

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

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

This version: July 23, 2014

Abstract

This paper studies the effects of enforcement on illegal behavior in the context of coca cultivation in Colombia. We explore the deterrent effects of a large aerial spraying program designed to curb cocaine supply. We exploit variation induced by a diplomatic friction between the governments of Colombia and Ecuador over the possible negative effects of spraying campaigns over the Ecuadorian territory in the border with Colombia. As a result of this friction, Colombia pledged to stop spraying campaigns in a 10 km band around the border with Ecuador in 2006. We estimate the effects of spraying on cultivation by regression discontinuity around the 10km threshold and conditional-differences-in-differences, using satellite data for 1-squared-km cells. Our results suggest that spraying one additional hectare reduces coca cultivation by about 0.02 to 0.065 hectares, consistent with the view that enforcement reduces illegal behavior. However, these effects are too small to make aerial spraying a cost-effective anti-narcotic strategy.

Keywords: Illegal Markets, War on drugs, Crime.

JEL Classification: O13, O33, O54 and Q18.

*We would like to thank SIMCI at the United Nations Office of Crime and Drugs (UNODC) in Colombia for their invaluable collaboration and for making the data available for this study. We would also like to thank Adriana Lleras-Muney, Adriana Camacho, Paola Guiliano, and Leah Boustan for their suggestions. We are also grateful with participants of the economics proseminar at UCLA and Universidad de los Andes for the comments received. All remaining errors are our own.

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

[‡]MIT, email: pascual@mit.edu

[§]UCLA, email: sandrarozo@ucla.edu

1 Introduction

Illegal activities such as counterfeiting, tax evasion, and the operation of illegal drug markets remain a serious problem throughout the globe. Yet, there is still an open debate on how to design and implement public policies to reduce their extent. The economics analysis of crime suggests that the decision to engage in illegal activities is rational, and as such, is shaped by incentives and penalties (see Becker, 1968 and Stigler, 1970). The central prediction is that enforcement reduces crime by increasing its costs. Despite its theoretical appeal, social scientists and pundits have raised several concerns against this view. In particular, critics have argued 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). Apart from the theoretical controversies, evidence on the role of enforcement in reducing illegal behavior is not abundant; in part, due to the lack of the exogenous sources of variation in enforcement required to uncover its causal impact.

Our paper contributes to the growing literature that attempts to estimate the causal effect of enforcement on illegal activities. We use the war on drugs in Colombia as a case study, focusing on the role of aerial spraying with herbicides on the extent of illegal coca cultivation. At least since 1996 Colombia has been the world largest cocaine producer and grower of coca crops (the raw input for cocaine production). Coca cultivation takes place in remote areas of the country with little institutional presence, where farmers face the risk of being detected and sprayed with herbicides by the government. When coca crops are sprayed with herbicides they are partially lost, which increases the cost of this illegal activity and reduces the farmer's incentives to pursue it. Our goal in this paper is to assess the effectiveness of this form of enforcement, and explore how it affects farmers' behavior.

For this purpose we exploit the geographic and time variation on aerial spraying induced by a diplomatic friction between the governments of Ecuador and Colombia around 2006. Around the year 2000, the Ecuadorian government alleged that Colombian aerial spraying campaigns near the frontier were causing health problems, productivity losses, and environmental damages in their territory. In response, the Colombian government committed to completely stop aerial spraying campaigns in a 10 km band around the international frontier with Ecuador at the beginning of 2006. The Colombian government relented to its commitment at the end of 2006 and continued spraying the band during 2007. Yet, aerial spraying was stopped definitely in the band since 2008, in response to a lawsuit filed by the Ecuadorian government at international courts.

We use satellite and geo-referenced data of coca cultivation and aerial spraying on 1-squared-km (100hectares) cells between 2000 to 2010, and estimate the effect of aerial spraying using two methodologies: First, we use a fuzzy regression discontinuity design and compare coca cultivation in cells near both sides of the 10km threshold. We show that aerial spraying changes discontinuously at the 10km threshold during the years in which Colombia committed not to spray the exclusion area (except in 2009), while all other covariates, including coca cultivation in previous years or 2008, do not. This provides us with a discontinuous change in the likelihood of enforcement that we can use to identify its effect on illegal coca cultivation. Importantly, we document that the reduction in spraying since 2006 near Ecuador was not compensated with other types of enforcement (i.e., manual eradication of

coca crops).

Additionally, we report results obtained by conditional-difference-in-differences. In particular, we compare the cultivation of illicit crops on cells within the exclusion area to similar cells located 10 to 20km away from the frontier – the area that continued to be sprayed during all years in our sample. Both groups of cells were exposed to aerial spraying before 2006, and since this year, only the latter cells continued to be sprayed (except in 2007). Thus, the difference in coca cultivation between both regions since 2006 can be attributed to the change in enforcement. To guarantee the comparability of both groups, we control for coca cultivation and spraying before the intervention, and use a variety of techniques to control non-parametrically for these predetermined characteristics.

Consistent with the view that illegal behavior is a rational choice, we find significant (but small) deterrent effects of spraying on coca cultivation. The regression discontinuity estimates imply that cells in the sprayed area near the cutoff had a 10% higher likelihood of being sprayed than cells in the exclusion area near the cutoff. As a result, cultivation was reduced from 0.3 to 0.6 hectares per squared kilometer in the former group relative to the later. Similarly, our estimates using the conditional-difference-in-differences methodology suggest that the areas that were exposed to aerial spraying after 2006 faced 10% higher likelihood of being sprayed and that, as a result, had on average 0.2 less hectares of coca per squared kilometer (relative to cells in the region not-sprayed). Both methodologies suggest that an increase in coca cultivation is reduced between 0.02 and 0.065 hectares per additional hectare sprayed (recall that 1 squared km equals 100 hectares, so one spraying one additional hectare increases the likelihood of spraying by 1%).

Our findings confirm the key insight from the economics of crime. Namely that 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 when compared to the costs of this policy. In particular, our largest point estimates suggest that to reduce coca cultivation by 1 hectare approximately 15.4 have to be sprayed. Moreover, it is possible that coca cultivation is in part displaced by aerial spraying campaigns, making the 15.4 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 spent by the U.S. in 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 15.4 hectares of coca, and reducing cultivation by 1 hectare per year, is about \$37,000 dollars, out of which the U.S. pays at least \$11,550. As we show in greater detail in the paper, these numbers imply that the marginal cost to the U.S. of reducing coca supply in retail markets by 1kg subsidizing aerial spraying policies in Colombia is at least \$462,000 dollars, which is in the eyeball park of the costs reported by Mejia and Restrepo (2013) using a different methodology. This is significantly larger than the same cost for other policies such as interdiction in Colombia (\$181,000 dollars; see Mejia and Restrepo, 2013), or treatment and prevention in the U.S. (\$8,250 and \$68,750 dollars, respectively; see MacCoun and Reuter, 2001). It is also large when compared to the retail price of 1kg 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 is a key case in terms of anti-narcotic policy. During our period of analysis, it was the largest cocaine producer nation covering nearly 70% of the total supply and a similar proportion of total coca cultivation in the Andean region. Second, effective supply reduction policies in Colombia have the potential to reduce the availability of cocaine and its associated harms throughout the globe. In fact, most of the cocaine produced in Colombia is exported and between 60% and 70% of the cocaine consumed worldwide is produced in this country (UNODC, 2012). Third, aerial spraying is the largest anti-drug program implemented in Colombia. It not only entails resources from the local government but also from the U.S. In particular, since the start 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 – aerial spraying has been the most significant component, with both countries spending more than \$3 billion dollars. Finally, illegal behavior in Colombia is pervasive, and understanding how to reduce it 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 used 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. 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 in these models 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 estimates of the effects 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. 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 hired¹. Corman and Mocan (2000) use high frequency changes in the number of police officers arguing this variation is exogenous due to administrative burdens in the hiring process of police officers. They find that more police officers cause a reduction in burglaries, but no effect on other crime categories. Di Tella and Schargrofsky (2004) exploit the exogenous reallocation of police forces across Buenos Aires that came as a result of a terrorist attack in 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 exploiting

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

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 Bogota, Colombia. They find that the intervention significantly reduces crime, not by increasing the police force, but improving its quality and engagement with the community.

Another literature focus on the effects of enforcement, or characteristics related to the likelihood of detection, on soft crime or tax evasion². 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 larger among people whose income is directly reported to the IRS, and therefore, have fewer chances 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 clean identification strategy, but also, on the precision of our data on illicit crop cultivation and enforcement activities (i.e., aerial spraying and manual eradication of crops). In particular, we observe 1-square-kilometer cells on coca cultivation from satellite data, and spraying and manual eradication are recovered from GPS devices that record the exact geographic coordinates where the policies were implemented. Consistent with the previous finding on the literature, our results suggest that enforcement, in the form of a greater likelihood of being aerielly sprayed, reduces illicit coca cultivation by farmers.

Our paper is also related to the branch of applied economic literature on the cost-effectiveness of anti-drug policies. The main challenge in this area is that anti-drug interventions typically take place at large scales, and 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 costs, effectiveness, and efficiency of different anti-drug strategies in producer and transit countries. Their main result is that 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 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.

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

However, these results are likely to be driven by simultaneity bias in their estimates. In particular, spraying is higher in areas with more cultivation, creating an upward bias in OLS estimates. Recent studies have attempted to address these endogeneity concerns. For example, Moya (2005) uses matching techniques employing municipal data from Colombia and finds spraying does not have a significant effect on coca crops. Yet, the comparison between municipalities may still be subject to omitted variable bias even after matching on observables. Reyes (2011) instruments spraying with the distance between sprayed areas and the closest military base, and finds evidence that aerial spraying increases illicit crops. However, the main limitation of this study is that the exogeneity of the location of military bases is hard to justify. Finally, Rozo (2014) instruments spraying with the interaction between the distance between each 1-squared-km cells (or coca producer) and the nearest border of a protected area and the U.S. international anti-drug expenditures. The author exploits the fact that by governmental mandate protected areas cannot be sprayed with herbicides due to environmental and social concerns. Her results indicate that aerial spraying has a negative and significant effect on the hectares of coca cropped and coca producer's productivity, but that, at the same time, it causes 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. We also 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 useful anti-narcotic policy. 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 1996 and the increasing involvement of illegal armed groups in these activities, in September of 1999 the governments of Colombia and the U.S. launched a joint strategy which would come to be known as the *Plan Colombia*. According to official figures, the United States government disbursed close to \$470 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 \$710 million dollars per year during the same period in the fight against illegal drug production and trafficking under *Plan Colombia* (see DNP, 2006). Altogether, between 2000 and 2008 the total expenses on the military component of *Plan Colombia* represented close to \$1.2 billion dollars per year, corresponding to 1.1% of the country's annual GDP. As such, *Plan Colombia* is the largest anti-drug intervention that has ever been made in a producing country.

The strategies implemented under *Plan Colombia* included aerial spraying campaigns, manual eradication campaigns, control of chemical precursors used in the processing of coca leaf into cocaine, the detection and destruction of cocaine processing laboratories, and seizing 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 thousand hectares have been sprayed with herbicides per year, of which almost half are located in Putumayo

and Nariño, the two Colombian departments (states) bordering Ecuador and where our empirical analysis is centered. Figure 1 shows the evolution of hectares with coca cultivation, aerial spraying with herbicides and manual eradication for the whole country and for the departments of Nariño and Putumayo in the last years. About a third of total coca cultivation and half of overall aerial spraying in Colombia between 2000 and 2010 took place in the departments of Putumayo and Nariño.

Spraying campaigns are carried out by American contractors such as DynCorp using small aircraft. Coca crops are sprayed with substances such as Roundup, whose main active ingredient is glyphosate. Glyphosate is absorbed through the plant foliage and is only effective on actively 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, implying that areas with coca crops are much more likely to be sprayed and destroyed by this enforcement strategy.

Hence, farmers that grow coca bushes face the risk of having their crops destroyed by herbicides used in aerial spraying campaigns. 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, some farmers spray molasses on the coca bushes to prevent the herbicide from penetrating the foliage and killing the plant. Also, they can cut the stem of the plant a few hours after the fumigation event and the plant grows back a few months later. Finally, some farmers reallocate their crops to areas less likely to be sprayed. However, these alternatives are costly, which forces some farmers to start cultivating solely legal crops that are not targeted by spraying campaigns. This is the intended effect of aerial spraying that we measure in this paper.

Because aerial spraying campaigns typically target areas with 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 the exogenous variation induced by a natural experiment resulting from a diplomatic friction between the governments of Colombia and Ecuador. The friction concluded in the compromise by the Colombian government not to carry out spraying campaigns in a 10 km strip along the international border with Ecuador starting in 2006.

Since the beginning of fumigation under *Plan Colombia*, the Ecuadorian government complained of alleged adverse effects of spraying on the health of their population, the environment, livestock, and legal crops near the bordering area. In 2006, the Colombian government announced that it would discontinue aerial spraying in a 10 kilometer band around the international frontier with Ecuador within the Colombian territory. However, the Colombian government recanted at the end of 2006 and continued the spraying campaigns in the area. As a result of this noncompliance with the initial agreement, at the end of 2007 the Ecuadorian government filed a lawsuit against Colombia in the International Court of Justice in The Hague. The suit was filed on March 31st, 2008, and ever since then the Colombian government stopped all spraying campaigns in the 10 kilometer strip.

The implementation of this exclusion area generated geographical and time exogenous variation that we exploit to identify the effects of aerial spraying. Figure 2 describes the

geographic location of the exclusion strip.

4 Data

We employ a unique panel data on the location of coca crops in each of 1-square-kilometer (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. The data on coca cultivation comes from satellite images collected annually since 2000. The satellite images show the number of hectares with coca cultivation detected on each grid by the end of that year. We also use cell level data on the number of hectares sprayed for the same period. The data is collected from GPS devices installed in the aircraft used in aerial spraying campaigns, and it records the exact location of the plane when the spraying valves are activated and closed. Using these observations we code a dummy of whether a grid was sprayed or not, for each year from 2000 to 2010. Moreover, we use data on whether manual eradication campaigns took place on each grid, covering the 2007-2010 sub-period. These data is obtained from GPS devices used by manual eradication teams.

We restrict our sample to all grid points with centroids located within 20km of the international frontier with Ecuador. Our sample includes 10,880 cells. We refer to the 5,613 cells within 10km of the frontier as the exclusion region, since this is the area that Colombia committed not to spray. In contrast, we refer to the 5,275 cells located 10km to 20km 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 and coca crops per squared kilometer in both regions from 2000 to 2010. As anticipated above, the figure reveals similar patterns in aerial spraying until 2006, when the Colombian government reduced spraying in the exclusion region. The difference disappears again in 2007, consistent with the fact that the government recanted on its compromise during this year, and a significant gap opens up again since 2008, when the likelihood of spraying is reduced basically to zero in the exclusion region (though some cells in this region were sprayed) and increased in the sprayed area. The data on cultivation reveals a sharp decline from 2000 to 2004, during the first years of *Plan Colombia*, from about 3 hectares per squared kilometer to about 0.5. However, in 2006, and later in 2009 and 2010, cultivation increased in the exclusion region relative to the sprayed area.

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 in 2006, and implemented again from 2008 onwards, to stop aerial spraying 10kms around the international frontier.

In our setting, the forcing variable is the distance from the centroid of each cell i to the international frontier with Ecuador. We normalize the forcing variable to take the value of zero at the 10km cutoff, and denote it by \widehat{D}_i , where $\widehat{D}_i = D_i - 10km$. In this exercise,

we exclude from our sample all cells that had their centroid in the first 500m around the cutoff value, since they have a significant portion of their territory on both the exclusion and the sprayed area. 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. On the contrary, there should be no discontinuity for 2007 or years before 2006.

Let S_{it} be a dummy equal to 1 if grid i was sprayed during year t . Our discussion above implies that in 2006, and from 2008 onwards, all coca plantations with $\widehat{D}_i < 0$ have a probability of treatment near zero, whereas, for those coca plantations located above $\widehat{D}_i > 0$ the probability of treatment jumps to positive values. That is:

$$\lim_{d \downarrow -500} \Pr[S_{it} = 1 | \widehat{D} = d] < \lim_{d \uparrow 500} \Pr[S_{it} = 1 | \widehat{D} = d] \forall t \in \{2006, 2008, 2009, 2010\} \quad (1)$$

The existence of a discontinuity during these years also depends on the precise implementation of the exclusion area. For instance, there may be imperfect compliance by Colombian authorities around the cutoff, or spraying on the sprayed area near the cutoff may be reduced as well. In both cases, cells in the exclusion area will be less likely to be sprayed, but there need not be a discontinuity.

We start by exploring if the policy was implemented in such a way as to create a discontinuity in aerial spraying. We investigate this question by restricting our sample to several bandwidths around the cutoff, including cells with centroids $2.5km$, $2.75km$ or $3km$ away from the cutoff. For each choice of bandwidth, we estimate the model:

$$S_{it} = \pi_{0t} + \pi_{1t}1\{\widehat{D}_i > 0\} + f_t(\widehat{D}_i) + \varepsilon_{it}, \quad (2)$$

for different years, t , or pooling years. Here, f_t is a polynomial in the forcing variable. We compute standard errors clustering at the cell level and robust against heteroskedasticity. The coefficient on π_{1t} measures any discontinuity around the cutoff. Our discussion above implies that we expect $\pi_{1t} = 0$ for $t = 2001, \dots, 2005, 2007$, and $\pi_{1t} > 0$ for $t = 2006, 2008, 2009, 2010$.

The left panel of Table 1 (columns 1 to 3) presents estimates of the difference in enforcement at the discontinuity, π_{it} , for several years and pooled years. The results in each column correspond to different choices of bandwidths specified in the top rows. The estimates in column 1 show that, before 2006 there is no significant discontinuity. In contrast, during 2006, cells in the sprayed area near the 10km cutoff were 7.6% more likely to be sprayed than similar cells in the exclusion region. In 2007 the difference is not significant and it reappears after 2008, consistent with our narrative. The discontinuity in the likelihood of spraying is a robust finding for 2006, 2008 and 2010, as suggested by columns 2 and 3, but less so for 2009, when the Colombian government may have enforced the 10km limit in a continuous way. All the same, when we pool the years 2006, 2008, 2009 and 2010 together – which are the years in which Colombia agree not to spray the exclusion area – we find that the likelihood of spraying was 10.1% higher in the sprayed region (standard error=2.5%), relative to close cells in the exclusion area, and this discontinuity at the cutoff is precisely estimated. Removing the year 2009 does not change our conclusions. The estimates in columns 2 and 3, obtained with different choices of bandwidths, support these findings.

The previous results can be seen graphically. Figure 4 shows the local behavior of the likelihood of aerial spraying on both sides of the 10km cutoff and for each choice of bandwidth (which can be seen in the horizontal axis of each figure). To ease the graphical analysis, we plot cells by the distance of their border to the cutoff, 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 cubic polynomial to approximate the local behavior on each side of the 10km cutoff. The plots reveal a clear discontinuity in the likelihood of spraying during 2008 and 2010, and less so for 2006 and 2009. When we pool the years 2006, 2008, 2009, and 2010 (in which Colombia committed not to spray the exclusion area) we find a clear graphical discontinuity in the likelihood of spraying. In contrast, in the two top panels of Figure 6, which is constructed in a similar way, we see no apparent discontinuity in the likelihood of enforcement for all years before 2005 or in 2007; years during which Colombia sprayed both regions.

We now investigate the consequences of the discontinuity in enforcement 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 the year. We estimate the following specification:

$$Y_{it} = \gamma_{0t} + \gamma_{1t}1\{\widehat{D}_i > 0\} + f_t(\widehat{D}_i) + \epsilon_{it}, \quad (3)$$

for different years, t . Here, f_t is a polynomial. The coefficient on γ_{1t} measures any discontinuity around the cutoff. As for the previous estimates, we compute standard errors clustering at the grid level and robust against heteroskedasticity.

The right panel of Table 1 (columns 4 to 6) presents estimates of the difference in coca cultivation at the discontinuity, γ_{it} , for several years and pooled years. Each column presents estimates obtained with a different choice of bandwidth, indicated in the top row. Consistent with the results on spraying, we find no difference in cultivation from 2000 to 2005 in the first row, or during 2007, when there was regular spraying in the exclusion area. In contrast, for the years with no spraying of the exclusion area (2006, 2008, 2009, and 2010), we find reductions in cultivation in the sprayed area. 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. In 2009 we find negative and imprecise effects, consistent with the weak discontinuity estimated in the left panel. We obtain more precise estimates by pooling the years 2006, 2008, 2009 and 2010. In this case, the estimates in column 4 suggest that cultivation was reduced by about 0.43 hectares per squared kilometer in cells near the cutoff in the sprayed region relative to the exclusion area. This conclusion holds if we remove the data on cultivation for 2009. The estimates in columns 5 and 6 confirm our findings, but are less precise, given that they are obtained for smaller bandwidths with less observations.

Figure 5 and the bottom panels of Figure 6 also present these results graphically. Though it is hard to see the discontinuity in cultivation graphically during the years in which Colombia committed not to spray the exclusion area, the figures show some decline in cultivation in the sprayed region near the discontinuity (Figure 4). In contrast, Figure 6 shows no apparent discontinuity in cultivation during other years in which both areas were sprayed equally.

The above results suggest that the enforcement of the 10km exclusion area created a discontinuity in enforcement around the cutoff during 2006, 2008 and 2010, and less so for 2009. The discontinuity in spraying caused divergent illegal behavior in both sides of the

cutoff. In particular, near the cutoff, the 10 percentage points increase in the likelihood of spraying caused a statistically significant reduction of about 0.42 hectares per squared kilometer.

The fact that we do not find any discontinuity before 2005, or during 2007 in either cultivation or spraying is reassuring: It implies that time invariant characteristics do not vary discontinuously around the cutoff, and the dynamics of cultivation on both sides of the 10km line were balanced before 2005. This suggests that the particular choice of the exclusion area was rather arbitrary and was not done strategically aiming at certain cells with particular cultivation dynamics. Remarkably, when the policy was reversed in 2007, the difference in cultivation between both regions stopped being significant and was reduced. Though still negative, this can be fully accounted for persistence in cultivation.³

To quantify the exact effect of spraying on illegal coca cultivation, we compute 2SLS estimate using the discontinuity as the instrument. That is, we estimate equation:

$$Y_{it} = \beta_{0t} + \beta_{1t}S_{it} + f_t(\widehat{D}_i) + v_{it}, \quad (4)$$

for different years separately, instrumenting S_{it} with the dummy $1\{\widehat{D}_i > 0\}$ (so that equation 2 corresponds to the first stage).

Table 2 presents the fuzzy RD estimates for different choices of bandwidth presented in the different columns, and different degrees of the polynomial f_t in different rows. The estimates in the left panel (columns 1 to 3) pool the years 2006, 2008, and 2010; while the right panel (columns 4 to 6) adds 2009. Our estimates indicate a negative LATE of the likelihood of aerial spraying on total hectares of coca planted per squared km. Our estimates suggest that a 10 percentage points increase in the likelihood of aerial spraying reduces cultivation from 0.35 to 0.65 hectares depending on the different choices of bandwidth, or degree of the polynomial used to approximate the local behavior of spraying and cultivation near the cutoff.

To ensure the consistency of our estimates, we require other covariates determining cultivation to vary smoothly around the 10km cutoff. We have showed this is indeed the case for coca cultivation and spraying before 2006, which is reassuring. Table 3 shows that the same holds for manual eradication since 2007 (we only have data on manual eradication beginning this year). These results imply that the decrease of aerial spraying in the exclusion area was not compensated by increasing manual eradication, and hence our estimates only reflect the causal effect of the policy change in spraying. Moreover, though not reported to save space, we find no discontinuity in terms of altitude around the cutoff.⁴ This allows us to reject the idea that the government adjusted the exclusion area as a function of height so that only those areas where altitude prevented coca production were part of the no spraying area.⁵

³In fact, we estimate an auto-correlation coefficient of 0.3, which implies that the estimate in 2007 can be fully explained by persistence in cultivation.

⁴As a last robustness check on our results we estimate a placebo test using a fake cutoff value of the forcing variable of 15kms around the international border and including all the observations 5 to 25kms away from the international frontier with Ecuador. Given that this is an arbitrary threshold, and no policy change was being implemented around it, we expect no significant discontinuities. Though not reported to save space, we did not find any discontinuity in spraying or cultivation in this case. This results are available upon request.

⁵In fact, coca is more productive at medium altitudes; see Mejia and Restrepo, 2013b.

The evidence in this section suggests that increases in the likelihood of enforcement does result in less illegal behavior. We take the estimates in this section as evidence that an increase of 10 percentage points in the likelihood of spraying, causes a reduction between 0.35 and 0.65 hectares of coca per squared kilometer. In practice, we believe that part of our estimate captures the possibility that coca farmers reallocate their crops to the exclusion region, which seems reasonable given the proximity between cells. However, this simply 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 suggest that cultivation in the exclusion region increases within cells, and not because farmers grow coca bushes in new cells. This suggests little reallocation of farmers from the sprayed region to cells in the exclusion area that did not have coca before. In any case, we cannot entirely rule out the extent of reallocation of coca crops, and our point estimates remain an upper bound on the effects of enforcement.

6 Conditional-differences-in-differences estimates

In the previous section we exploited the geographical discontinuity in enforcement around the 10km cutoff. In this section, we exploit within cell variation to estimate the causal effect of aerial spraying on coca cultivation using a conditional difference-in-differences strategy. Figure 7 plots the difference in spraying and cultivation between the sprayed and exclusion areas for each year. As can be seen, there is a differential increase in the likelihood of spraying in the sprayed region during 2006 and since 2008 relative to the exclusion area and with respect to the years 2000 to 2005. However, the problem of directly exploiting this variation is that there are unbalanced dynamics in cultivation before 2006 that may confound our estimates, as can be seen in the right panel.

Formally, let $T_i = 1\{\widehat{D}_i > 0\}$, and let \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 $t = 2006, 2008, 2009, \text{ and } 2010$ (recall that in 2007 both regions were sprayed). The traditional differences in differences methodology, estimates it by a regression of $Y_{it} - \bar{Y}_i$ on T_i . The problem with the traditional estimate is that, as shown in Figure 7, T_i is correlated with $Y_{it'}$ for $t' \leq 2005$. If there are dynamic linkages in cultivation, this would bias the traditional differences in differences estimate.

Thus, any consistent estimate exploiting within cell variation and comparing both regions must control for the dynamics of cultivation and spraying before the diplomatic friction (that is, the years 2000 to 2005). Our regression discontinuity setup did not face this issue because there was no discontinuity in cultivation and spraying from 2000 to 2005. The difference only appears when comparing both regions as a whole, and not simply cells around the cutoff.

In this section, instead, we posit the following conditional independence assumption:

$$Y_{it} - \bar{Y}_i \perp T_i | Z_i, \{Y_{it'}, S_{it'}\}_{t' \leq 2005}, \quad (5)$$

where Y_{itd} is the potential cultivation in cell i and year t , for a fixed enforcement regime d . Here $d = 1$ means the cell is in a sprayed area, and $d = 0$ means it is not. This 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 covariates Z_i (including a polynomial in altitude and municipality fixed effects), potential cultivation would be equal for cells in the sprayed and exclusion areas (in the absence of a difference in enforcement). We use the change in potential cultivation instead of its level to remove any permanent difference between cells not captured by the conditioning set.⁶ We believe this is a plausible assumption, as coca cultivation in two cells with the exact same path of cultivation, spraying, and manual eradication from 2000 to 2005, should follow very similar trajectories in the absence of a difference in enforcement, even if they are on different sides of the 10km threshold.

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

$$Y_{it} - \bar{Y}_i = \beta_t T_i + \delta_t + \sum_{t' \leq 2006} \Gamma \cdot (Y_{it'}, S_{it'})' + \Theta Z_i' + \varepsilon_{it}, \forall t \geq 2006 \quad (6)$$

Here, β_t identifies the effect of being in the sprayed region (relative to an area without spraying) 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 the top panel of Table 4 presents the regression estimates of being in the sprayed region on cultivation. We present separately the estimates by year, and also pool together the years 2006, 2008, 2009 and 2010 when Colombia committed not to spray the exclusion area. We refer to these years as post-treatment years. Our estimates show that cultivation was reduced in all years after 2006; the effects are all significant at all traditional levels; and the average effect pooling the post-treatment years together is a reduction of 0.233 hectares per squared km (s.e=0.015). The bottom panel presents estimates using the likelihood of spraying as dependent variable. Our estimates show a large increase in the likelihood of spraying in the sprayed region, specially after 2008, and a very small increase in 2007, consistent with our narrative that this year both regions were sprayed though the exclusion region less so. When pooling the post-treatment years together (2006, 2008, 2009 and 2010), we find that the likelihood of aerial spraying was 10 percentage points larger in the sprayed area, which is similar to the estimate obtained using our regression discontinuity design. We do not estimate pre-treatment behavior in these variables to check for equal trends because these are zero once we control for pre-treatment characteristics.

Our estimates reveal a significant decline in cultivation during 2007 in the sprayed area despite the very small difference in the likelihood of spraying during this year. However, this could simply reflect the persistence of cultivation in the exclusion region. In fact, we estimate an autocorrelation coefficient in the series for coca cultivation for each grid of 0.3. This implies that the reduction of 0.197 hectares per squared kilometer in 2007 can be fully

⁶In theory, we do not even have 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 estimated propensity score.

explaining by the persistent effects of the sharp decline in cultivation in 2006.⁷

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.

In column 2 we reweight the regression in equation 6 by the propensity score (see Hirano, Imbens and Ridder, 2003, for more on this approach). 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 $\hat{\lambda}_i$ 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 on the treated (that is, the sprayed area).⁸ Besides reweighting by the propensity score, we also control linearly for all covariates in the regression. This is known as a double-robust regression: on the one hand, reweighting the data controls non-parametrically for the influence of covariates; on the other one, the covariates in the regression control linearly for any source of misspecification in the propensity score. 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.⁹

The role of weighting the data by the propensity score can be seen graphically in Figure 8. We plot estimates of β_t for $t = 2000, \dots, 2010$ 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 area before 2006. On the contrary, the raw data on cultivation presented in Figure 7 exhibited unbalanced dynamics before 2006 that could confound our estimates. A similar pattern emerges for spraying, although dynamics were already roughly balanced in the raw data.¹⁰

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 estimated propensity score between $(j-1) \times 0.05$ and $j \times 0.05$. For each bin, we estimate equation 6 separately, and use a weighted averages of all these estimates to obtain an estimate for β_t . We obtained the variance of β_t as a weighted sum of the variances as well.¹¹

⁷If we subtract $0.3Y_{it-1}$ from the dependent variable to control for dynamics, the effect for 2007 disappears and we still find an effect for all other years, consistent with our analysis above.

⁸We 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 score that make this exercise imprecise. In any case, our results are similar.

⁹We 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 computed an alternative set of bootstrapped standard errors taking into account the estimation of the propensity score and obtained slightly smaller standard errors not reported.

¹⁰When computing these estimates we do not control linearly for the covariates in the regression presented in equation 6. Doing so mechanically sets the estimates of $\beta_t = 0$ for $t \leq 2005$.

¹¹Again, bootstrapping the whole procedure resulted in similar standard errors. Thus we present standard errors that assume the propensity score is known.

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 as well (results for the ATE were similar). 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 as well when estimating equation 6 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 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 scores in the exclusion region, and thus controls non-parametrically for the propensity score. We present estimates of the standard errors assuming that the propensity score is known.¹² 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.

In columns 5 to 8 of Table 4 we conduct another exercise. Instead of focusing on the period since 2006, when Colombia first committed not to spray the exclusion area, we focus on the period since 2008, when it was forced not to do so by international pressure. In this case, we condition on all covariates until 2007, and replicate the estimates from columns 1 to 4. We do not report the estimates for 2006 and 2007 because these are mechanically zero, showing that this methodology weights the data in such a way that, by construction, the sprayed and exclusion areas have equal cultivation and spraying trends until 2007. These estimates also reveal a negative effect on cultivation of being in the sprayed area, and a similar positive effect on the likelihood of spraying of around 10 percentage points. The fact that the estimate on cultivation is less negative could reflect the possibility of weaker responses to enforcement during these years. However, we prefer our first set of estimates in columns 1 to 4, since arguably, including the endogenous response of cultivation in 2006 as a control is not entirely satisfactory.

All the same, the estimates in this section suggest that grids in the sprayed region were 10 percentage points more likely to be sprayed during years in which Colombia committed not to spray the exclusion area. As a consequence, farmers in this region reduced cultivation by 0.2 hectares per squared kilometer. The 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 4 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 10km cutoff. This implies that this indirect margin is more relevant for the regression discontinuity estimates, making them overstate the deterrent effect

¹²In general, it is not known whether bootstrapping them produces consistent estimates, so we keep the naive ones. In any case, bootstrapped standard errors were very similar.

of spraying more. In any case, both methodologies reveal significant responses by households to a differential likelihood of spraying, consistent with the view that enforcement reduces illegal behavior.

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 increase in the likelihood of aerial spraying reduces coca cultivation between 0.35 and 0.65 hectares per squared kilometer. Our conditional difference in differences points to a reduction of approximately 0.2 hectares. Since 1 squared kilometer contains 100 hectares, these estimates imply that to reduce cultivation by 1 hectare during a given year, the Colombian government has to spray 15.4-50 hectares per year (to increase the likelihood of spraying in one squared kilometer by 15.4 percentage points or 50 percentage points).

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 financing spraying campaigns costs the U.S. \$11,550-\$37,500 dollars. Additionally, for every dollar spent by the U.S., Colombia spends about 2.2 dollars (aerial spraying campaigns are jointly financed by both countries), making the overall cost \$37,000-\$120,000. To put these numbers in perspective, that hectare is able to produce enough coca bushes to synthesize 6 kilograms of cocaine per year, which would cost \$12,000 dollars to buy in Colombian farms.

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 follow 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 0.004%.¹³This elasticity is small for several reasons. First, because cultivation represents only a small fraction of the total market value of cocaine in consumer markets. Thus, an increase in the price of coca bushes 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 more chemical precursors and technology to produce more cocaine per hectare, demand more cocaine from other source countries or switch to better transportation techniques, partially offsetting the effect of the shock on the supply of cocaine.

¹³We use the estimate they report of $\Lambda_q = 0.004$, which corresponds to the elasticity of final supply with respect to reductions in cultivated area in Colombia. We take the estimate for an elasticity of demand of 0.75 in Table 5.

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. \$9.24-\$30 million dollars per year. However, this investment would only reduce the supply of cocaine in its territory by 0.004%, which equals 20kg. This implies that the marginal cost to the U.S. of reducing retail quantities of cocaine by 1kg by subsidizing aerial spraying in Colombia is \$462,000-\$1,500,000 dollars. These are large magnitudes, but are similar to the estimates reported by Mejia and Restrepo (2013) using an entirely different methodology. To put them in perspective, the price of 1kg 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 1kg is estimated at \$181,000 dollars 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 only has 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, 2014); other legal crops and the environment (see Relyea, 2005 and Dávalos et al., 2011); and on the socio-economic conditions of coca-producing areas (see Roza (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 partially destroying the illicit goods – 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, it responds to the likelihood of enforcement.

Our main contribution is to present a clean and credible source of identification to identify 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 in a 10km band along the border with Ecuador in 2006.

We use a regression discontinuity design, exploiting the arbitrary 10km line and a conditional-differences-in-difference 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 10 percentage points more likely to be sprayed than cells in the exclusion area. In consequence, coca cultivation decreased in this region between 0.35 to 0.65 hectares (regression discontinuity estimates) or 0.2 hectares (conditional differences in differences estimate).

Despite reducing coca cultivation aerial spraying in Colombia only has small effects in downstream markets. We estimate that reducing the Colombian coca cultivation between 2000 and 2010 by 1% (800 hectares), would cost the U.S. between \$9.24-\$30 million dollars per year. However, this investment would only reduce the supply of cocaine in its territory by 0.004%, which equals 20kg. Hence, the cost of reducing cocaine retail supply by 1kg is at least \$462,000 dollars per year if resources are used to subsidize spraying campaigns in Colombia. Other policies such as treatment and prevention, or even subsidizing interdiction efforts abroad, would be significantly more cost effective in curbing drug supply.

9 References

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

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

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

Abadie, A., and Imbens, G. (2011) “Matching on the Estimated Propensity Score,” NBER.

Andreoni, James, Brian Erard, and Jonathan Feinstein (1998) “Tax Compliance,” *Journal of Economic Literature*, 36(2), 818-860.

Angrist, Joshua D. (1998) “Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants,” *Econometrica* Vol. 66, No. 2 , pp. 249-288

Bar-Ilan, Avner and Sacerdote, Bruce (2001) “The Response to Fines and Probability of Detection in a Series of Experiments,” National Bureau of Economic Research (Cambridge, MA), Working Paper No. 8638.

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

Benabou, Roland, and Jean Tirole (2003) “Intrinsic and extrinsic motivation,” *The Review of Economic Studies* 70.3: 489-520.

Benabou, Roland, and Jean Tirole (2006) “Incentives and Prosocial Behavior,” *The American Economic Review* 96.5: 1652-1678.

Beron, Kurt J., Helen V. Tauchen, and Anne D. Witte (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.

Buonanno, P. and Mastrobuoni, G. (2012) “Police and crime: Evidence from dictated delays in centralized police hiring”, IZA working paper # 6477.

Caliendo, M., and Kopeinig, S. (2005) “Some Practical Guidance for the Implementation of Propensity Score Matching,” IZA Discussion Paper No. 1588.

Camacho, Adriana and Mejia, Daniel. (2014) “The Health Consequences of Aerial Spraying of Illicit Crops: The Case of Colombia,” Mimeo, Universidad de los Andes.

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

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

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

Dehejia, Rajeev H., and Sadek Wahba (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), pp. 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, pp. 399-421.

Draca, Mirko, Stephen Machin, and Robert Witt (2011) “Panic on the Streets of London: Police, Crime and the July 2005 Terror Attacks,” *American Economic Review* 101(5): 2157-81.

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

Eck, John and Maguire, Edward (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), pp. 181-201.

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

Garcia, Juan Felipe, Daniel Mejia, and Daniel Ortega (2013) “Police Reform, Training and Crime: Experimental evidence from Colombia’s Plan Cuadrantes,” *Documento CEDE*.

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

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

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

Imbens and Wooldridge (2007a) “What’s new in Econometrics: Estimation of Average Treatment Effects Under Uncounfoundedness,” NBER Summer 2007.

Imbens and Wooldridge (2007b) “What’s New in Econometrics,” *Lecture Notes Summer 2007 NBER*.

- Kagan, Robert A. (1989)** “On the Visibility of Income Tax Law Violations,” in *Taxpayer Compliance: Social science perspectives* 107, Jeffrey A. Roth & John T. Scholz eds.
- Klepper, Steven, and Daniel Nagin (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 ,” NBER. Working Paper 14723.
- Ludwig, J., and Miller, D. (2005)** “Does Head Start Improve Children’s Life Chances? Evidence from a Regression Discontinuity Design,” Working Paper 11702, National Bureau of Economic Research..
- Levitt, Steven (1997)** “Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime,” *American Economic Review*, 87(3), pp. 270-90.
- Levitt, Steven (2002)** “Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime: Reply,” *American Economic Review*, 92(4), pp. 1244-50.
- Marvell, Thomas and Moody, Carlisle (1996)** “Specification Problems, Police Levels, and Crime Rates,” *Criminology*, 34(4), pp. 609-46.
- McCrary, Justin (2002)** “Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime: Comment,” *American Economic Review*, 92(4), pp. 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, Daniel and Restrepo, Pascual (2013)** “The Economics of the War on Illegal Drug Production and Trafficking,” Documento CEDE No. 54.
- Mejia, Daniel and Restrepo, Pascual (2013b)** “Bushes and Bullets: Illegal Cocaine Markets and Violence in Colombia,” Documento CEDE No. 53.
- Menniger, Karl (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*. Vol. 31, No. 2, pp. 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*: 618-627
- Reyes, L. (2011)** “Estimating the Causal Effect of Forced Eradication on Coca Cultivation in Colombian Municipalities,” Department of Economics. Job Market paper. Michigan State University.
- Rosenbaum, P., and D. Rubin (1983)** “The Central Role of the Propensity Score in Observational Studies for Causal Effects,” *Biometrika*, 70, 41-50.
- Rozo, S. (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* Vol. 25, N1, 1-21.




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

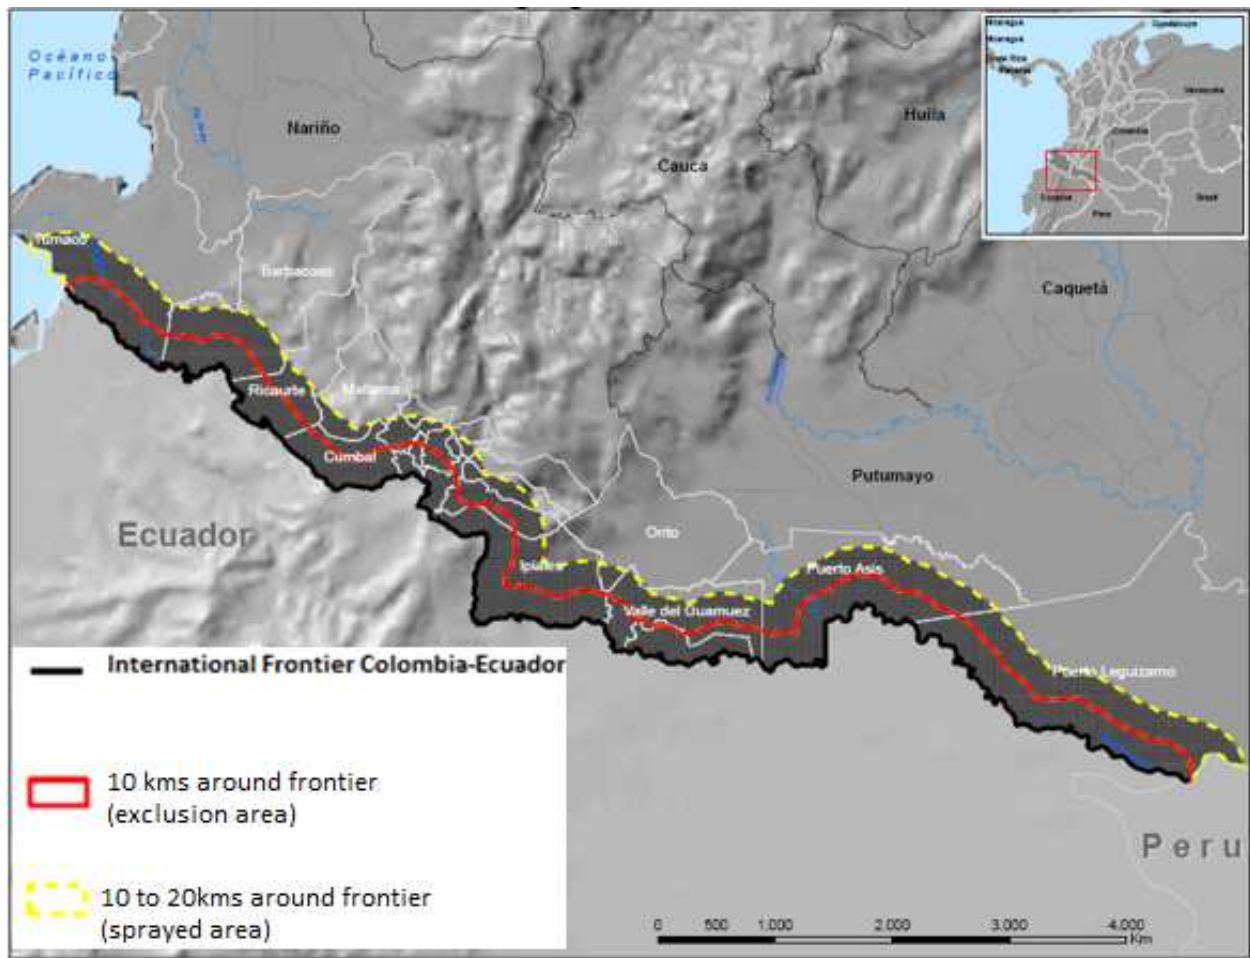


Figure 2: Map of the frontier between Colombia and Ecuador illustrating the sprayed and exclusion areas.

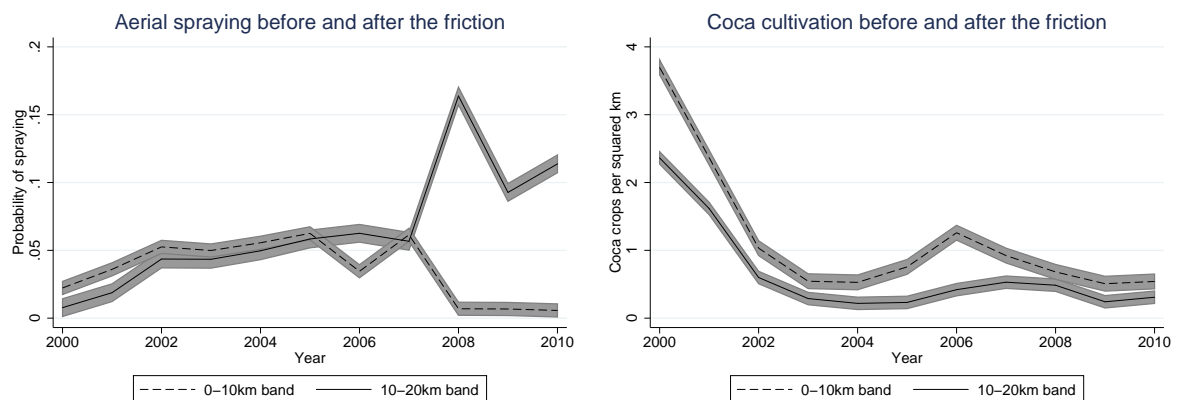


Figure 3: Coca cultivation and likelihood of aerial spraying in the exclusion (dotted line) and sprayed area (solid line) from 2000 to 2010. 95% confidence intervals for the averages in each group are presented as the gray area for each year.

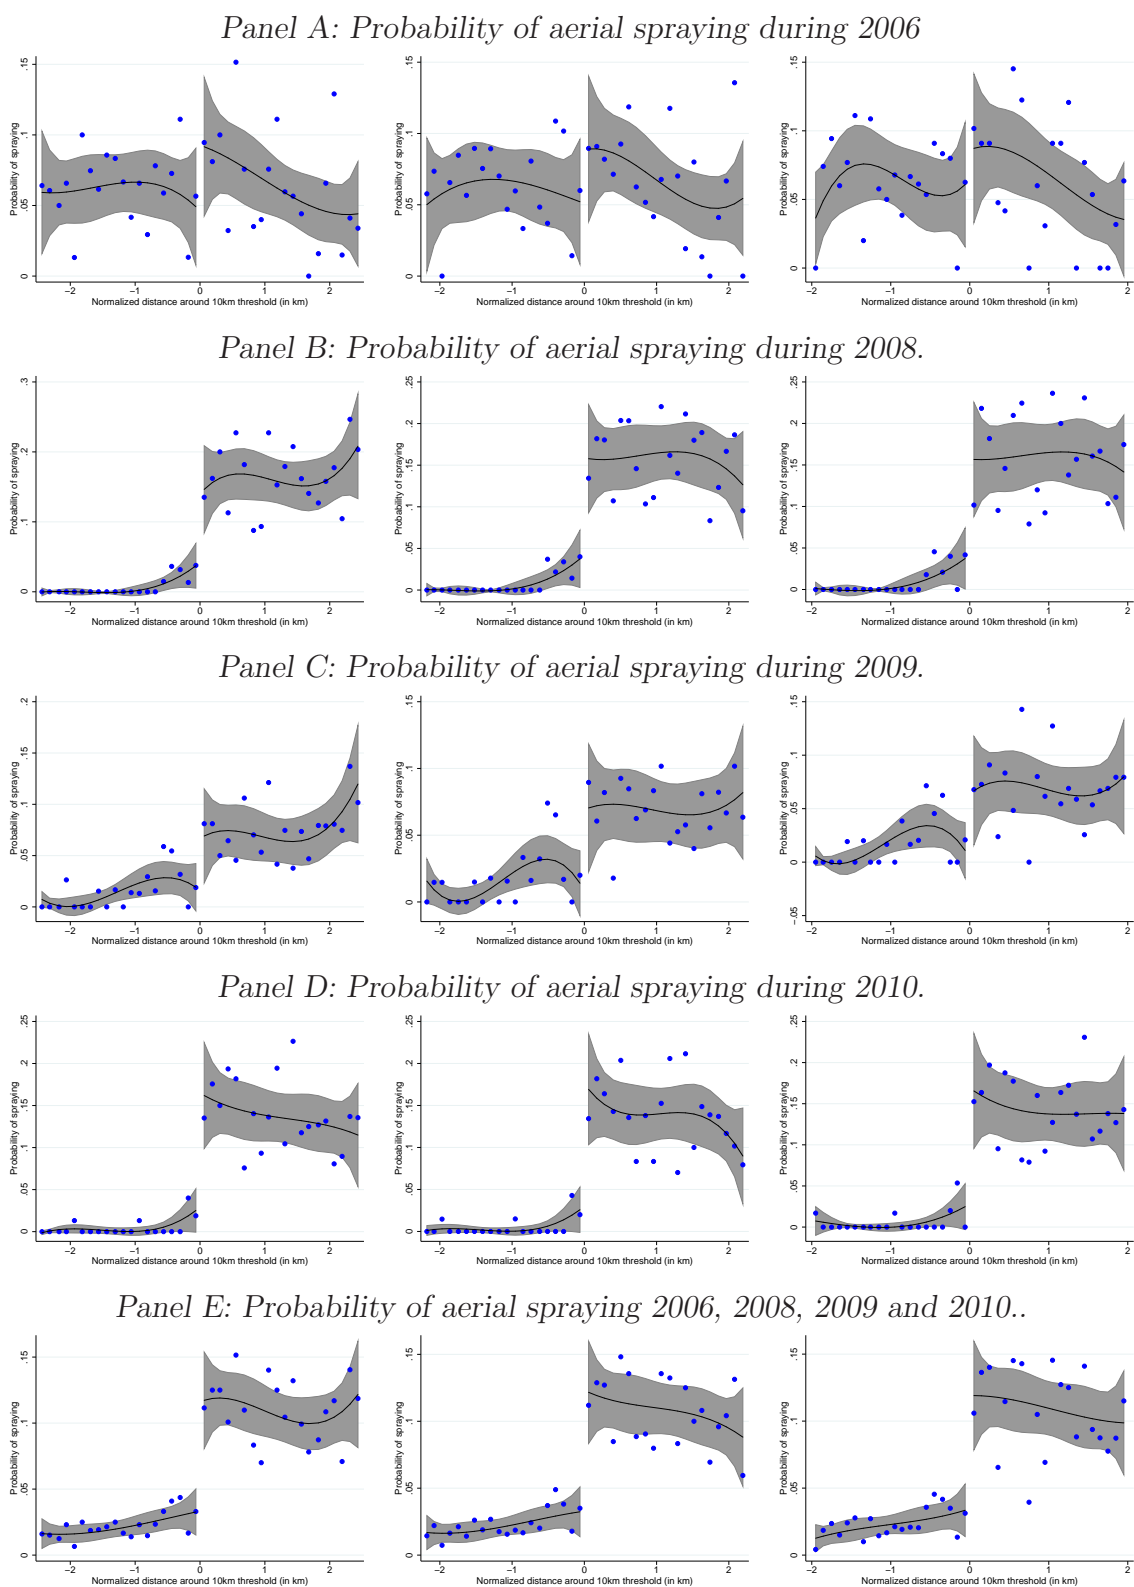




Figure 4: Probability of spraying around the 10km cutoff during years in which Colombia committed not to spray the excluded region.


Panel A: coca cultivation during 2006.




Panel B: coca cultivation during 2008.



Panel C: coca cultivation during 2009.



Panel D: coca cultivation during 2010.



Panel E: coca cultivation during 2006, 2008, 2009 and 2010.




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

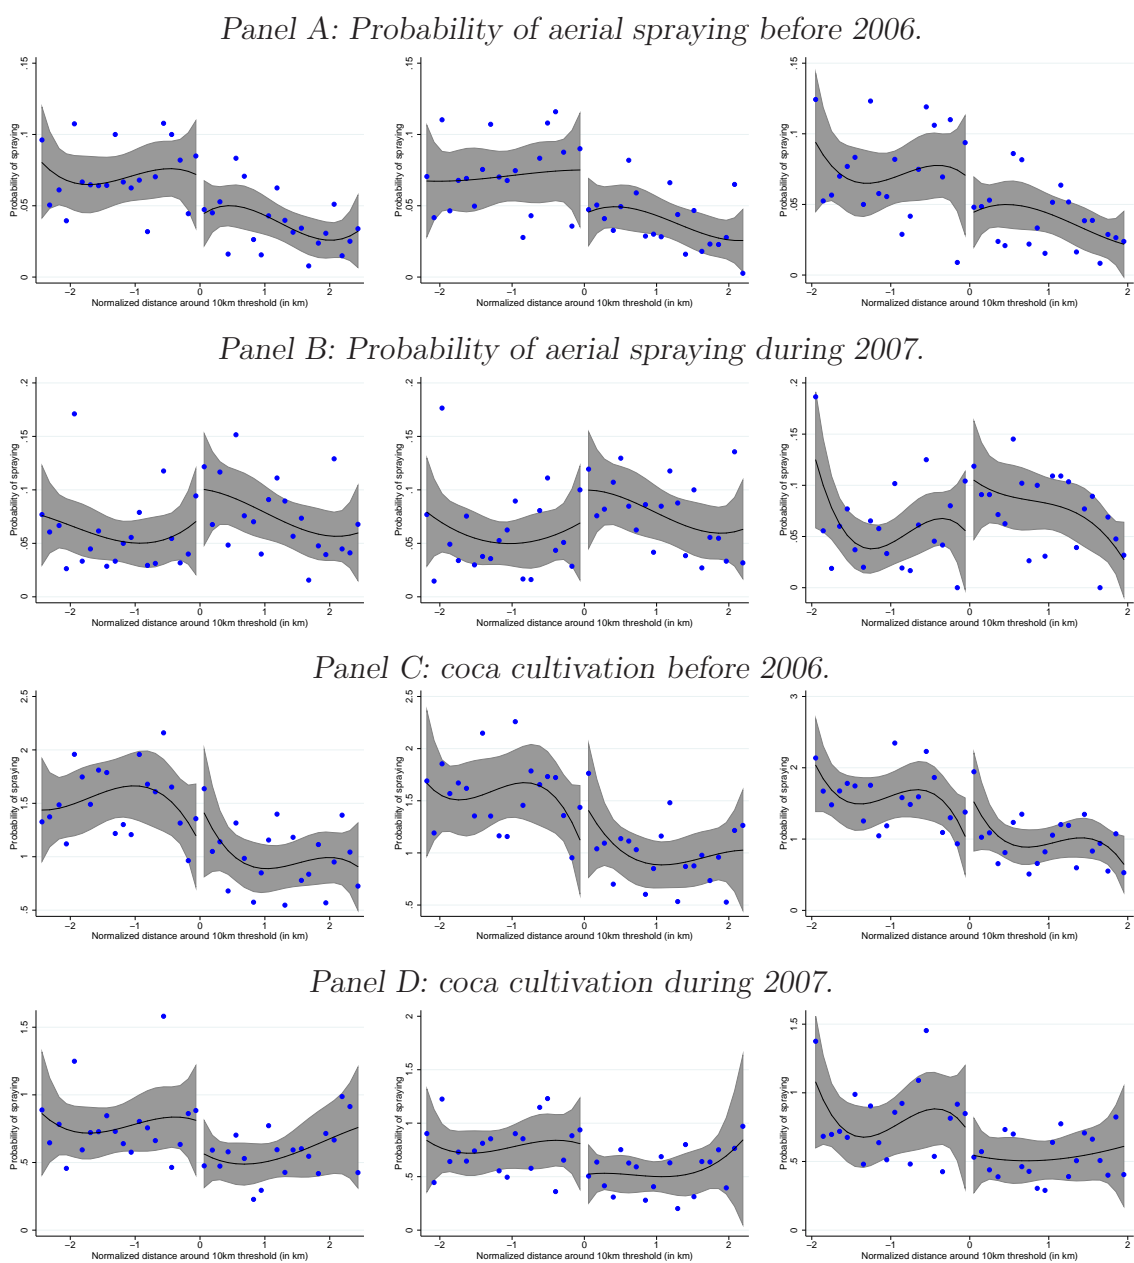


Figure 6: Probability of aerial spraying and coca cultivation around the 10km cutoff during years in which Colombia formally sprayed the exclusion area.




Figure 7: Difference in coca cultivation and spraying between the sprayed and exclusion area from 2000 to 2010.




Figure 8: Re-weighted difference in coca cultivation and spraying between the sprayed and exclusion area from 2000 to 2010 relative to the average before 2006. We weight observations in the exclusion area by the estimated odds ratio based on cultivation and spraying from 2000 to 2005, so that the distribution of these covariates in the exclusion area matches that of the sprayed area.

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

	<i>Difference in spraying</i>			<i>Difference in cultivation</i>		
	(1) 3km	(2) 2.75km	(3) 2.5km	(4) 3km	(5) 2.75km	(6) 2.5km
<i>Bandwidth:</i>						
Difference before 2006:	-0.013 (0.026)	-0.006 (0.029)	-0.025 (0.032)	0.406 (0.458)	0.495 (0.538)	0.297 (0.593)
Observations	15918	14262	12672	15918	14262	12672
Difference in 2006:	0.076* (0.043)	0.092* (0.049)	0.090* (0.054)	-0.729 (0.514)	-0.847 (0.589)	-1.304** (0.658)
Observations	2653	2377	2112	2653	2377	2112
Difference in 2007:	0.063 (0.044)	0.061 (0.048)	0.007 (0.056)	-0.423 (0.318)	-0.340 (0.386)	-0.434 (0.391)
Observations	2653	2377	2112	2653	2377	2112
Difference in 2008:	0.126*** (0.047)	0.088* (0.051)	0.091 (0.058)	-0.407* (0.218)	-0.628*** (0.242)	-0.602** (0.255)
Observations	2653	2377	2112	2653	2377	2112
Difference in 2009:	0.064* (0.035)	0.058 (0.038)	0.066 (0.042)	-0.172 (0.168)	-0.255 (0.186)	-0.230 (0.202)
Observations	2653	2377	2112	2653	2377	2112
Difference in 2010:	0.139*** (0.046)	0.121** (0.051)	0.147** (0.058)	-0.420* (0.235)	-0.419 (0.269)	-0.412 (0.293)
Observations	2653	2377	2112	2653	2377	2112
Difference in 2006, 2008, 2010:	0.114*** (0.030)	0.100*** (0.033)	0.109*** (0.037)	-0.519** (0.254)	-0.631** (0.290)	-0.773** (0.317)
Observations	7959	7131	6336	7959	7131	6336
Difference in 2006, 2008, 2009 2010:	0.101*** (0.025)	0.090*** (0.028)	0.098*** (0.032)	-0.432** (0.220)	-0.537** (0.250)	-0.637** (0.271)
Observations	10612	9508	8448	10612	9508	8448

Notes: The table present regression discontinuity estimates of the difference in spraying (left panel) and cultivation (right panel) around the 10km cutoff. Each row has a different model estimated for different years or pooled years. The bandwidth of cells included in the sample varies from 3km (columns 1 and 4), 2.75km (columns 2 and 5) and 2.5km (columns 3 and 6). 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%.

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

	Years 2006, 2008, 2010			Years 2006, 2008, 2009, 2010		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Bandwidth:</i>	3km	2.75km	2.5km	3km	2.75km	2.5km
<i>Using quadratic polynomial</i>						
Effect of spraying:	-4.011***	-3.582**	-3.447**	-4.197**	-3.608**	-3.536*
	(1.500)	(1.523)	(1.710)	(1.667)	(1.642)	(1.834)
Observations	7959	7131	6336	10612	9508	8448
Instrument F-stat.	42.1	40.2	31.7	37.5	37.7	30.2
<i>Using cubic polynomial</i>						
Effect of spraying:	-4.553*	-6.170*	-6.934*	-4.260*	-5.839*	-6.349*
	(2.587)	(3.594)	(3.809)	(2.509)	(3.453)	(3.596)
Observations	7959	7131	6336	10612	9508	8448
Instrument F-stat.	14.8	9.1	8.5	16.4	10.2	9.5
<i>Using quartic polynomial</i>						
Effect of spraying:	-4.560*	-6.292*	-7.076*	-4.270*	-5.985*	-6.479*
	(2.594)	(3.644)	(3.848)	(2.518)	(3.523)	(3.637)
Observations	7959	7131	6336	10612	9508	8448
Instrument F-stat.	14.9	9.4	8.7	16.4	10.4	9.7

Notes: The table present fuzzy regression discontinuity estimates of the effect of differential aerial spraying on cultivation around the 10km cutoff. Columns 1 to 3 pool the years 2006, 2008 and 2010; while columns 4 to 6 add 2009. Each panel presents estimates controlling for a different polynomial in the forcing variable. The bandwidth of cells included in the sample varies from 3km (columns 1 and 4), 2.75km (columns 2 and 5) and 2.5km (columns 3 and 6). 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%.

Table 3: Estimates of the local difference in manual eradication around the 10km cutoff (sprayed minus exclusion area).

<i>Bandwidth:</i>	(1)	(2)	(3)
	3km	2.75km	2.5km
Difference in 2007:	0.018	0.005	-0.001
	(0.036)	(0.040)	(0.045)
Observations	2653	2377	2112
Difference in 2008:	0.065	0.047	0.017
	(0.058)	(0.064)	(0.072)
Observations	2653	2377	2112
Difference in 2009:	-0.044	-0.064	-0.061
	(0.048)	(0.054)	(0.060)
Observations	2653	2377	2112
Difference in 2010:	0.015	0.014	0.010
	(0.017)	(0.017)	(0.019)
Observations	2653	2377	2112
Difference in 2006, 2008, 2010:	0.040	0.030	0.014
	(0.030)	(0.033)	(0.038)
Observations	5306	4754	4224
Difference in 2006, 2008, 2009, 2010:	0.012	-0.001	-0.011
	(0.028)	(0.031)	(0.034)
Observations	7959	7131	6336

Notes: The table present regression discontinuity estimates of the difference in manual eradication around the 10km cutoff. Each row has a different model estimated for different years or pooled years. The bandwidth of cells included in the sample varies from 3km (column 1), 2.75km (column 2) and 2.5km (column 3). 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%.

Table 4: Conditional differences in differences estimate of being in sprayed region.

Covariates:	Before 2005				Before 2007			
	Propensity score methods				Propensity score methods			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Dependent variable: coca cultivation.</i>								
Estimate for 2006:	-0.527*** (0.044)	-0.491*** (0.044)	-0.459*** (0.048)	-0.483*** (0.043)
Estimate for 2007:	-0.197*** (0.042)	-0.206*** (0.046)	-0.126** (0.056)	-0.193*** (0.045)
Estimate for 2008:	-0.073*** (0.024)	-0.041* (0.023)	0.044* (0.026)	-0.026 (0.023)	-0.063*** (0.024)	-0.017 (0.024)	0.012 (0.025)	-0.008 (0.024)
Estimate for 2009:	-0.196*** (0.017)	-0.178*** (0.018)	-0.110*** (0.019)	-0.165*** (0.017)	-0.142*** (0.017)	-0.146*** (0.022)	-0.105*** (0.020)	-0.129*** (0.018)
Estimate for 2010:	-0.137*** (0.021)	-0.101*** (0.020)	-0.068*** (0.022)	-0.089*** (0.020)	-0.058*** (0.020)	-0.026 (0.019)	-0.007 (0.019)	-0.014 (0.018)
Pooling post-treatment years:	-0.233*** (0.015)	-0.203*** (0.018)	-0.147*** (0.019)	-0.191*** (0.018)	-0.087*** (0.012)	-0.063*** (0.017)	-0.034** (0.016)	-0.051*** (0.016)
<i>Dependent variable: aerial spraying.</i>								
Estimate for 2006:	0.049*** (0.004)	0.045*** (0.004)	0.044*** (0.004)	0.047*** (0.004)
Estimate for 2007:	0.009** (0.004)	0.008** (0.004)	0.014** (0.006)	0.008** (0.004)
Estimate for 2008:	0.148*** (0.005)	0.163*** (0.005)	0.182*** (0.008)	0.165*** (0.005)	0.131*** (0.005)	0.150*** (0.006)	0.155*** (0.007)	0.151*** (0.006)
Estimate for 2009:	0.078*** (0.004)	0.089*** (0.004)	0.091*** (0.006)	0.089*** (0.005)	0.074*** (0.004)	0.086*** (0.005)	0.092*** (0.006)	0.086*** (0.005)
Estimate for 2010:	0.112*** (0.005)	0.111*** (0.005)	0.126*** (0.007)	0.113*** (0.005)	0.105*** (0.005)	0.109*** (0.005)	0.116*** (0.007)	0.109*** (0.005)
Pooling post-treatment years:	0.097*** (0.002)	0.102*** (0.003)	0.116*** (0.004)	0.104*** (0.003)	0.103*** (0.003)	0.115*** (0.004)	0.121*** (0.005)	0.115*** (0.004)

Notes: The table present conditional differences in differences estimates of the effect of being in the sprayed region (relative to the exclusion area) on cultivation and spraying. Each row has a different model estimated for different years or pooled years. Columns 1 and 5 are usual linear regressions. In columns 2 and 6 we estimate the ATT by reweighting the regression using the estimated propensity score. In columns 3 and 7 we estimate the ATT by stratifying on the estimated propensity score. Finally, in columns 4 and 8 we match observations on the propensity score. Columns 1 to 4 condition on cultivation and spraying from 2000 to 2005. Columns 5 to 7 condition on cultivation and spraying from 2000 to 2007, and hence no estimate is reported for 2006 and 2007. Each year has 10,880 observations. Standard errors robust against heteroskedasticity and serial correlation within cells are reported in parenthesis (however, the errors ignore the estimation of the propensity score in a previous stage). Estimates with *** are significant at the 1%, those with ** are significant at the 5% and those with * are significant at the 10%.