Timothy Guinnane, Kara Dimitruk, Wayne Sandholtz, and Briggs Depew for their comments. I also thank Jake Johnston for sharing data and insights on the elections.

## ENDNOTES

<sup>1</sup> Similarly, many states disenfranchise ex-felons, but even before they were convicted of crimes these voters had some of the lowest turnout rates (Burch, 2011).

<sup>2</sup> Miles and Feeney (2001) also report that CARICOM reported 20% turnout and Aristide's opposition claimed turnout was only 5%. But there were approximately 2.6 million votes in a country of 8.5 million, so the lowest turnout could be was 30%, and that assumes every Haitian (including children) was registered to vote.

<sup>3</sup> For more information on the controversy, see the summary by Jake Johnston at the Center for Economic and Policy Research, available at <https://web.archive.org/web/20190809152541/http://cepr.net/blogs/haiti-relief-and-reconstruction-watch/presidential-elections-in-haiti-the-most-votes-money-can-buy>.

<sup>4</sup> There were 614 sections with turnout data in the 2010 election. Of those, 609 had turnout data in both the 2015 and 2016 election. Of the remaining five, two had turnout data in 2015 and three had turnout data in 2016.

<sup>5</sup> Klimek et al. (2012) also recommend looking at the distribution of the logarithmic vote rate, which I do in Figure A2 of Supporting Information S2. Haiti seems to pass this test too.

<sup>6</sup> Specifically, the following districts are omitted: Port-au-Prince, Carrefour, Cite Soleil, Delmas, Gressier, Kenscoff, Petionville, and Tabarre.

<sup>7</sup> On August 14, 2018, Haitian filmmaker Gilbert Mirambeau Jr. posted a picture of himself blindfolded holding a piece of cardboard that read, "Kot kòb Petwo Karibe a???" which translates to "Where is the PetroCaribe money???" a reference to the missing \$2 billion

Venezuelan loans. Mirambeau's picture was like Mohamed Bouazizi's self-immolation in December 2010, sparking a political movement that engulfed the country. Soon other Haitians posted similar pictures, and the hashtag #KotKobPetroKaribeA became a rallying cry for those discontent with governmental disfunction.

## REFERENCES

- Acemoglu, D., Naidu, S., Restrepo, P. & Robinson, J.A. (2019) Democracy does cause growth. *Journal of Political Economy*, 127(1), 47–100.
- Ali, S.N. & Lin, C. (2013) Why people vote: ethical motives and social incentives. *American Economic Journal: Microeconomics*, 5(2), 73–98.
- Allcott, H. & Gentzkow, M. (2017) Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2), 211–236.
- Battaglini, M., Morton, R.B. & Palfrey, T.R. (2010) The swing voter's curse in the laboratory. *Review of Economic Studies*, 77, 66–89.
- Biggers, D.R. & Smith, D.A. (2020) Does threatening their franchise make registered voters more likely to participate? Evidence from an aborted voter purge. *British Journal of Political Science*, 50(3), 933–954.
- Burch, T. (2011) Turnout and party registration among criminal offenders in the 2008 general election. *Law and Society Review*, 45(3), 699–730.
- Cantoni, E. & Pons, V. (2019) Strict ID laws don't stop voters: evidence from a U.S. nationwide panel, 2008–2016, *NBER Working Paper No. 25522*.
- Cascio, E. & Washington, E. (2014) Valuing the vote: the redistribution of voting rights and state funds following the Voting Rights Act of 1965. *Quarterly Journal of Economics*, 129(1), 379–433.
- Coate, S., Conlin, M. & Moro, A. (2008) The performance of pivotal-voter models in small-scale elections: evidence from Texas liquor referenda. *Journal of Public Economics*, 92(3–4), 582–596.
- Collier, P. & Vicente, P.C. (2012) Violence, bribery, and fraud: the political economy of elections in Sub-Saharan Africa. *Public Choice*, 153(1–2), 117–147.
- Corstange, D. & Marinov, N. (2012) Taking sides in other people's elections: the polarizing effect of foreign intervention. *American Journal of Political Science*, 56(3), 655–670.
- Dellavigna, S., List, J.A., Malmendier, U. & Rao, G. (2017) Voting to tell others. *Review of Economic Studies*, 84, 143–181.
- Downs, A. (1957) *An economic theory of democracy*. Harper.
- Duffy, J. & Tavits, M. (2008) Beliefs and voting decisions: a test of the pivotal voter model. *American Journal of Political Science*, 52(3), 603–618.
- Dupuy, A. (2006) *The prophet and power : Jean-Bertrand Aristide, the international community, and Haiti*. Rowman & Littlefield Publishers.
- Gerber, A.S., Green, D.P. & Larimer, C.W. (2008) Social pressure and voter turnout: evidence from a large-scale field experiment. *American Political Science Review*, 102(1), 33–48.
- Hoekstra, M. & Koppa, V. (2021) (almost) no one votes without ID, even when they can, *Working Paper*.
- Johnston, J. & Weisbrot, M. (2011) *Haiti's fatally flawed election*. Center for Economic and Policy Research (CEPR).
- Jones, D.B., Troesken, W. & Walsh, R. (2017) Political participation in a violent society: the impact of lynching on voter turnout in the post-reconstruction south. *Journal of Development Economics*, 129, 29–46.
- Katz, J.M. (2013) *The big truck that went by*. Palgrave Macmillan.
- Klimek, P., Yegorov, Y., Hanel, R. & Thurner, S. (2012) Statistical detection of systematic election irregularities. *Proceedings of the National Academy of Sciences*, 109(41), 16469–16473.
- Kolbe, A. & Muggah, R. (2016) *Assessing Haiti's electoral legitimacy crisis—results of a 2016 survey, strategic note 21*. Igarapé Institute.
- Kudamatsu, M. (2012) Has democratization reduced infant mortality in sub-saharan Africa? Evidence from micro data. *Journal of the European Economic Association*, 10(6), 1294–1317.
- Levin, D.H. (2016) When the great power gets a vote: the effects of great power electoral interventions on election results. *International Studies Quarterly*, 60, 189–202.
- Miles, M. & Feeney, M. (2001) *Elections 2000: participatory democracy in Haiti*.
- Mobekk, E. (2001) Enforcement of democracy in Haiti. *Democratization*, 8(3), 173–188.
- Naidu, S. (2012) Suffrage, schooling, and sorting in the post-bellum US South, *Working Paper no. 18129 (June)*, NBER.
- Nohlen, D. (2005) *Elections in the Americas: a data handbook* (Vol. 1). North America, Central America, and the Caribbean: Oxford University Press.
- Organization of American States. (2010) *Statement by the OAS-CARICOM joint electoral observation mission on Haiti's presidential and legislative elections of 28 November 2010* [Press Release].
- Organization of American States. (2011a) Report of the Joint Electoral Observation Mission of the Organization of American States (OAS) and the Caribbean Community (CARICOM). *Electoral Observations in the Americas Series, no. 73*.
- Organization of American States. (2011b) *Report of the OAS-CARICOM Joint electoral observation mission on the first round of the presidential and legislative elections of 28 November 2010 in Haiti* [Press Release].
- Riker, W.H. & Ordeshook, P.C. (1968) A theory of the calculus of voting. *The American Political Science Review*, 62(1), 25–42.
- Rose, G. (2000) Democracy promotion and American foreign policy. *International Security*, 25(3), 186–203.
- Rosnick, D. (2011) *The Organization of American States in Haiti: Election monitoring or political intervention?*. Center for Economic and Policy Research. Available from: <https://cepr.net/documents/publications/haiti-oas-2011-10.pdf> [Accessed 3rd December 2021].
- Shulman, S. & Bloom, S. (2012) The legitimacy of foreign intervention in elections: the Ukrainian response. *Review of International Studies*, 38, 445–471.
- Ubert, L.J. & Jackson, D. (2020) Does aid promote electoral integrity? *Journal of Development Studies*, 56(6), 1067–1094.

U.S. Government Accountability Office. (2020) *Democracy assistance: state should improve information sharing with embassies*, GAO-20-173.  
von Hippel, K. (1995) Democratisation as Foreign Policy: the Case of Haiti. *The World Today*, 51(1), 11–14.  
Williams, J. (2021) Historical lynchings and the contemporary voting behavior of blacks. *American Economic Journal: Applied Economics*.

### SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

**How to cite this article:** Palsson, C. (2021) The medium-run effects of a foreign election intervention: Haiti's presidential elections, 2010–2015. *Contemporary Economic Policy*, 1–22. Available from: <https://doi.org/10.1111/coep.12559>