

On the Effects of Enforcement on Illegal Markets: Evidence from a Quasi-experiment in Colombia*

Daniel Mejía[†] Pascual Restrepo[‡] Sandra V. Rozo[§]

This version: June 9, 2015

Abstract

This paper studies the effects of enforcement on illegal behavior in the context of a large aerial spraying program designed to curb coca cultivation in Colombia. In 2006, the Colombian government pledged not to spray a 10km band around the frontier with Ecuador due to diplomatic frictions arising from the possibly negative collateral effects of this policy on the Ecuadorian side of the border. We exploit this variation to estimate the effect of spraying on coca cultivation by regression discontinuity around the 10 km threshold and by conditional differences in differences. Our results suggest that spraying one additional hectare reduces coca cultivation by 0.022 to 0.03 hectares; these effects are too small to make aerial spraying a cost-effective policy for reducing cocaine production in Colombia.

Keywords: Illegal Markets, War on drugs, Crime, Colombia.

JEL Classification: O13, O33, O54 and Q18.

*We thank the editor, Andrew Foster, and two anonymous referees for their valuable comments and suggestions. We are indebted to SIMCI, at the United Nations Office of Crime and Drugs (UNODC) in Colombia, for their invaluable collaboration in providing the data used in this study. We would also like to thank Adriana Lleras-Muney, Adriana Camacho, Leopoldo Ferguson, Paola Guiliano, and Leah Boustan for their suggestions. We are also grateful to the participants of the economics seminar at UCLA, Universidad de los Andes for the comments received.

[†]Universidad de los Andes, email: dmejia@uniandes.edu.co

[‡]MIT, email: pascual@mit.edu

[§]UCLA, email: sandra.rozo@marshall.usc.edu

1 Introduction

Illegal activities such as counterfeiting, tax evasion, and the operation of illegal drug markets remain a serious problem throughout the world. Yet, there is still an open debate about how to fight these problems, what strategies to use and the costs that different enforcement strategies entail. On one side, the economic analysis of crime sees the decision to engage in illegal activities as rational and, as such, shaped by incentives and penalties (see Becker, 1968 and Stigler, 1970). The central prediction is that enforcement reduces crime to the extent that it increases its costs. On the other side of the debate, social scientists and pundits have raised several concerns about the role of enforcement. In particular, critics argue that criminals may be irrational, myopic, or predisposed to illegal behavior (Menninger, 1968); that extrinsic penalties crowd out intrinsic motivations (Frey, 1997); or that enforcement may backfire if it conveys information about widespread illegal behavior (Benabou and Tirole, 2003, 2006). The debate also suffers from lack of abundant evidence, in part due to the lack of exogenous sources of variation in enforcement strategies required to uncover their causal impact on crime and illegal activities.

Our paper contributes to the growing literature aimed at disentangling and quantifying the causal effects of enforcement on crime and illicit activities. We use the war on drugs in Colombia as a case study, focusing on the role that aerial spraying with herbicides has on illegal coca cultivation. At least since 1996, Colombia has been the world’s largest cocaine producer and grower of coca crops (the raw input for producing cocaine). Coca cultivation takes place in remote areas of the country with little institutional presence, where farmers face the risk of being detected and their illegal crops being aurally sprayed with herbicides. When coca crops are sprayed with herbicides they are partially lost, which in principle should increase the cost of this illegal activity and reduce the farmers’ incentive to pursue it. Our goal in this paper is to assess the effectiveness of this form of enforcement in reducing coca cultivation.

We exploit the geographic and time variation on aerial spraying induced by a diplomatic friction between the governments of Ecuador and Colombia around 2006. In 2000, the Ecuadorian government alleged that Colombian aerial spraying campaigns near the frontier were causing health problems, productivity losses, and environmental damage in their territory. In response, the Colombian government committed to completely stop aerial spraying campaigns within a 10 km band around the international frontier with Ecuador at the beginning of 2006. The Colombian government broke its commitment at the end of 2006 and continued spraying within the band throughout 2007— though considerably less so than before. However, aerial spraying in the 10km band from the frontier completely stopped in 2008, in response to a lawsuit filed by the Ecuadorian government against its Colombian counterpart in international courts.

We use satellite and geo-referenced data of coca cultivation and aerial spraying on 1-square-km (100 hectares) grid cells between 2000 to 2010, and estimate the effect of aerial spraying using two complementary methodologies. First, we use a fuzzy regression discontinuity design and compare coca cultivation in cells near both sides of the 10 km threshold. We show that aerial spraying changed discontinuously at the 10 km threshold during the

years in which Colombia agreed not to spray the exclusion area, while other covariates and policies did not. We use this geographic discontinuity in enforcement to identify its effect on coca cultivation.

Additionally, we report results obtained using a conditional differences in differences estimation. In particular, we compare the cultivation of illicit crops in cells within the exclusion area to that in similar cells located 10 to 20 km away from the frontier, an area that continued to be sprayed throughout the years in our sample. Both groups of cells were exposed to aerial spraying before 2006, but after that, only the latter group of grid cells continued to be sprayed. Our estimator attributes the differential change in cultivation after 2006 across both areas to the change in enforcement. To guarantee the comparability of both groups, we control for coca cultivation and spraying before the intervention using a variety of parametric and non-parametric estimators.

Consistent with the view that illegal behavior is a rational choice, we find significant (but very small) deterrent effects of spraying on coca cultivation. The regression discontinuity estimates imply that cells in the sprayed area near the cutoff were 10% more likely to be sprayed than close cells in the exclusion area. Farmers responded by planting 0.3 hectares of cocaine less per square kilometer in the sprayed region. Similarly, our estimates using the conditional differences in differences estimator suggest that the areas that were exposed to aerial spraying after 2006 faced approximately 10% higher likelihood of being sprayed and, as a result, had on average 0.22 fewer hectares of coca per square kilometer (relative to cells in the region not sprayed). Both methodologies suggest that spraying one additional hectare reduces coca cultivation by 0.022 to 0.03 hectares in a given year.

Our findings confirm the key insight from the economics of crime for this particular context: Enforcement in the form of a higher likelihood of being sprayed with herbicides dissuades farmers from growing illegal crops. However, these effects are too small and aerial spraying of illicit crops is too costly. In particular, our largest point estimates suggest that to reduce coca cultivation by 1 hectare, approximately 33 additional hectares must be sprayed every year. However, it is very likely that coca cultivation is in part displaced by aerial spraying campaigns, making the 33 hectares a lower bound. The average direct cost to the U.S. of spraying one hectare of coca crops in Colombia is estimated to be about \$750 dollars (DNE, 2004, cited in Walsh et al., 2008). According to official sources, for each dollar the U.S. spends on the spraying program, the Colombian government spends about \$2.2 dollars protecting the spraying crews and cleaning up the area before they carry out these campaigns. Thus, the joint cost of spraying 33 hectares of coca, and reducing coca cultivation by 1 hectare per year, is about \$79,200 dollars, out of which the U.S. pays at least \$24,750. As we show in greater detail in the paper, these numbers imply that the marginal cost to the U.S. of reducing cocaine supply in retail markets by 1 kg by subsidizing aerial spraying policies in Colombia is about \$1.6 million dollars, which is in the ballpark of the costs reported by Mejia and Restrepo (2013) using a different methodology. This is significantly higher than the marginal cost of reducing cocaine consumption in the U.S. using other policies, such as interdiction efforts in Colombia (\$175,000 dollars; see Mejia and Restrepo, 2013), or treatment and prevention policies in the U.S. (\$8,250 and \$68,750 dollars, respectively; see MacCoun and Reuter, 2001). It is also high when compared to the

retail price of 1 kg of pure cocaine in U.S. retail markets, which ranges from \$100,000 to \$150,000 dollars.

In addition to providing evidence on the link between enforcement and illegal behavior, estimating the impact of aerial spraying on coca cultivation is important for several reasons. First, Colombia has been a key player in the international drug trade during the last 30 years. During our period of analysis, it was the world's main cocaine producer, supplying nearly 70% of cocaine and holding a similar share of coca crops among the Andean countries (see UNODC, 2012). Thus, reducing the supply of Colombian cocaine has the potential to decrease the availability of cocaine and its associated harms throughout the world. Second, aerial spraying is the largest anti-drug program implemented in Colombia. It entails not only resources from the local government but also from the U.S. In particular, since the beginning of *Plan Colombia* in 2000 – the largest cooperative effort between the U.S. and a source country to curb drug supply and improve security conditions – both countries have spent more than \$3 billion dollars in this program. Finally, illegal behavior in Colombia is pervasive, and understanding its nature and how to contain it is a key policy challenge.

The rest of the paper is organized as follows. Section 2 describes the related literature; section 3 describes the Colombian context and the natural experiment that we exploit in order to identify the causal impact of aerial spraying on coca cultivation. Section 4 presents the data and estimates the effects of spraying. Section 5 discusses the main results and presents a cost-benefit analysis of the aerial spraying program in Colombia. Finally, section 6 concludes.

2 Related literature

Our paper is related to two branches of the economics literature. First, it is related to the literature on the effects of enforcement on crime. This topic goes back to the seminal contributions of Becker (1969), Stigler (1970), and Ehrlich (1973). The main implication of these papers is that enforcement— in the form of fines, tighter punishments, or a higher probability of detection— reduces crime and illegal behavior. Yet, testing this proposition is challenging, as it requires credible sources of exogenous variation in enforcement. Otherwise, the fact that enforcement reacts to crime creates a misleading upward bias in the estimated effect of enforcement on crime. Initially, the economics literature failed to find empirical support for this proposition (see Cameron, 1988, Marvell and Moody, 1996, and Eck and Maguire, 2000, for surveys of the early literature), but many of these contributions were plagued with endogeneity issues.

Recent studies have addressed identification more carefully and find some support for the deterrence effect of enforcement on some types of crimes and behavior. For instance, Marvell and Moody (1996) find that within-state increases in the number of police officers reduce crime in the U.S. Levitt (1997) uses electoral cycles as an instrument for police hiring and finds significant reductions in crime when more policemen are hired.¹ Corman and Mocan (2000) use high frequency changes in the number of police officers and find that police

¹See also the criticism by McCrary, (2002), and the reply by Levitt (2002).

reduces burglaries, but has no effect on other crime categories. Di Tella and Schargrodsky (2004) exploit the exogenous reallocation of police forces across Buenos Aires as a result of a terrorist attack on a Jewish Center. They find a large and localized deterrent effect of more police presence on car thefts. A similar strategy is used by Draca, Machin, and Witt (2011), who also find evidence of deterrence effects by exploiting police reallocation in London after the terrorist attacks of 2005. Evans and Owens (2007) use state grants to fund Community Oriented Policing Services (COPS) as an instrument for the number of police officers. They find that higher police presence reduces auto thefts, burglaries, robberies, and aggravated assaults. Buonanno and Mastrobuoni (2012) exploit delays created by a centralized police hiring system in Italy to estimate the effect of police officers on local crime, finding deterrence effects in some crime categories. Finally, Garcia, Mejia, and Ortega (2012) study the randomized introduction of a police training program among small localities in the four largest cities in Colombia. They find that the intervention significantly reduces crime, not by increasing the police force, but by improving its quality and engagement with the community.

Another body of literature focuses on the effects of enforcement, or characteristics related to the likelihood of detection, on soft crime or tax evasion.² For example, Bar-Ilan and Sacerdote (2001) find that the introduction of traffic cameras and changes in fines reduced driving infractions. Dubin, Graetz, and Wilde (1987) and Beron, Tauchen, and Witte (1992) present evidence that higher audit rates modestly increase reported income for some groups of taxpayers. In this same area, Klepper and Nagin (1989) find that noncompliance rates are related to the traceability, deniability, and ambiguity of the items being declared, which are in turn related to the probability that evasion will be detected and punished; and Kagan (1989) presents evidence that compliance is greater among people whose income is directly reported to the IRS and who therefore have fewer opportunities to cheat.

We contribute to this literature by cleanly identifying the effect of enforcement on illegal behavior in the context of the war on drugs and illicit crop cultivation in Colombia. The strength of our empirical exercise relies not only on our identification strategy, but also on the precision of our data on illicit crop cultivation and enforcement activities. In particular, we observe satellite data on coca cultivation in small 1-squared-km cells, and information on the exact location of aerial spraying campaigns, closely monitored by the army and police using GPS devices that are built-in in the aircraft used in the aerial spraying program in Colombia. Consistent with the previous findings on the literature, our results suggest that farmers respond to a greater likelihood of enforcement in an area by reducing illicit coca cultivation there, but we show that the effects are too small to make the spraying program a cost-effective policy to reduce cocaine supply.

Our paper is also related to the applied economics literature on the cost-effectiveness of anti-drug policies. The main challenge in this area is that anti-drug interventions typically take place on a large scale; hence, it is difficult to obtain appropriate counterfactuals. One approach, followed by Mejía and Restrepo (2011, 2013), is to construct and calibrate economic models of illegal drug markets to understand and quantify the main forces and determinants of the cost-effectiveness of different supply-reduction strategies. They find that

²For a thorough review of the empirical literature, see Andreoni, Erard, and Feinstein (1998).

spraying illicit crops is costly and ineffective relative to policies aimed at seizing drug shipments. However, both strategies are costly relative to demand reduction policies in consumer countries.

Other papers in the literature have focused on estimating the impact of spraying campaigns on coca cultivation by using geographic and time variation. For example, Moreno-Sanchez et al. (2003) and Dion and Russler (2008) use departmental data from Colombia and find a positive correlation between the levels of spraying and the presence of coca crops. However, these results are likely to be driven by simultaneity bias in their estimates. Recent studies have attempted to address these endogeneity concerns. For example, Moya (2005) deals with selection on observables using matching techniques, and finds that spraying does not have a significant effect on coca crops in Colombian municipalities. Bogliacino and Naranjo (2012), also exploit within municipality variation and find that eradication does not reduce coca production. They estimate a system of equations in which aerial spraying and cultivation are determined simultaneously, identified by the restrictive assumption that crime rates and the number of internal refugees are uncorrelated with coca cultivation, but do create political pressure for spraying campaigns. Reyes (2014) instruments spraying with the distance to the closest military base, and finds evidence that aerial spraying increases illicit crops. However, his results require the location of military bases to be exogenous. A different approach is implemented by Ibanez and Carlsson (2010), who elicit farmers' preferences over risk and the profitability of other crops by using hypothetical questions. Consistent with our findings, and the economic model of crime, farmers in Putumayo claim they would reduce coca plantations if they faced a higher risk of eradication. However, the implied reductions are small, as we confirm using a different methodology. Finally, Roza (2014) instruments spraying with the interaction of the distance between each 1-square-km cell (or coca producer) and the nearest border of a protected area and U.S. international anti-drug expenditures. She exploits the fact that, by governmental mandate, protected areas cannot be sprayed with herbicides due to environmental and social concerns, and finds that aerial spraying has a negative and significant effect on the hectares of coca cropped and their yield. In addition, she documents negative unintended consequences on the socio-economic conditions of coca-producing areas.

This paper contributes to the existing evidence by estimating the effects of aerial spraying programs using a sharp natural experiment, which in our view provides a credible source of exogenous variation. In addition, we use cost figures to back up a lower bound for the cost-effectiveness of these programs. We find that despite reducing cultivation, aerial spraying is too costly to be a cost-effective anti-narcotic strategy. In particular, demand reduction policies in the U.S. or interdiction campaigns provide the same benefits in terms of supply reduction at much lower costs.

3 The Colombian context and the natural experiment

Following the large increase in coca cultivation that took place in Colombia after 1994 and the increasing involvement of illegal armed groups in these activities, the governments of

Colombia and the U.S. launched *Plan Colombia* in September of 1999, a large anti-narcotics program aimed at reducing the Colombian cocaine supply. Under this program, the United States government disbursed close to \$540 million dollars per year between 2000 and 2008 in subsidies to the Colombian armed forces to fight against the production and trafficking of drugs. Additionally, the Colombian government spent close to \$810 million dollars per year during the same period in the fight against illegal drug production and trafficking (GAO, 2008). Total expenditures on the military component of *Plan Colombia* represented close to \$1.35 billion dollars per year, corresponding to about 1.2% of the country’s annual GDP, making it the largest anti-drug intervention in a producing country.

The strategies implemented under *Plan Colombia* included aerial spraying campaigns, manual eradication, control of chemical precursors used in the processing of coca leaf into cocaine, detection and destruction of cocaine processing laboratories, and seizure of drug shipments en route to foreign countries. Aerial spraying has been by far the main anti-drug strategy in terms of financial resources invested. On average, 128,000 hectares per year were sprayed with herbicides, of which almost half are located in Putumayo and Nariño, the two Colombian departments (states) bordering Ecuador, where our empirical analysis is centered. Figure 1 shows the evolution of coca cultivation, aerial spraying with herbicides, and manual eradication between 2000 and 2010, for the whole country (Panel A) and for the departments of Nariño and Putumayo (Panel B).

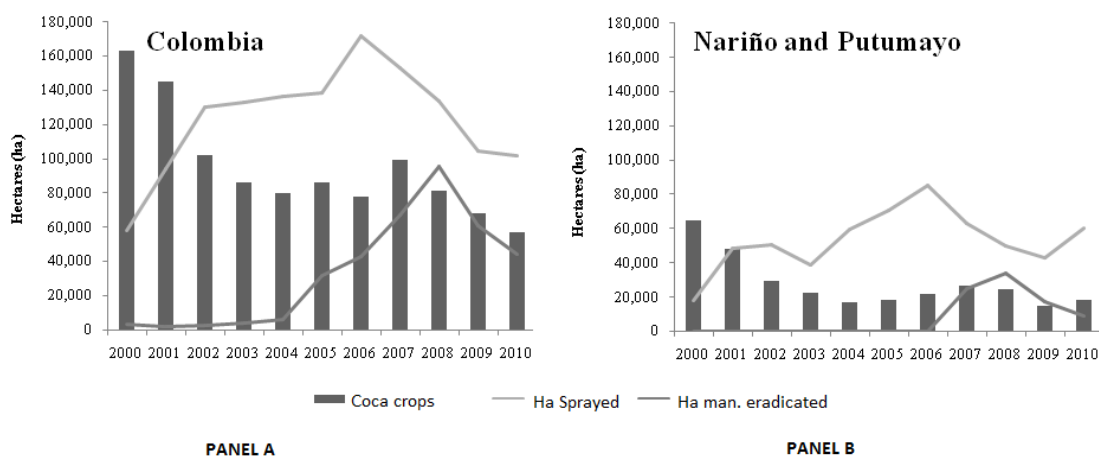


Figure 1: Coca cultivation and aerial spraying in Colombia (left panel) and Narino and Putumayo (right panel). Data from the United Nations Office of Drugs and Crime, UNODC.

Spraying campaigns are carried out by American contractors, such as *DynCorp*, using small aircraft. Coca crops are sprayed with chemicals containing glyphosate, such as *Roundup*, the commercial name of the herbicide used in the spraying program in Colombia. The herbicide is absorbed through the plant foliage and is effective only on growing plants (e.g., it is not effective in preventing seeds from germinating). It kills the plant by inhibiting its growth. Though *Roundup* was designed to kill weeds and grasses, including coca bushes, it may also affect other legal crops that are not glyphosate-resistant. Aerial spraying with glyphosate is targeted at areas where coca crops have been detected using satellite images.

When planting coca crops, farmers face the risk of having their crops destroyed by the

herbicides used in aerial spraying. Given this risk, they may still grow coca bushes and play their luck, or mitigate the effects of the herbicide using a variety of techniques. For instance, farmers can spray molasses on the coca bushes to prevent the herbicide from penetrating the foliage and killing the plant. In addition, they can cut the stem of the plant a few hours after the fumigation event, enabling the plant to grow back a few months later. Finally, farmers can relocate their crops to areas less likely to be sprayed. However, these alternatives are costly and force some farmers to cultivate legal crops instead, which are, in principle, not targeted by spraying campaigns.

Because aerial spraying campaigns typically target areas with a high prevalence of coca plantations, traditional estimates of the effect of spraying on cultivation are biased upwards. In this paper, we solve this problem and identify the effects of aerial spraying using a natural experiment. In particular, we exploit a diplomatic friction between the governments of Colombia and Ecuador, which concluded with the compromise by the Colombian government not to carry out spraying campaigns within a 10 km strip along the international border with Ecuador, starting in 2006.

From the beginning of *Plan Colombia*, the Ecuadorian government complained of alleged adverse effects of spraying on the health of its population, the environment, livestock, and legal crops near the bordering area. In 2006, the Colombian government announced that it would discontinue aerial spraying within a 10 km band along the international frontier with Ecuador within Colombian territory. However, the Colombian government recanted at the end of 2006 and continued spraying the area. As a result of this noncompliance with the initial agreement, the Ecuadorian government filed a lawsuit against Colombia in the International Court of Justice in The Hague. Since the suit was filed, on March 31st 2008, the Colombian government stopped all spraying campaigns within the 10 km strip.

The implementation of this exclusion area generated geographical and time variation in the likelihood of aerial spraying, which we exploit to identify its effects. Figure 2 shows a map of the exclusion strip and its location in Colombia.

4 Data

We employ a unique panel data on the location of coca crops within 1-square-kilometer (or 100 hectares) cells from 2000 to 2010. The data is collected and processed by the United Nations Office for Drugs and Crime (UNODC) in Colombia, and comes from satellite images. Using these images, UNODC estimates the number of hectares with coca cultivation detected on each grid by the end of each year. We also use cell level data on the location of aerial spraying campaigns for the same period. The data is collected by the Colombian police using GPS devices, and it records the exact location of the plane when the spraying valves are open. Using these observations we code a dummy of whether a grid was sprayed or not, for each year from 2000 to 2010.³

³Moreover, we use data on whether manual eradication campaigns took place on each grid, covering the 2007-2010 sub-period. These data are obtained from GPS devices used by manual eradication teams.

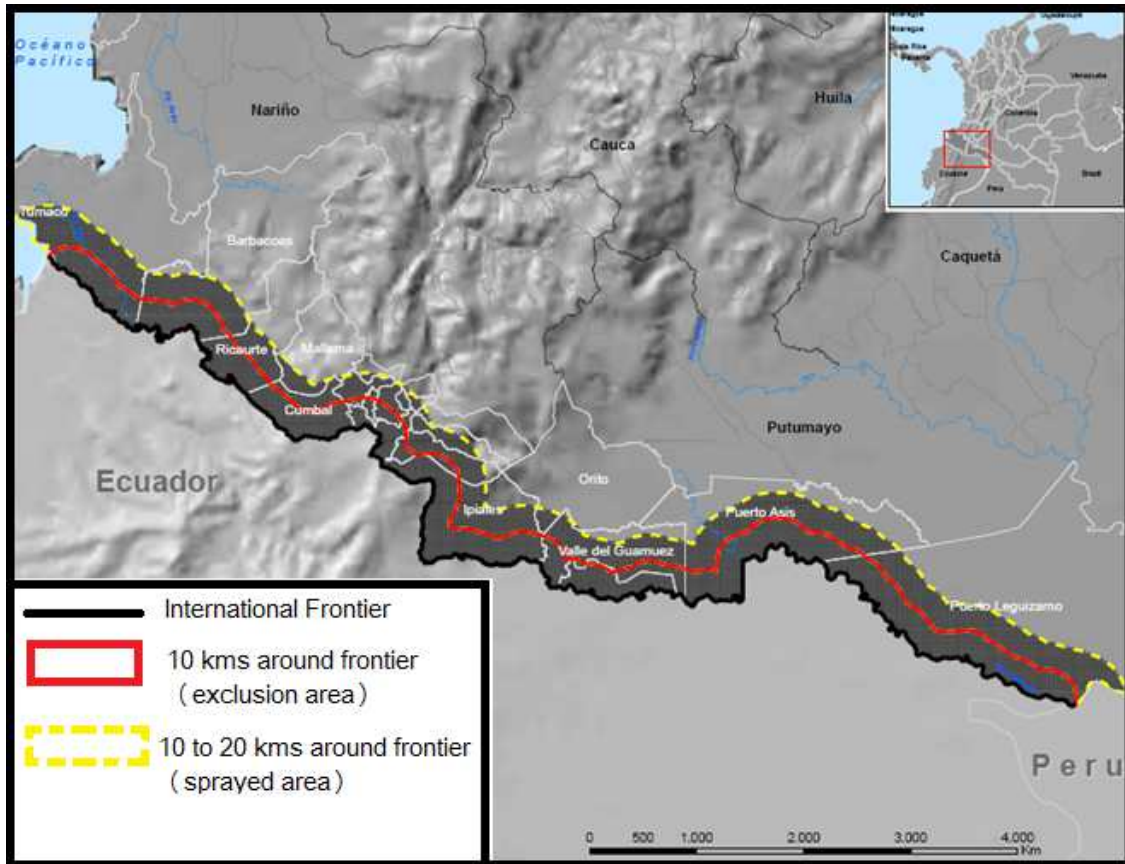


Figure 2: Map of the frontier between Colombia and Ecuador illustrating the sprayed and exclusion areas.

We restrict our sample to all grid points with centroids located within 20 km of the international frontier with Ecuador. Our sample includes 10,880 cells. We refer to the 5,613 cells within 10 km of the frontier as the exclusion region, since this is the area that Colombia agreed not to spray. In contrast, we refer to the 5,275 cells located 10 km to 20 km from the frontier as the sprayed area, as these cells were sprayed throughout our period of analysis. Both regions are depicted in Figure 2.

To summarize the data, Figure 3 presents the likelihood of aerial spraying (left panel) and the average number of hectares with coca crops per square kilometer (right panel) in both regions from 2000 to 2010. As anticipated above, the figure reveals similar patterns before 2006, with the exception of 2004. However, and consistent with the description of the natural experiment we exploit, a significant gap opens up beginning in 2006, when the Colombian government first agreed to reduce the spraying campaigns in the exclusion area. The data shows that the Colombian government still sprayed the exclusion area in 2006 and 2007, but considerably less than the sprayed region, where about 20% of the cells were sprayed. However, spraying of the exclusion area falls to virtually zero, effectively since 2008—the year in which Ecuador filed the lawsuit against Colombia in international courts.

The data on cultivation reveals a sharp decline from 2000 to 2004, during the first years

of *Plan Colombia*, from about 3 hectares per square kilometer to about 0.6. However, in 2006, and later in 2009 and 2010, cultivation increased in the exclusion region relative to the sprayed area; though only mildly.

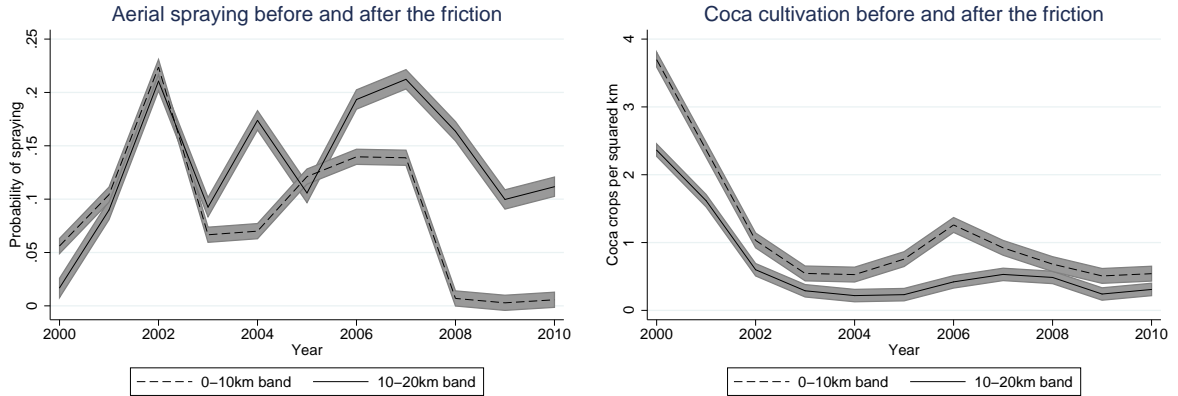


Figure 3: Coca cultivation and likelihood of aerial spraying in the exclusion (dotted line) and sprayed areas (solid line) from 2000 to 2010. 95% confidence intervals for the mean in each group are presented in gray.

5 Fuzzy regression discontinuity approach

In this section we employ a fuzzy regression discontinuity design to evaluate the impact of aerial spraying on coca cultivation. We exploit the exogenous rule applied by the Colombian government since 2006 to stop aerial spraying 10 kms around the international frontier, and which became particularly binding from 2008 onward.

In our setting, the forcing variable is the distance from the centroid of each cell i to the international frontier with Ecuador, capturing its geographic location. We normalize the forcing variable to take the value of zero at the 10 km cutoff, and denote it by \widehat{D}_i , where $\widehat{D}_i = D_i - 10km$.⁴ Our discussion above implies that in the remaining sample of cells, there should be a discontinuity in aerial spraying around $\widehat{D}_i = 0$ in 2006, and from 2008 onwards, assuming that the Colombian government fulfilled its commitment strictly during these years.

Let S_{it} be a dummy equal to 1 if grid i was sprayed during year t . We start by exploring whether the diplomatic friction led to a discontinuity in aerial spraying near the 10km band.

⁴In this exercise, we exclude from our sample all cells that had their centroid in the first 500 m around the cutoff value, since they have a significant portion of their territory in both the exclusion and the sprayed area. Thus, we only compare cells near the 10 km cutoff lying entirely on one side or the other. We obtain similar results using the cells within 500 m of the cutoff to estimate the conditional expectation of cultivation and the likelihood of spraying as a function of the distance to the cutoff. In these models, we add separate dummies for cells within 0 to 500 m away from the cutoff in the exclusion area, and cells within 0 to 500 m away from the cutoff in the sprayed area.

To do so, we estimate the model:

$$S_{it} = \pi_{0t} + \pi_{1t}1\{\widehat{D}_i > 0\} + f_t(\widehat{D}_i, X_i) + \varepsilon_{it}, \quad (1)$$

for specific years, t , or pooling different years together. Here, f_t is a polynomial in the forcing variable and other geographic characteristics of the cell (including longitude and latitude), which captures the continuous variation of the likelihood of enforcement over different cells. Essentially, it represents the conditional expectation of the policy based solely on the geographic characteristics of a cell. The coefficient on π_{1t} measures any policy discontinuity around the cutoff. Our discussion above implies that we expect $\pi_{1t} > 0$ for $t \geq 2006$; in particular since 2008, when spraying was reduced virtually to zero in the exclusion area. In addition we control for municipality fixed effects, year effects (when necessary), and report standard errors robust against heteroskedasticity and serial correlation within cells throughout.

The main challenge in estimating equation 1 is to specify a flexible model for the conditional expectation function $f_t(\widehat{D}_i, X_i)$. Our first approach is to use a cubic polynomial in D_i , together with quadratic terms for latitude and longitude, aimed at capturing the variation of S_{it} over space. We present estimates using this approach in columns 1 and 2 of Table 1. In column 1 we estimate equation 1 on the sample of cells with centroids within 3km of the discontinuity, and in column 2 we further restrict the sample to cells within 2.5km. By restricting the sample of cells used, we rely less on the particular parametrization of the conditional expectation function. Alternatively, in columns 3 and 4 we approximate the conditional expectation using a local quadratic regression with different choices of bandwidth and a triangular Kernel. In column 3 we use Imbens and Kalyanaraman (2012) optimal bandwidth (labeled *IK* throughout), which in our setting equals $b = 4.7km$.⁵ Since one may be concerned that this is too large, we also set the bandwidth to $b = 3km$ in column 4.

The estimates in column 1 show that, since 2006, cells in the sprayed area near the 10 km cutoff were 11.2 percentage points more likely to be sprayed than similar cells in the exclusion region (standard error=0.023). We find similar estimates when we pool the years 2008 to 2010 together — when the spraying campaigns were reduced to nearly zero in the exclusion area— and in the remaining models in columns 2 to 4. Importantly, we find a similar pattern year by year in Panel B. Though the estimates are less precise, they do point out to a discontinuity in the likelihood of aerial spraying since 2006 around the 10km band near Ecuador. Consistent with our discussion, the discontinuity becomes more clear since 2008, when Colombia stopped aerial spraying campaigns in the exclusion band altogether.

The previous results can be seen graphically. Figure 4 shows the local behavior of the likelihood of aerial spraying on both sides of the 10 km cutoff for cells within 2.5 km of the discontinuity. To ease the graphical analysis, we plot cells by the distance of their border to the cutoff.⁶ The plots reveal a clear discontinuity in the likelihood of spraying after 2008,

⁵Though the optimal bandwidth varies estimate by estimate, all are very close to the 4.7km one used here. Thus, we use this bandwidth and label it as the optimal bandwidth throughout.

⁶This is defined as $\widehat{D}_i - 500$ on the right of the cutoff and $\widehat{D}_i + 500$ on the left. By doing so, we remove from the figure the 500 meter band around the cutoff that we excluded from the estimation sample. We use a quadratic polynomial to approximate the local behavior on each side of the 10 km cutoff.

Table 1: Estimates of the discontinuity in spraying and coca cultivation around the 10km cutoff (sprayed minus exclusion area).

Model for conditional expectation:	Discontinuity in spraying				Discontinuity in cultivation			
	Cubic polynomial for grids between:		Local model with bandwidth:		Cubic polynomial for grids between:		Local model with bandwidth:	
	$\pm 3km$ (1)	$\pm 2.5km$ (2)	$b = 1K$ (3)	$b = 3km$ (4)	$\pm 3km$ (5)	$\pm 2.5km$ (6)	$b = 1K$ (7)	$b = 3km$ (8)
<i>Panel A: Pooled-years estimates.</i>								
Difference after 2006:	0.112*** (0.023) [0.023]	0.086*** (0.029) [0.026]	0.107*** (0.019) [0.021]	0.098*** (0.033) [0.027]	-0.479*** (0.138) [0.147]	-0.644*** (0.172) [0.153]	-0.473*** (0.121) [0.132]	-0.501** (0.204) [0.183]
Observations	13265	10560	22340	13265	13265	10560	22340	13265
Difference after 2008:	0.117*** (0.025) [0.026]	0.098*** (0.032) [0.029]	0.107*** (0.021) [0.023]	0.089** (0.036) [0.029]	-0.371*** (0.117) [0.131]	-0.454*** (0.140) [0.137]	-0.337*** (0.104) [0.111]	-0.424** (0.168) [0.157]
Observations	7959	6336	13404	7959	7959	6336	13404	7959
<i>Panel B: Year-by-year estimates.</i>								
Difference in 2006:	0.079 (0.051)	0.079 (0.066)	0.085* (0.045)	0.109 (0.076)	-0.840* (0.463)	-1.403** (0.584)	-0.794** (0.403)	-0.966 (0.691)
Difference in 2007:	0.128** (0.059)	0.057 (0.074)	0.130** (0.051)	0.116 (0.085)	-0.444 (0.305)	-0.455 (0.374)	-0.558** (0.268)	-0.268 (0.428)
Difference in 2008:	0.119*** (0.046)	0.086 (0.058)	0.108*** (0.039)	0.068 (0.066)	-0.440** (0.213)	-0.641*** (0.247)	-0.397** (0.186)	-0.644** (0.287)
Difference in 2009:	0.096*** (0.036)	0.077* (0.044)	0.068** (0.030)	0.080 (0.050)	-0.210 (0.161)	-0.263 (0.190)	-0.195 (0.142)	-0.140 (0.227)
Difference in 2010:	0.137*** (0.045)	0.132** (0.058)	0.144*** (0.039)	0.120* (0.065)	-0.463** (0.216)	-0.460* (0.266)	-0.421** (0.194)	-0.488 (0.325)
Observations per year	2653	2112	4468	2653	2653	2112	4468	2653

Notes: The table presents regression discontinuity estimates of the difference in spraying (columns 1 to 4) and cultivation (columns 5 to 8) around the 10 km cutoff. Panel A presents estimates pooling several years together, as indicated by the row labels. Panel B presents year-by-year estimates. Columns 1,2, 5, and 6 use a global approximation to the conditional expectation based on a cubic polynomial, and restrict the sample to grids within 3km and 2km of the discontinuity. Columns 3,4, 7, and 8 use a quadratic local regression to approximate the conditional expectation, and sets a bandwidth of 4.7km (following Imbens and Kalyanaraman, 2012) or 3km. Standard errors robust against heteroskedasticity and serial correlation within cells are reported in parenthesis. Estimates with *** are significant at the 1%, those with ** are significant at the 5%, and those with * are significant at the 10%.

and to a lesser extent for 2006 and 2007. When we pool the years 2006 to 2010 (after the diplomatic friction began), we observe a clear discontinuity in the likelihood of spraying.

We now investigate the consequences of the discontinuity in spraying on coca cultivation. Let Y_{it} be the hectares with coca crops in cell i in year t , measured with satellite images at the end of each year. We estimate the following specification:

$$Y_{it} = \gamma_{0t} + \gamma_{1t}1\{\widehat{D}_i > 0\} + f_t^Y(\widehat{D}_i, X_i) + \epsilon_{it}, \quad (2)$$

for different years, t . Here, f_t^Y is another polynomial approximating the conditional expectation of cultivation, based solely on geographic characteristics (distance to the frontier, latitude, longitude). The coefficient on γ_{1t} measures any discontinuity around the cutoff.

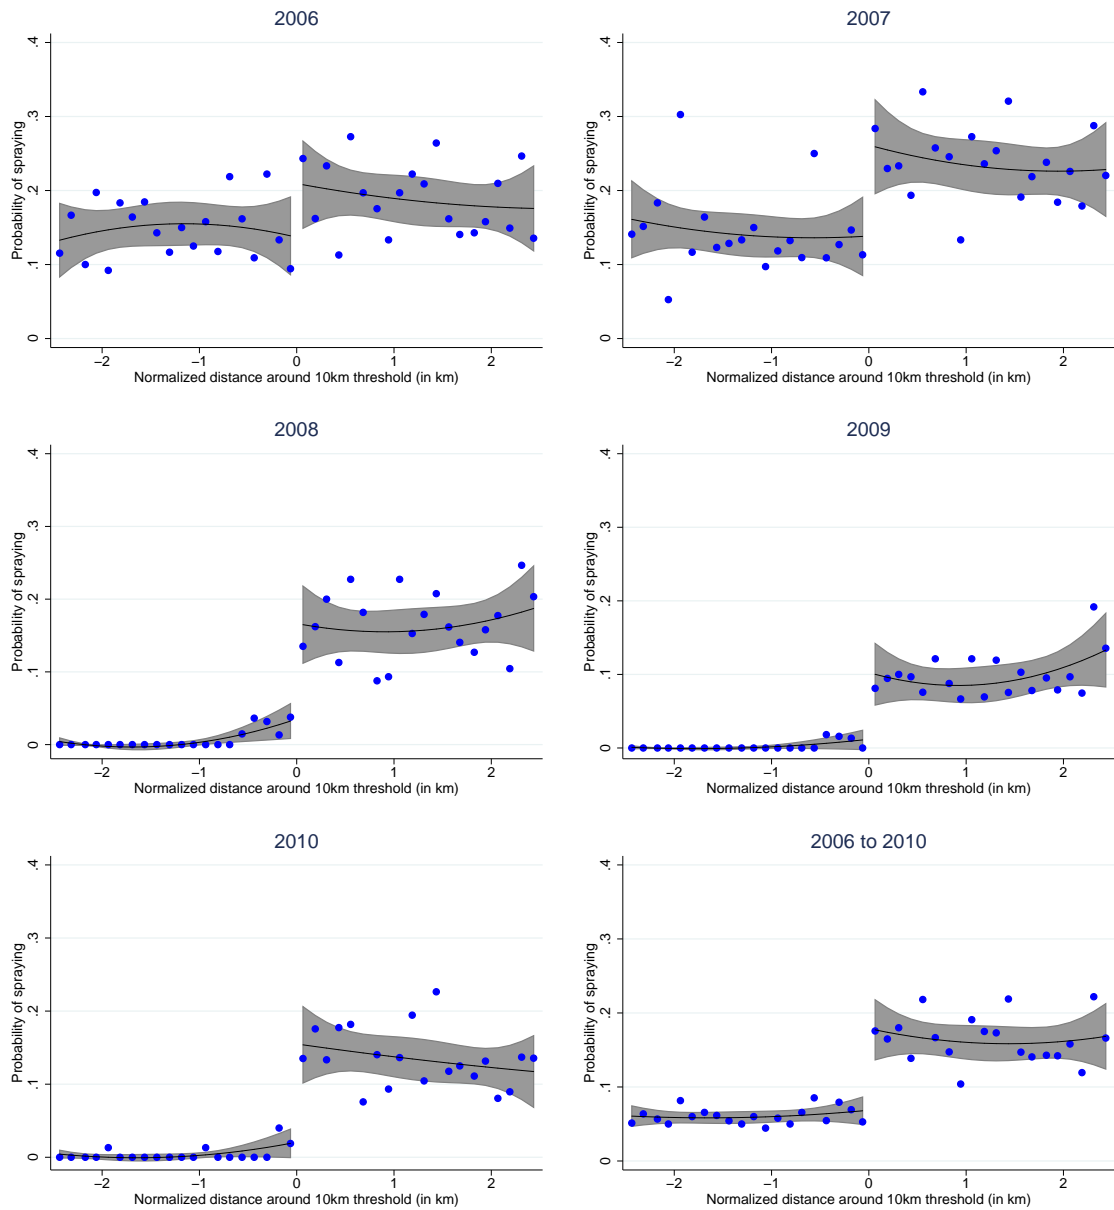


Figure 4: Probability of aerial spraying around the 10 km cutoff during years in which Colombia committed not to spray the excluded region (2006 to 2010).

Columns 5 to 8 of Table 1 present estimates of the difference in coca cultivation at the discontinuity, γ_{it} , for several years separately and for pooled years. Each column presents estimates based on different approximations of the conditional expectation function analogous to the ones used in columns 1 to 4. Consistent with the results on spraying, we find evidence of reductions in cultivation in the sprayed area since 2006, relative to the exclusion area. In this case, the estimates in column 5 suggest that cultivation was reduced by about 0.48 hectares per square kilometer in cells near the cutoff in the sprayed region relative to the exclusion area since 2006 (standard error=0.14). The estimates in columns 5 through 8 confirm our findings. When we focus on the years from 2008 onward, we also find significant,

though slightly smaller, differences in cultivation. The results are not very precise when the sample is divided by years, but they are mostly significant at the 10% confidence level, or near significant, and all point to a reduction in cultivation in the sprayed relative to the exclusion area.

Figure 5 also present these results graphically (the construction of these figures is analogous to that of Figure 4). The above results suggest that the enforcement of the 10 km exclusion area created a discontinuity in enforcement around the cutoff after and during 2008, and less so for 2006 and 2007. The discontinuity in spraying caused divergent illegal behavior on both sides of the cutoff.

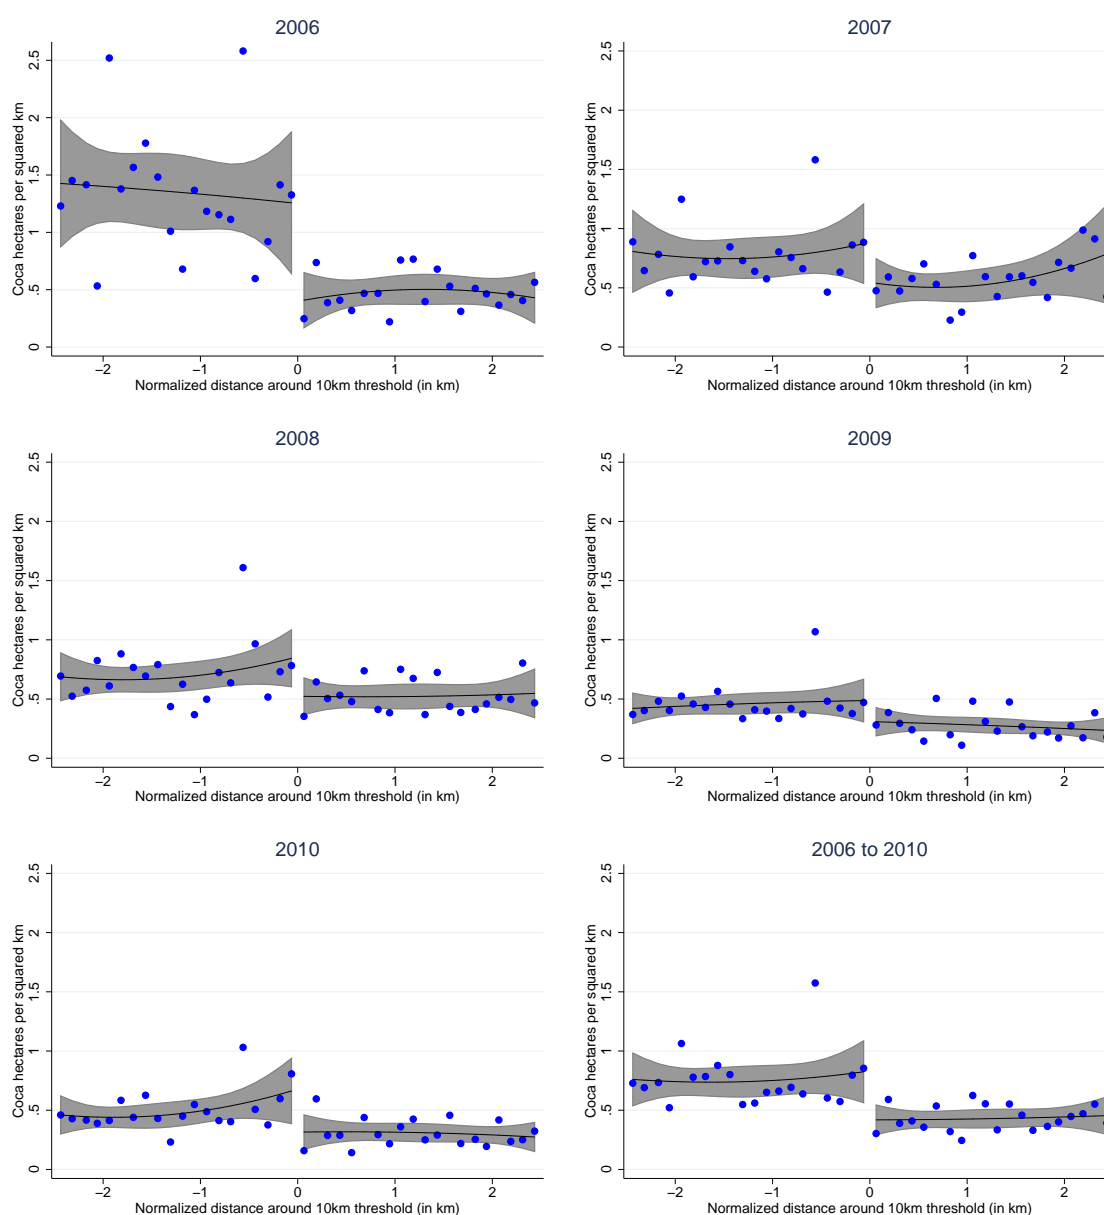


Figure 5: Coca cultivation around the 10 km cutoff during years in which Colombia committed not to spray the excluded region.

Throughout, we present standard errors clustering at the cell level and robust against heteroskedasticity. In this case, the validity of our standard errors requires the error terms ε_{it} and ϵ_{it} to be uncorrelated across cells. We believe this is a good starting point, since the flexible polynomials f_t and f_t^Y — which include detailed geographical information— already purge the errors from sources of spatial correlation depending on geographical proximity. Moreover, municipality fixed effects also remove sources of spatial correlation related to unobserved differences across municipalities. In addition, we also present standard errors allowing for spatial correlation in square brackets below our main estimates in the top panel. These standard errors, based in Conley (1999), allow ε_{it} and ϵ_{it} to be correlated among cells within 5km of each other, and also permit these error terms to be serially correlated. Overall, they are close to our traditional standard errors, lending support for our inference.⁷

The validity of our estimates requires our approximation of the conditional expectations to be valid, so that we are capturing a true discontinuity. To support our approach, we take advantage of the timing of the diplomatic friction and show that there is no discontinuity in spraying nor in cultivation before 2006. Figure 6 shows that, when we pool all years before 2006 together, there is no evidence of a discontinuity. This suggests that our estimates in Table 1 are not driven by unobserved characteristics of cells or a misspecification of the conditional expectation— as these would show up as a discontinuity in the years before 2006. Moreover, these findings suggest that the particular choice of the exclusion area was rather arbitrary and was not done strategically aiming at certain cells with particular cultivation dynamics, at least not near the 10km cutoff.

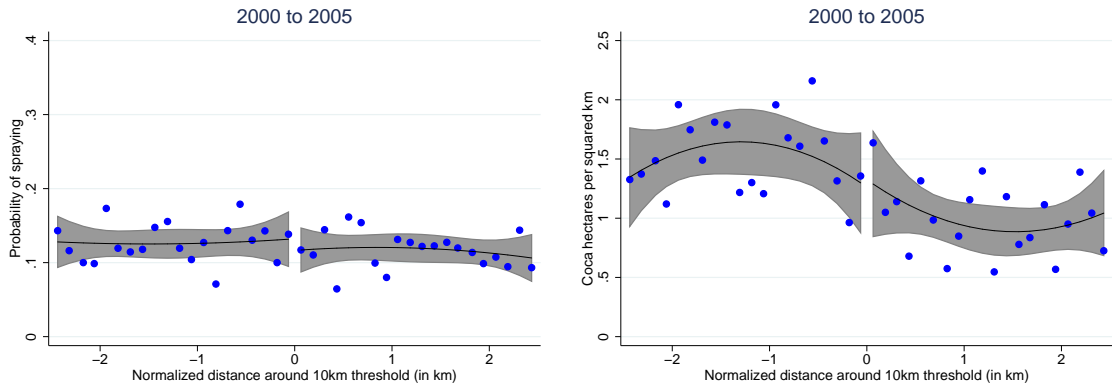


Figure 6: Probability of spraying and coca cultivation around the 10 km cutoff before the diplomatic friction leading to a suspension of aerial spraying campaigns near Ecuador.

We further explore this point by plotting year-by-year estimates of the difference around the 10km threshold in spraying and cultivation, together with their standard errors. We focus on the specification obtained with the local quadratic polynomial and $b = 3km$ —

⁷ We use the implementation and codes by Hsiang (2010). These are faster than Conley’s and also allow for some restricted forms of serial correlation. We also computed, but do not report, Conley (1999) standard errors that allow for spatial correlation between a cell and the 8 cells located in a 3×3 square of grids around it; or the 24 cells located in a 5×5 square of grids around it. These standard errors were almost identical to the reported standard errors as well.

which is the more demanding one. Consistent with our motivation, we find no evidence of discontinuities in any year before 2006; with most point estimates being small (relative to the yearly estimates for 2006 to 2010) and close to zero.⁸

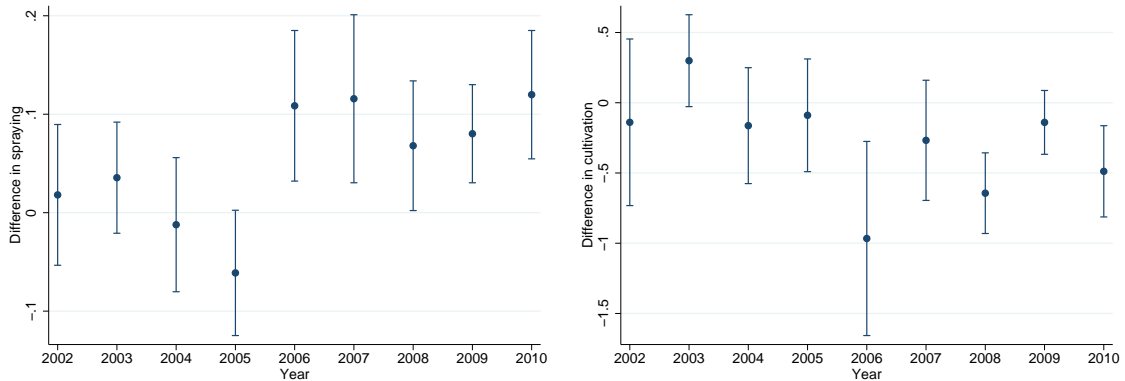


Figure 7: Year-by-year estimates of the difference in cultivation and aerial spraying near the 10km cutoff, both before and after the diplomatic friction.

An alternative test of our specification consists of conducting Placebo tests around random cutoffs which experienced no policy changes. To do so, we draw a random sample of 10,000 cutoffs from 10 to 20km away from the frontier. We then estimated a discontinuity in cultivation and spraying at each cutoff pooling data from 2006 to 2010.⁹ Figure 8 plots the empirical distribution of these estimates. As expected, they have mean zero, but are also remarkably close to zero for all artificially created cutoffs. For both spraying and cultivation, our estimates in Table 1 pooling the years 2006 to 2010 are clearly on the tails of both empirical distributions, indicating that they are unlikely to be spurious.

Finally, we conduct two additional tests showing that other observed covariates do not exhibit discontinuities near the 10km cutoff. In particular, Figure 9 shows there is no discontinuity in the likelihood of manual eradication, an alternative policy for which we have geo-referenced data since 2007. Thus, the decrease of aerial spraying on the exclusion area was not compensated for by an increase in manual eradication, and hence our estimates reflect only the causal effect of the policy change in spraying. Moreover, the right panel of this figure shows there is no discontinuity in terms of altitude around the cutoff, which is a key predictor of yields (see Mejia and Restrepo, 2013b).

All these tests make us confident that our estimates provide a reasonable approximation to the conditional expectation of spraying and cultivation based on the location of cells, and that we are capturing a real discontinuity near the 10km threshold induced by the diplomatic friction since 2006.

⁸Other specifications yield similar patterns. However, as we allow larger bandwidths or use polynomials with lower degrees, we obtain larger point estimates for the difference in cultivation during 2004. Our view is that, although on average our approximation for the conditional expectation function is valid, it may fail for particular years. This potential failure lends support for complementary approaches, as the conditional difference in differences estimator we analyze below.

⁹Again, we focus on the specification using the local quadratic approximation to the conditional expectation. We obtain very similar results for the other specifications and for different choices of bandwidths.

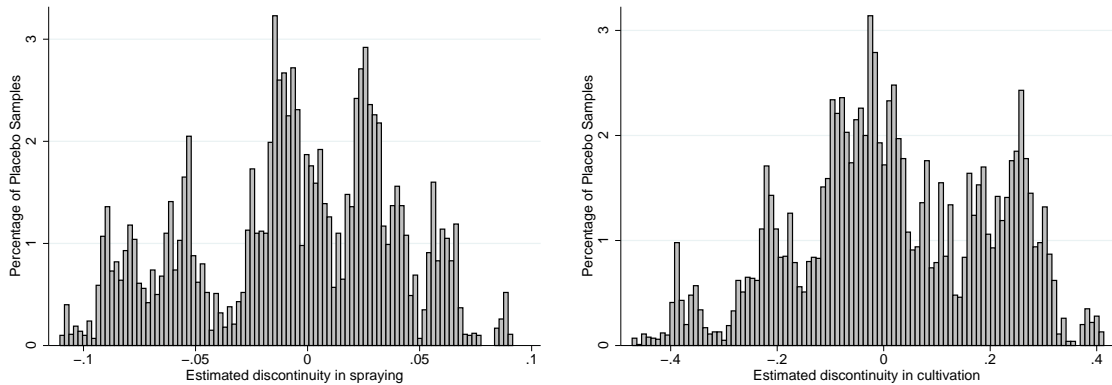


Figure 8: Distribution of discontinuity estimates obtained from Placebo cutoffs located 10 to 20km away from the frontier

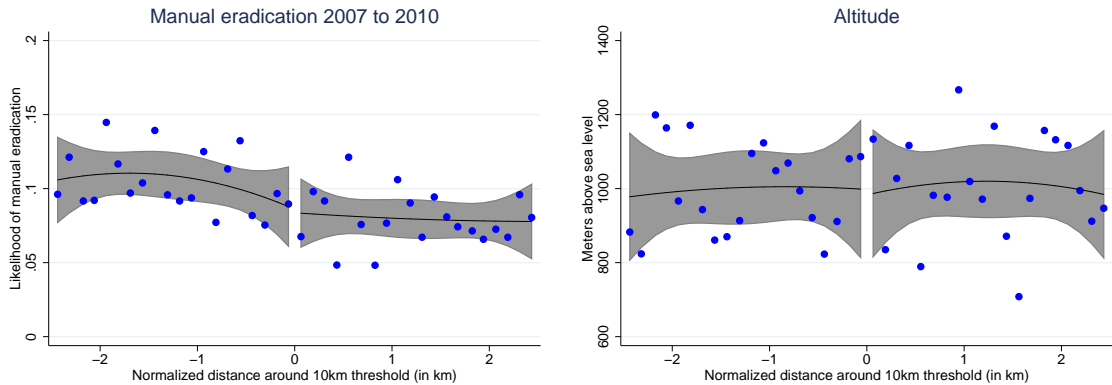


Figure 9: Probability of manual eradication and altitude around the 10 km cutoff.

5.1 Estimating local treatment effects

To quantify the treatment effect of spraying on illegal coca cultivation, we compute 2SLS estimates using the discontinuity in spraying as the instrument. In particular, we estimate the fuzzy regression discontinuity model:

$$Y_{it} = \beta_0 + \beta_1 S_{it} + f(\widehat{D}_i, X_i) + v_{it}, \quad (3)$$

instrumenting S_{it} with the dummy $1\{\widehat{D}_i > 0\}$ (so that equation 1 corresponds to the first stage).

Table 2 presents our estimates using different approximations of the conditional expectation $f(\widehat{D}_i, X_i)$. Columns 1 to 4 present results obtained by pooling all years from 2006 to 2010; while columns 5 to 8 present estimates for the years 2008 to 2010— when Colombia effectively stopped spraying in the exclusion area. In panel A we present estimates approximating f with a linear function of distance to the 10 km cutoff, latitude and longitude. We restrict the sample to cells within 3km, 2.75km and 2.5km of the cutoff in different columns as indicated in the bottom rows of the panel. Our results in column 1 indicate that a 10

percentage point increase in the likelihood of aerial spraying reduces cultivation by 0.35 hectares per squared kilometer (standard error=0.14); this result is in line with the reduced form estimates in the previous section. The remaining columns yield similar estimates.

In panel B we use a quadratic polynomial and in panel C a cubic polynomial. Despite the similar findings and the fact that we restrict our sample to be close to the cutoff, one may be concerned that these high order polynomials end up giving more weight to observations far from the cutoff (see Imbens and Geilman, 2012). To address these concerns we present results using a linear and a quadratic local approximation to the conditional expectation. We use Imbens and Kalyanaraman (2012) optimal bandwidth (4.7 km) in columns 1 and 4 and smaller bandwidths of 4km and 3km in the remaining columns, as indicated in the bottom row of the panel. Reassuringly, we find similar estimates in this case and for different choices of bandwidths. If anything, only the precision seems to change, as expected, when we use smaller bandwidths.

Overall, our estimates indicate a negative local average treatment effect of the likelihood of aerial spraying on coca cultivation. Our estimates suggest that a 10 percentage point increase in the likelihood of aerial spraying reduces coca cultivation by about 0.3 hectares per squared kilometer, though the results vary slightly depending on the specification.

In practice, we believe that part of our estimate captures the possibility that coca farmers reallocate their crops to the exclusion area, which seems reasonable given the proximity between cells. However, this implies we are over-stating the real effect of spraying on overall cultivation, and does not rule out our conclusion that farmers respond rationally to the increase in enforcement; it simply suggests another margin of response. However, one additional piece of evidence suggests that reallocation may not be that pervasive. When we estimate the effect of the discontinuity in enforcement on the likelihood of cultivation (the extensive margin, rather than the intensive margin), we find no effects (not reported to save space). This suggests that cultivation in the exclusion region increases within cells, and not because farmers move to new cells in this area. In any case, we cannot entirely rule out the extent of reallocation of coca crops, and our point estimates remain an upper bound of the effects of spraying on coca cultivation.

6 Conditional differences in differences estimates

Let $T_i = 1\{\widehat{D}_i > 0\}$ and \bar{Y}_i be the average cultivation in grid i from 2000 to 2005. We are interested in estimating the treatment effect of T_i on $Y_{it} - \bar{Y}_i$, for $t \geq 2006$ (or $t \geq 2008$ —when Colombia effectively stopped spraying in the exclusion area). This effect informs us about changes in cultivation brought about by the differences in enforcement induced by the diplomatic friction.

In the previous section we exploited the geographical discontinuity in enforcement around the 10 km cutoff to identify this effect. We relied on comparing cells near the cutoff and controlling for potential differences using a smooth function of their geographic location.

In this section, we present complementary estimates exploiting within-cell variation and

Table 2: Fuzzy RD estimates of the local average treatment effect of spraying on cultivation around the 10 km cutoff.

<i>Sample:</i>	From 2006 to 2010			From 2008 to 2010		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Linear global polynomial</i>						
Effect of spraying:	-3.508** (1.403)	-3.084** (1.369)	-2.918** (1.455)	-2.257** (1.050)	-1.869* (1.014)	-2.313* (1.203)
Excluded instrument F statistic	41.6	42.7	35.7	44.8	47.7	35.0
<i>Panel B: Quadratic global polynomial</i>						
Effect of spraying:	-3.756*** (1.345)	-3.472*** (1.338)	-3.260** (1.418)	-2.392** (1.019)	-2.102** (0.996)	-2.561** (1.186)
Excluded instrument F statistic	41.8	42.7	35.8	44.6	47.7	35.1
<i>Panel C: Cubic global polynomial</i>						
Effect of spraying:	-4.144** (2.112)	-5.394* (3.229)	-7.298* (4.383)	-3.085* (1.673)	-4.843* (2.803)	-4.539 (2.781)
Excluded instrument F statistic	15.1	8.0	5.6	15.9	7.8	7.0
<i>Sample of grids:</i>	$\pm 3km$	$\pm 2.75km$	$\pm 2.5km$	$\pm 3km$	$\pm 2.75km$	$\pm 2.5km$
<i>Observations:</i>	13265	11885	10560	7959	7131	6336
<i>Panel D: Local linear regression</i>						
Effect of spraying:	-2.980*** (1.019)	-3.183*** (1.163)	-3.628** (1.537)	-1.790** (0.765)	-1.933** (0.889)	-2.526** (1.206)
Excluded instrument F statistic	74.0	58.2	34.5	84.8	64.8	37.6
<i>Panel E: Local quadratic regression</i>						
Effect of spraying:	-3.179*** (1.233)	-3.591** (1.516)	-3.791* (2.178)	-1.900** (0.944)	-2.402** (1.167)	-3.241* (1.926)
Excluded instrument F statistic	30.0	20.8	9.8	30.5	21.5	9.1
<i>Bandwidth:</i>	$b = IK$	$b = 4km$	$b = 3km$	$b = IK$	$b = 4km$	$b = 3km$
<i>Observations:</i>	22340	18495	13265	13404	11097	7959

Notes: The table presents fuzzy regression discontinuity estimates of the effect of differential aerial spraying on cultivation around the 10 km cutoff. Columns 1 to 4 pool the years 2006 to 2010; while columns 5 to 8 pool the years 2008 to 2010. Results use different approximations of the conditional expectation, as indicated in the top of each panel. Moreover, sample restrictions and the choice of bandwidth is indicated at the bottom of each panel. In all models we include municipality and year specific intercepts. Standard errors robust against heteroskedasticity and serial correlation within cells are reported in parenthesis. Estimates with *** are significant at the 1%, those with ** are significant at the 5%, and those with * are significant at the 10%.

use the data from 2000 to 2005 to construct counterfactuals for cultivation and spraying in a given cell. In particular, we exploit the following conditional independence assumption:

$$Y_{itd} - \bar{Y}_i \perp T_i | Z_i, \{Y_{it'}, S_{it'}\}_{t' \leq 2005}, \quad (4)$$

where Y_{itd} is the potential cultivation in cell i and year t for cells in the sprayed area, $d = 1$, and cells in the exclusion area, $d = 0$. The assumption states that once we condition on the whole history of cultivation and spraying in a cell ($\{Y_{it'}, S_{it'}\}_{t' \leq 2005}$), and cell characteristics,

Z_i (including a polynomial in altitude—determining yields—and municipality fixed effects), the *change* in potential cultivation from 2006 onward would be equal for cells in the sprayed and exclusion areas (in the absence of changes in spraying).¹⁰ In other words, the assumption requires that all differences between cells in the exclusion and sprayed areas are captured by their history of spraying and cultivation, and the observed geographic characteristics. Any change in cultivation not predicted by these observables is attributed to changes in aerial spraying since 2006.

We believe this is a plausible assumption. Including lags of cultivation and spraying takes into account the fact that, before 2006, there was a differential behavior of cultivation in both areas (see Figure 10, which plots the difference in spraying and cultivation between the sprayed and exclusion areas for each year). This implies a violation of traditional difference in difference estimates and requires us to condition on the observed paths of cultivation and spraying, as we do here.

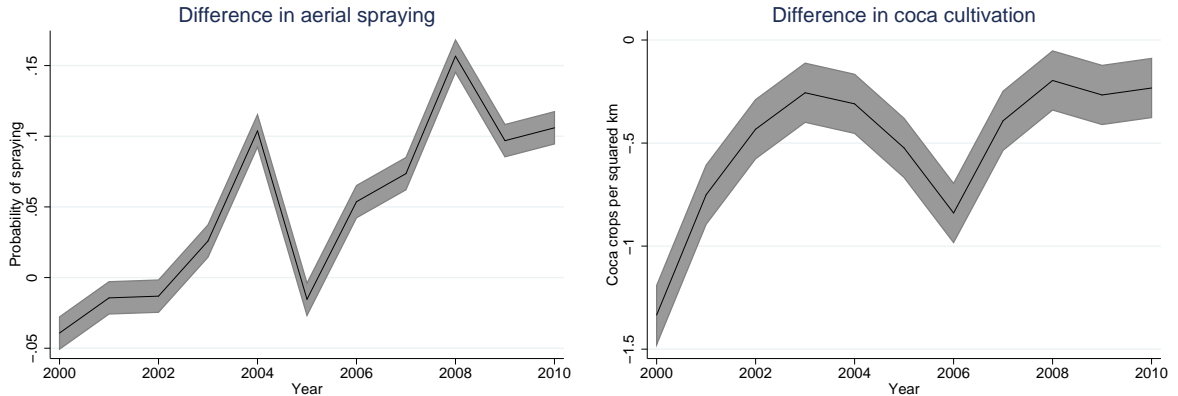


Figure 10: Difference in coca cultivation and spraying between the sprayed and exclusion areas from 2000 to 2010.

We exploit the above CIA in several ways to estimate the effect of being in the sprayed region during years in which Colombia agreed not to spray the exclusion area. First, we start by running the regression:

$$Y_{it} - \bar{Y}_i = \beta_t T_i + \delta_t + \sum_{t' \leq 2006} \Gamma \cdot (Y_{it'}, S_{it'})' + \Theta Z_i' + \varepsilon_{it}, \forall t \geq 2006 \quad (5)$$

Here, β_t identifies the effect of being in the sprayed region (relative to the exclusion region) during year t as long as the conditional expectation of the outcome is linear in the covariates. We compute standard errors clustering at the cell level and robust to heteroskedasticity.

Column 1 in Table 3 presents the regression estimates on cultivation separately by year, and also pooling together the years 2006 to 2010 and 2008 to 2010—when Colombia effectively stopped spraying the exclusion area. Our estimates show that cultivation fell in the

¹⁰We use the change in potential cultivation instead of its level to remove any permanent difference between cells not captured by the conditioning set. In theory, we do not need to remove the average cultivation, as this is already in the conditioning set. In practice, this helps to control for potential sources of misspecification in the model for the propensity score used below.

Table 3: Conditional differences in differences estimate of being in sprayed region.

	Estimates for cultivation				Estimates for spraying			
	Controlling for propensity score				Controlling for propensity score			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A: Pooled-years estimates.</i>								
Pooling 2006 to 2010:	-0.239*** (0.015) [0.053]	-0.227*** (0.022) [0.052]	-0.150*** (0.021) [0.024]	-0.222*** (0.021) [0.052]	0.077*** (0.002) [0.011]	0.100*** (0.004) [0.011]	0.104*** (0.004) [0.007]	0.095*** (0.004) [0.011]
<i>Observations</i>	54400	54400	54400	54400	54400	54400	54400	54400
Pooling 2008 to 2010:	-0.153*** (0.012) [0.037]	-0.140*** (0.019) [0.036]	-0.067*** (0.020) [0.022]	-0.132*** (0.018) [0.036]	0.104*** (0.003) [0.011]	0.124*** (0.004) [0.011]	0.122*** (0.005) [0.007]	0.119*** (0.004) [0.011]
<i>Observations</i>	32640	32640	32640	32640	32640	32640	32640	32640
<i>Panel B: Year-by-year estimates.</i>								
Estimate for 2006:	-0.581*** (0.045)	-0.571*** (0.051)	-0.517*** (0.048)	-0.579*** (0.050)	0.057*** (0.005)	0.072*** (0.006)	0.082*** (0.007)	0.072*** (0.006)
Estimate for 2007:	-0.154*** (0.042)	-0.146*** (0.043)	-0.032 (0.039)	-0.137*** (0.040)	0.019*** (0.006)	0.054*** (0.009)	0.073*** (0.009)	0.047*** (0.008)
Estimate for 2008:	-0.091*** (0.023)	-0.059** (0.029)	0.022 (0.033)	-0.049* (0.028)	0.129*** (0.005)	0.159*** (0.007)	0.155*** (0.008)	0.150*** (0.006)
Estimate for 2009:	-0.204*** (0.017)	-0.218*** (0.022)	-0.156*** (0.022)	-0.206*** (0.020)	0.079*** (0.004)	0.105*** (0.005)	0.101*** (0.006)	0.100*** (0.005)
Estimate for 2010:	-0.163*** (0.022)	-0.143*** (0.026)	-0.066** (0.027)	-0.141*** (0.025)	0.102*** (0.005)	0.110*** (0.005)	0.111*** (0.006)	0.108*** (0.005)
<i>Observations per year</i>	10880	10880	10880	10880	10880	10880	10880	10880

Notes: The table presents conditional differences in differences estimates of the effect of being in the sprayed region (relative to the exclusion areas) on cultivation and spraying. Columns 1 and 5 present linear regressions. In columns 2 and 6 we reweigh the data using the estimated propensity score. In columns 3 and 7 we stratify on the estimated propensity score. Finally, in columns 4 and 8 we match observations on the propensity score. The propensity score is estimated with a Probit model using data for cultivation and spraying from 2000 to 2005 as explanatory variables. The top panel presents results pooling several years together; while the bottom panel presents year by year estimates. Standard errors robust against heteroskedasticity and serial correlation within cells are reported in parenthesis. Estimates with *** are significant at the 1%, those with ** are significant at the 5%, and those with * are significant at the 10%.

sprayed area relative to the exclusion region in all years since 2006 and the effects are all significant at traditional levels. Pooling the years 2006 to 2010 together (or 2008 to 2010), we find that cells in the sprayed area had 0.24 less hectares of coca per squared km (standard error=0.015) as a consequence of the spraying. As before, we also present standard errors robust against spatial correlation among cells within 5 km of each other. These are presented in squared brackets only below the main estimates in Table 3. Though, in this case, they are considerably larger than the traditional ones, they do not change any of our conclusions.

Consistency of the previous estimates requires the conditional expectation of cultivation and aerial spraying to be linear in the covariates. To relax this assumption, we follow several strategies in which we control non-parametrically for the propensity score, $\lambda_i = P[T_i = 1|Z_i, \{Y_{it'}, S_{it'}\}_{t' \leq 2005}]$. We estimate the propensity score, $\hat{\lambda}_i$, using a Probit model (not reported to save space), making this approach semi-parametric.

In column 2 we reweight the regression in equation 5 by the propensity score (see Hirano, Imbens, and Ridder, 2003). In particular, we weight observations in the sprayed area by $p/(1-p)$, where p is the fraction of cells in this area, and observations in the exclusion area by $\hat{\lambda}_i/(1-\hat{\lambda}_i)$, with $\hat{\lambda}_i$ being the estimated propensity score of the grid. This method ensures that all covariates are balanced and set to the distribution of the sprayed region. Once reweighted, the regression estimate equals the average treatment effect of being in an area currently sprayed.¹¹ Besides reweighting by the propensity score, we also control linearly for all covariates in the regression. This is known as a double-robust regression, as it provides consistent estimates as long as the propensity score is correctly specified or the linear covariates are an accurate model for the conditional expectation of spraying and cultivation (see Imbens and Wooldridge, 2009). As can be seen from the results in column 2, the results change little relative to column 1, suggesting that the linear controls were already capturing most of the relevant heterogeneity in cultivation and spraying dynamics before the intervention.¹²

The role of reweighting the data by the propensity score can be seen graphically in Figure 11. We plot the difference in cultivation and spraying for each year after reweighting the data using the propensity score as described above. The right panel shows that now, cultivation is balanced between the sprayed and exclusion areas before 2006. A similar pattern emerges for spraying, although dynamics were already roughly balanced in the raw data.

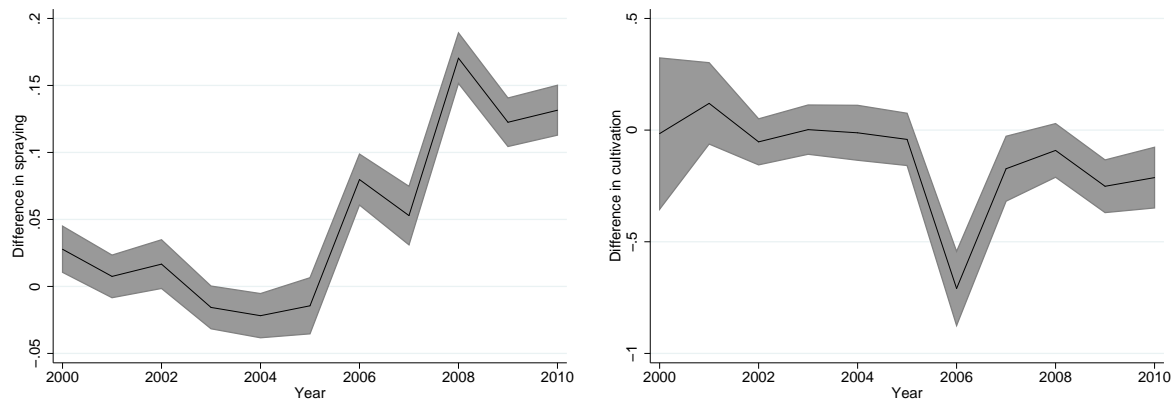


Figure 11: Re-weighted difference in coca cultivation and spraying between the sprayed and exclusion areas from 2000 to 2010 relative to the average before 2006. We weight observations in the exclusion area by the estimated likelihood ratio based on the estimated propensity score.

¹¹We can also estimate the average treatment effect, but this requires weighting by $1/\hat{\lambda}_i$ the observations in the sprayed area. However, there are values with very low estimated propensity scores that make this exercise imprecise. In any case, our results are similar. Moreover, the average treatment effect on the treated seems like the relevant object to evaluate the actual effect of spraying on cultivation.

¹²We report the usual regression standard errors clustering at the grid level. These errors ignore the fact that the propensity score is estimated in a previous stage. However, as suggested by Hirano, Imbens, and Ridder (2003), these standard errors are actually conservative, relative to adjusted ones. We compute an alternative set of bootstrapped standard errors taking into account the estimation of the propensity score and obtained slightly smaller standard errors (not reported).

In column 3 we follow another strategy and stratify on the propensity score as in Angrist (1998) and Dehejia and Wahba (1999). In particular, we group observations by their propensity score in 20 equal bins covering the $(0, 1)$ interval. The j -th bin contains grids with an estimated propensity score between $(j - 1) \times 0.05$ and $j \times 0.05$. For each bin, we estimate equation 5 separately, and use weighted averages of all these estimates to obtain an estimate for β_t . We obtain the variance of β_t as a weighted sum of the variances for each bin as well. We weight each estimate by the number of observations in the bin from the sprayed region. This guarantees that we estimate the average treatment effect on the treated. This approach has the advantage of not imposing any functional form on the conditional expectation as a function of the propensity score, but of course is limited by the size of our bins. Again, we control locally for all covariates when estimating equation 5 for each bin. This partly controls for differences in the propensity score within bins and misspecification of the propensity score. Our results are similar to the basic regression estimates in column 1, though we find a smaller reduction in cultivation.

Finally, in column 4 we do Kernel matching on the propensity score. This works by finding, for each grid in the sprayed region, others in the exclusion area within a band around its estimated propensity score, and weighting them by a triangular Kernel that assigns less weight to distant grids. Reweighting the regression using these weights produces an estimate of the average treatment effect on the treated. The reweighting guarantees that every grid in the sprayed region is compared to an average of grids with similar propensity score in the exclusion region, and thus controls non-parametrically for the propensity score. Again, we also include the covariates in the regression linearly, which control partly for differences in the propensity score within the kernel of an observation. Our results vary little with respect to the traditional regression estimates in column 1.

Columns 5 to 8 present analogous estimates using the likelihood of spraying as the dependent variable. We find a large increase in the likelihood of spraying in the sprayed region, especially after 2008, and a smaller increase in 2006 and 2007, consistent with the fact that the diplomatic friction began in 2006, but became binding only from 2008 onwards. When pooling the years 2006 to 2010 together in column 5, we find that cells in the sprayed area were indeed 7.7 percentage points more likely to be sprayed (standard error=0.002), independently of their past paths of cultivation. Thus, the diplomatic friction caused a clear change in the likelihood of enforcement among cells in different sides of the 10km band near the frontier.

All the same, the estimates in this section suggest that grids in the sprayed region were approximately 10 percentage point more likely to be sprayed from 2006 to 2010, independently of their past levels of cultivation or observable characteristics. As a consequence, farmers in the region reduced cultivation by approximately 0.22 hectares per square kilometer.¹³

¹³The fact that this is smaller than our regression discontinuity estimates could be due to two things. First, because they are, in theory, different objects. In columns 2 to 4 of Table 3 we estimate an average treatment effect on the sprayed grids (and the regression in column 1 produces a mix between the ATT and ATE), while regression discontinuity estimates a local effect. Second, there may be more reallocation of crops to the exclusion region near the 10 km cutoff. This implies that this indirect margin is more relevant for the regression discontinuity estimates, making them overstate the deterrent effect.

7 Cost benefit analysis of aerial spraying

As discussed in the introduction, the aerial spraying program is the largest component among the supply reduction efforts implemented under *Plan Colombia*. Between 2000 and 2008, \$585 million were allocated to the eradication program, whereas \$62.5 were allocated to air interdiction and \$89.3 to coastal and river interdiction by the military forces, and \$152.7 to interdiction activities carried out by the Colombian Police (see the U.S. Government Accountability Office - GAO, 2008).

Our regression discontinuity estimates suggest that a 10 percentage points increase in the likelihood of spraying a 1 squared km grid leads to a reduction of approximately 0.3 hectares of coca. Our conditional differences in differences points to a reduction of approximately 0.22 hectares. Since 1 square kilometer contains 100 hectares, these estimates imply that to reduce cultivation by 1 hectare during a given year, between 33 and 45 hectares would have to be sprayed.

It is estimated that the average direct cost to the U.S. per hectare sprayed is about \$750 (see Walsh et al., 2008). Thus, reducing cultivation by 1 hectare through financing spraying campaigns costs the U.S. between \$24,750 and \$33,750 dollars. Additionally, for every dollar spent by the U.S., Colombia spends about 2.2 dollars (aerial spraying campaigns are jointly financed by the countries), making the overall total cost range between \$79,200 and \$108,000 per hectare of coca crops reduced. To put these numbers in perspective, the coca leaf in one hectare produces about 1.2 kgs of cocaine per harvest, with a farmgate market value of about \$4,200 dollars.

From a drug policy perspective, it is more informative to calculate the benefits in terms of the reduction of kilograms of cocaine in consumer markets. We do not have estimates of the social benefits of such reductions, but at least we can compare the cost to that of other policies achieving a similar objective. To do so, we use the estimates in Mejia and Restrepo (2013) obtained by calibrating a model of downstream cocaine markets. The authors find that a 1% reduction in coca cultivation reduces cocaine in consumer markets by 0.0025%. This elasticity is small for several reasons. First, cultivation represents only a small fraction of the total market value of cocaine in consumer markets. Thus, an increase in the price of coca leaf caused by spraying translates into a small increase in consumer prices. Second, demand is inelastic, so the small increase in prices barely affects consumption. Finally, downstream markets adjust to the shock by substituting towards cheaper inputs of production, such as better chemical precursors and technologies to produce more cocaine per hectare, by demanding more cocaine from other source countries, or by switching to better transportation techniques, partially offsetting the effect of the shock on the supply of cocaine.

Total coca cultivation in Colombia was about 80,000 hectares during our period of analysis. Reducing this by 1% (800 hectares), would cost the U.S. \$20-\$27 million dollars per year. However, this investment would reduce the supply of cocaine in its territory only by about 0.0025%, which equals 12.5 kg. This implies that the marginal cost to the U.S. of reducing retail quantities of cocaine by 1 kg by subsidizing aerial spraying in Colombia is in the order of \$1.6- \$2.16 million dollars. These are large magnitudes, but are similar to the

estimates reported by Mejia and Restrepo (2015) using an entirely different methodology. To put them in perspective, the price of 1 kg of cocaine in retail markets is about \$150,000 per kilogram.

The conclusion from this exercise is that aerial spraying is a very costly policy from a supply-reduction perspective. In particular, the policy is significantly more costly than other alternatives achieving the same objective. The estimated marginal cost to the U.S. of reducing retail quantities of cocaine by 1 kg is estimated at \$175,000 dollars by subsidizing interdiction policies in Colombia (Mejia and Restrepo, 2013), or \$8,250 and \$68,750 dollars by funding treatment and prevention efforts, respectively, in the U.S. (MacCoun and Reuter, 2001). Thus, despite being able to reduce coca cultivation by affecting farmers incentives, aerial spraying has only small effects on cultivation. These effects translate to even smaller effects on downstream markets for the reasons emphasized above, making it a costly supply-reduction policy. If on top of that we add the share of the costs paid by Colombia and the alleged negative effects on health (Camacho and Mejia, 2015), other legal crops, the environment (see Relyea, 2005 and Dávalos et al., 2011), and the socio-economic conditions of coca-producing areas (see Rozo, 2014), the policy looks even less favorable.

8 Concluding remarks

In this paper we explored the deterrent effects of enforcement on illegal behavior. We did so in the context of illegal coca cultivation in Colombia. We find that aerial spraying of coca crops – a particular type of enforcement aimed at disrupting the production of an illicit good (cocaine) – induces farmers to reduce coca cultivation. Our findings are aligned with the key insight from the economic analysis of crime, suggesting that the decision to engage in illegal activities is rational and, as such, responds to the likelihood of enforcement.

Our main contribution is to present a clean and credible source of identification for the effects of enforcement on illegal markets. In particular, we exploit a diplomatic friction between the governments of Colombia and Ecuador over the possible negative effects of spraying campaigns in the Colombian territory bordering Ecuador. This diplomatic friction ended in a compromise by the Colombian government to stop spraying campaigns with glyphosate within a 10 km band along the border with Ecuador in 2006.

We use a regression discontinuity design, exploiting the arbitrary 10 km cutoff and a conditional differences in differences estimator comparing similar cells with different treatment probabilities to uncover the causal effects of spraying on coca cultivation. Both methodologies point to a negative and significant effect of the program on coca production. In particular, both methodologies show that cells in the region that continued to be sprayed were approximately 10 percentage points more likely to be sprayed than cells in the exclusion area. In consequence, coca cultivation decreased in this region by about 0.3 hectares (regression discontinuity estimates) or 0.22 hectares (conditional differences in differences estimate) per squared kilometer.

Despite reducing coca cultivation, aerial spraying in Colombia has only small effects in downstream markets. We estimate that reducing the Colombian coca cultivation by 1%

(about 800 hectares) would cost the U.S. between \$20 and \$27 million dollars per year. However, this investment would reduce the supply of cocaine in its territory by only 12.5 kg of cocaine less per year. Hence, the cost of reducing cocaine retail supply by 1 kg via aerial spraying campaigns is at least \$1.6 billion dollars per year. Other policies, such as treatment and prevention, or interdiction efforts in Colombia, would be significantly more cost effective in curbing drug supply.

Remark: *While this version of the paper was being completed, the Colombian government announced that it will stop the aerial spraying program. The decision was taken based on the possible health effects that the program might be having on the populations exposed to the herbicide used in the aerial spraying campaigns. The findings in this paper indicate that, on top of its negative collateral consequences on health, the aerial spraying program is not a cost-effective strategy in reducing cocaine supply. Thus, the Colombian government decision is unlikely to cause a large surge in cocaine supply.*

9 References

Abadie, A., J., and Imbens, G. (2002). Instrumental Variables Estimates of the Effect of Subsidized Training on the Quantiles of Trainee Earnings. *Econometrica*, 70(1): 91-117.

Abadie, A., and Imbens, G. (2006). Large Sample Properties of Matching Estimators for Average Treatment Effects. *Econometrica*, 74(1): 235-267.

Abadie, A., Angrist, J., and Imbens, G. (2010). On the failure of Bootstrap for Matching Estimators. *Econometrica*, 76(6): 1537-1557.

Abadie, A., and Imbens, G. (2011). Matching on the Estimated Propensity Score. National Bureau of Economic Research, Working Paper N.15301.

Andreoni, J., Erard, B., and Feinstein, J. (1998). Tax Compliance. *Journal of Economic Literature*, 36(2): 818-860.

Angrist, J. (1998). Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants. *Econometrica*, 66(2): 249-288.

Angrist, J. and Pischke, J. (2009). *Mostly Harmless Econometrics*, Princeton University Press, New Jersey.

Bar-Ilan, A. and Sacerdote, B. (2001). The Response to Fines and Probability of Detection in a Series of Experiments. National Bureau of Economic Research, Working Paper No. 8638.

Becker, G. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy*, 76(2): 169-217.

Benabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *The Review of Economic Studies*, 70(3): 489-520.

Benabou, R. and Tirole, J. (2006). Incentives and Prosocial Behavior. *The American Economic Review*, 96(5): 1652-1678.

Beron, K., Tauchen, H., and Witte A. (1992). The Effect of Audits and Socioeconomic Variables on Compliance. In *Why People Pay Taxes*, edited by J. Slemrod. Ann Arbor: Univ. of Michigan Press.

Bogliacino, F. and Naranjo, A. (2012). Coca Leaves Production and Eradication: A General Equilibrium Analysis. *Economics Bulletin*, 32(1): 382-397.

Buonanno, P. and Mastrobuoni, G. (2012). Police and crime: Evidence from dictated delays in centralized police hiring. Institute for the Study of Labor, Discussion Paper N. 6477.

Caliendo, M. and Kopeinig, S. (2005). Some Practical Guidance for the Implementation of Propensity Score Matching. Institute for the Study of Labor, Discussion Paper N. 1588.

Camacho, A. and Mejia, D. (2015). The Health Consequences of Aerial Spraying of Illicit Crops: The Case of Colombia. Center for Global Development WP 408, June.

Cameron, S. (1988). The Economics of Crime Deterrence: A Survey of Theory and Evidence. *Kyklos*, 41(2): 301-23.

Conley, Timothy G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1): 1-45.

Corman, H. and Mocan, H. (2000). A Time- Series Analysis of Crime, Deterrence, and Drug Abuse in New York City. *American Economic Review*, 90(3): 584-604.

Davalos, L., Bejarano, A., Hall, M., Correa, H., Corthals, A., and Espejo, O. (2011). Forests and Drugs: Coca-Driven Deforestation in Tropical Biodiversity Hotspots. *Environ. Sci. Technol.*, 45(4):1219-1227.

Dehejia, R. and Sadek, W. (1999). Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs. *Journal of the American Statistical Association*, 94(448): 1053-1062.

Dehejia, R. (2004). Program Evaluation as a Decision Problem. *Journal of Econometrics*, 125: 141-173.

Di Tella, R. and Schargrodsy, E. (2004). Do Police Reduce Crime? Estimates Using the Allocation of Police Forces After a Terrorist Attack. *American Economic Review*, 94(1): 115-133.

Dion, M.L. and Russler, K. (2008). Eradication efforts, the State, Displacement and Poverty: Explaining Coca Cultivation in Colombia During Plan Colombia. *Journal of Latin American Studies*, 40(3): 399-421.

Draca, M., Machin, S., and Witt, R. (2011). Panic on the Streets of London: Police, Crime and the July 2005 Terror Attacks. *American Economic Review*, 101(5): 2157-81.

Dubin, J., Graetz, M., and Wilde, L. (1987). Are We a Nation of Tax Cheaters? New Econometric Evidence on Tax Compliance. *American Economic Review*, 77: 240-45.

Eck, J. and Maguire, E. (2000). Have Changes in Policing Reduced Violent Crime? An Assessment of the Evidence. in A. Blumstein and J. Wallman, eds., *The crime drop in America*. New York: Cambridge University Press, pp. 207-65.

Evans, W. and Owens, E. (2007). COPS and crime. *Journal of Public Economics*, 91(1-2): 181-201.

Frey, B. (1997). Not Just for the Money: An Economic Theory of Personal Motivation. Edward Elgar Pub; 1 Ed edition.

Garcia, J., Mejia, D., and Ortega, D. (2013). Police Reform, Training and Crime: Experimental evidence from Colombia's Plan Cuadrantes. Documento CEDE.

Gelman, A. and Imbens, G. (2014). Why High-Order Polynomials should not be used in Regression Discontinuity Designs. NBER, Working Paper N. 20405.

Hirano, K., Imbens, G. W. and Ridder, G. (2003). Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score. *Econometrica*, 71(4): 1161–1189.

Heckman, J., Ichimura, H., Smith, J. and Todd, P. (1998). Characterizing Selection Bias Using Experimental Data. *Econometrica*, 66: 1017-1098.

Hsiang, S. (2010). Temperatures and Cyclones Strongly Associated with Economic Production in the Caribbean and Central America. Proceedings of the National Academy of Sciences.

Ibanez, M. and Carlsson, F. (2010). A survey-based choice on coca cultivation. *Journal of Development Economics*, 93: 249-273.

Imbens, G. and Kalyanaraman (2012). Optimal Bandwidth Choice for the Regression Discontinuity Estimator. *The Review of Economic Studies Advance Access*, pp: 1-27.

Imbens, G. and Lemieux, T. (2008). Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics*, 142(2): 615-635.

Imbens, G. and J.M. Wooldridge. (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature*, 47(1): 5-86.

Kagan, R. (1989). On the Visibility of Income Tax Law Violations. in Taxpayer Compliance: Social science perspectives 107, Jeffrey A. Roth & John T. Scholz eds.

Klepper, S. and Nagin, D. (1989). Tax Compliance and Perceptions of the Risks of Detection and Criminal Prosecution. *Law & Society Review* 23: 209-40.

Lee, D. and Lemieux, T. (2009). Regression Discontinuity Design in Economics. National Bureau of Economic Research, Working Paper 14723.

Ludwig, J. and Miller, D. (2005). Does Head Start Improve Children's Life Chances? Evidence from a Regression Discontinuity Design. National Bureau of Economic Research, Working Paper 11702.

Levitt, S. (1997). Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime. *American Economic Review*, 87(3): 270-90.

Levitt, S.(2002). Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime: Reply. *American Economic Review*, 92(4): 1244-50.

Marvell, T. and Moody, C. (1996). Specification Problems, Police Levels, and Crime Rates. *Criminology*, 34(4): 609-46.

McCrary, J. (2002). Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime: Comment. *American Economic Review*, 92(4): 1236-43.

McCrary, J. (2008). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics*, 142(2): 698-714.

Mejia, D., and Rico, D. (2011). The Microeconomics of Cocaine Production and Trafficking in Colombia. in *Anti-drug policies in Colombia: Successes, failures and lost opportunities* (A. Gaviria and D. Mejia, eds), ch. 1. Ediciones UniAndes, Bogota.

Mejia, D. and Restrepo, P. (2013). The Economics of the War on Illegal Drug Production and Trafficking. Documento CEDE No. 54.

Mejia, D. and Restrepo, P. (2013b). Bushes and Bullets: Illegal Cocaine Markets and Violence in Colombia. Documento CEDE No. 53.

Menniger, K. (1968). *The Crime of Punishment*. Viking Press, New York.

Moreno-Sanchez, R., Kraybill, D. and Thompson, S. (2003). An Econometric Analysis of Coca Eradication Policy in Colombia. *World Development*, 31(2): 375-383.

Moya, A. (2005). Impacto de la Erradicación Forzosa y el Desarrollo Alternativo Sobre los Cultivos de Hoja de Coca. Facultad de Economía. Universidad de Los Andes, Bogotá, Colombia.

Relyea, R. (2005). The Impact of Insecticides and Herbicides on Biodiversity and Productivity of Aquatic Communities. *Ecological Society of America* 15(2): 618-627.

Reyes, L. (2014). Estimating the Causal Effect of Forced Eradication on Coca Cultivation in Colombian Municipalities. *World Development*, 61: 70-84.

Rosenbaum, P., and D. Rubin (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika*, 70(1): 41-50.

Rozo, S.V. (2014). On the Unintended Consequences of Anti-Drug Programs in Producing Countries. Online-paper collection Association for Public Policy Analysis and Management.

Stigler, George J. (1970). The Optimum Enforcement of Laws. *Journal of Political Economy*, 78: 526-536.

Stuart, E. (2010). Matching Methods for Casual Inference: A Review and a Look Forward. *Statistical Science*, 25(1):1-21.