Did sugarcane expansion lead to GDP growth in São Paulo state? An analysis at municipal level using propensity score-based estimators[*]

Annelies Deuss[†]
Cornell University

November, 2009

Abstract

Previous studies suggest that sugarcane expansion in Brazil is linked with economic growth. This study tests these findings in sugarcane-expanding municipalities in São Paulo state. In particular, I examine whether municipalities that have increased their sugarcane production have as a result experienced higher GDP per capita growth. Whereas previous studies only demonstrated associations between sugarcane expansion and economic growth, this study establishes a causal link. Using estimators based on the propensity score, I compare GDP per capita growth in sugarcane-expanding municipalities with suitable controls. The results indicate that there is no statistically significant impact of sugarcane expansion on GDP per capita growth in municipalities in São Paulo state.

JEL classification: C14, C21, O13, R13

Key words: Brazil, São Paulo, sugarcane expansion, economic growth, propensity score, ethanol

---

[*] This research was supported by grants from the Cornell University Department of Applied Economics and Management, the Mario Einaudi Center at Cornell University, the Cornell University Graduate School and the Latin American Studies Program at Cornell University. I would like to acknowledge valuable comments and suggestions from Steven Kyle, David Just, Antonio Bento, Heloisa Burnquist, Alceu Veiga Filho, Antonio Carlos Simões Florido, Carlos Edoardo Fredo, Antonio de Padua Rodrigues, Eduardo Leão de Sousa, Gerd Sparovek and Arnaldo César Walter. Any remaining errors and omissions are my own.
[†] Department of Applied Economics and Management, Cornell University, Ithaca, NY, 14853, USA. Email: ad328@cornell.edu

Contents

1 Introduction 3

2 The impact evaluation problem 5

2.1 Theoretical aspects 5

2.2 The propensity score 6

2.3 Estimations based on the propensity score 7

2.3.1 Blocking on the propensity score 8

2.3.2 Reweighting based on the propensity score 8

3 Description of variables and data 9

3.1 Period of analysis 9

3.2 Definition of treatment 10

3.3 Selection of control variables 11

3.4 Outcome variable 12

4 Causal effect of sugarcane expansion on GDP per capita growth 12

4.1 Estimation of the propensity score 12

4.2 Analysis of ATT estimates 14

4.2.1 ATT estimates from blocking on the propensity score 14

4.2.2 ATT estimates from reweighting on the propensity score 15

4.2.3 Comparison of ATT estimates obtained with different techniques 17

5 Conclusion 18

6 References 20

7 Appendix 22

1 Introduction

Brazil has experienced a sharp increase in sugarcane production since 2000. The main drivers behind this increase were the rising demand for sugar and ethanol on the domestic and international market. As a result, many municipalities in Brazil have converted their agricultural land to sugarcane plantations. This was especially the case in the municipalities in the state of São Paulo, which were responsible for 63 percent of the 1.53 million hectare national increase in sugarcane harvested area over the period 2000-2006.

The impact of the booming sugarcane sector on the economies of these municipalities is, however, not straightforward. On the one hand, sugarcane expansion leads to more employment and income opportunities in the sugar and ethanol sector (Macedo, 2005). On the other hand however, there is a growing concern that as sugarcane replaces other crops it monopolizes agricultural and economic activities (Ramos, 2008). The aim of this research is to analyze whether sugarcane expansion in São Paulo had a positive impact on the economies of these sugarcane-expanding municipalities.

The existing literature suggests that there is a positive link between sugarcane production and economic growth. Walter (2008) compares two groups of cities in São Paulo state: cities that together produced more than 90% of the total sugarcane output and the remaining cities. He shows that the former group of cities had a statistically significant higher monthly per capita income in 2000 than the latter group. Sparovek et al. (2008) analyze economic changes associated with sugarcane expansion in all municipalities in Brazil over the period 1996-2006. The authors categorize Brazilian municipalities into two groups: municipalities that expanded sugarcane production and municipalities that didn't expand sugarcane production. They show that GDP and GDP growth are higher in the group of municipalities that are classified as sugarcane-expanding compared to the group of municipalities without sugarcane expansion.

The set-up of this research is similar to the abovementioned studies. In particular, average annual GDP per capita growth from 2002 to 2006 is compared between two groups of municipalities in São Paulo state: municipalities that expanded sugarcane production between 2000 and 2006 and municipalities that didn't expand sugarcane production during this period.

This research, however, differs significantly from the previous studies in its methodological design. A main weakness of these studies is that they fail to establish a strong causal link between sugarcane expansion and GDP per capita growth. In order to evaluate this causal effect, it is necessary to assess what the situation would have been if no sugarcane expansion had taken place, i.e. the counterfactual situation. Both Walter (2008) and Sparovek et al. (2008) compare sugarcane-expanding with sugarcane-non expanding cities or municipalities, but only control for one variable in their counterfactual design. In particular, Walter (2008) controls for the cities' population while Sparovek et al. (2008) control for regional effects by exclusively comparing neighboring municipalities. There is, however, a broad set of other factors that might have influenced GDP per capita growth in these municipalities. Ignoring these key variables will hamper a solid construction of the counterfactual scenario.

This study controls for the effects of other covariates on GDP per capita growth by constructing counterfactual scenarios based on the propensity score. This technique was first developed by Rosenbaum and Rubin (1983) and has been widely used to estimate causal effects. Using estimators based on the propensity score, this research then establishes whether sugarcane expansion in São Paulo state did indeed cause economic growth.

A sound analysis of the potential impacts of sugarcane expansion on economic growth is especially crucial at this point. Brazil plans to expand its area devoted to sugarcane production from 6.2 million hectares in 2006 to 13.9 million hectares by 2020 (Jank, 2007). A better understanding of how sugarcane expansion has influenced local economies will give insights into the potential benefits of these projected increases.

This paper proceeds as follows. The next section presents the theoretical basis of the empirical analysis. It introduces the impact evaluation problem and discusses the estimation of average treatment effects on the treated (ATT) based on the propensity-score. Section 3 describes the variables and datasets used in this study. It explains how municipalities are classified into the treatment or control group and describes the selected control variables and outcome variable. Section 4 analyzes the causal effect of sugarcane production on GDP per capita growth in São Paulo state. First, the propensity score is estimated and consequently two types of propensity-score based estimators are constructed. Then, ATT is estimated with each estimator and the results are compared. Section 5 concludes.

2 The impact evaluation problem

I am interested in analyzing whether municipalities that have increased their sugarcane production have as a result experienced growth in GDP per capita. Unlike previous studies, which only show associations between increased sugarcane production and municipal GDP, I want to establish whether there is a *causal link* between sugarcane expansion and municipal GDP per capita growth.

In order to make such causal claims, I need to take into account two issues. First, I need to establish the direction of causality. That is, I need to ensure that sugarcane expansion influenced GDP per capita growth and not vice versa. I will address this potential endogeneity problem by considering lagged values of the control variables in the analysis. Second, I need to know what the situation would have been in these municipalities if no sugarcane expansion had taken place, i.e. the counterfactual situation. An obvious way of assessing the counterfactual situation is by comparing GDP per capita growth between two groups of municipalities: municipalities where sugarcane expansion has increased between 2002 and 2006, and municipalities where sugarcane production hasn't increased over that same period. What is crucial in this design is that I have to be sure that a difference in GDP per capita growth is due to sugarcane expansion and is not a result of prior differences between the two groups.

The problem that arises here is that randomization into the two groups is not possible since I am dealing with non-experimental data. As a result, I am confronted with the problem of "missing data" (Blundell and Costa Dias, 2000) and hence have to estimate the direct effect of sugarcane expansion from the variation in the outcome variable across the municipalities. There are several parametric and non-parametric models available to estimate these causal effects. In this research I will refrain from using a parametric approach with simple regression estimators because the estimators in these models can be very sensitive to differences in the covariate distributions between the two groups. Indeed, simple regression estimators rely heavily on extrapolation. If control units don't look similar to treated units, then the causal effect estimates become very sensitive to minor modifications in the statistical model (King and Zeng, 2006). I will overcome this problem by using estimators based on the propensity score, a non-parametric approach first proposed by Rosenbaum and Rubin (1983). In particular, I will apply two different techniques: blocking on the propensity score and reweighting based on the propensity score.

2.1 Theoretical aspects

My goal is to estimate the effect of sugarcane expansion on GDP per capita growth in the municipalities in the state of São Paulo. The unit of analysis i is hence the municipality, and the outcome variable is municipal GDP per capita growth. Following the notation in the evaluation literature, let the treatment status be represented by a dummy variable D, taking value 1 if the municipality expanded its sugarcane production and value 0 otherwise[1]. I then define the outcome (GDP per capita growth) for sugarcane-expanding (treated) municipalities, i.e. the municipalities for which D_i=1, as Y_{1i} and the outcome for sugarcane-non-expanding (non-treated or control) municipalities, i.e. the municipalities for which D_i=0, as Y_{0i}.

[1] I describe the classification into treatment or control group in detail in section 3.2.

$$Y_i \equiv \begin{cases} Y_{0i} & if\ D_i = 0 \\ Y_{1i} & if\ D_i = 1 \end{cases} \qquad [1]$$

The causal effect of treatment (sugarcane expansion) in a certain municipality is given by the difference in the potential outcomes with and without treatment, $Y_{1i} - Y_{0i}$. Since a municipality will either expand sugarcane production (D_i=1) or not expand sugarcane production (D_i=0), one of these potential outcomes is always a counterfactual and thus never observed. This is known as the "fundamental problem of causal inference" (Holland, 1986) and implies that I cannot compute the individual treatment effect. I can however estimate the average effect, which compares the average outcomes of the treated and non-treated groups.

In this study, I am interested in estimating the average treatment effect on the treated (ATT). In other words, I want to evaluate what would have happened to GDP per capita growth of the sugarcane-expanding municipalities if they hadn't expanded sugarcane. The ATT is defined as the difference between expected outcome values with and without treatment for those municipalities that actually participated in the treatment (Heckman, 1997):

$$ATT \equiv E(Y_1 - Y_0 | D = 1) = E(Y_1 | D = 1) - E(Y_0 | D = 1) \qquad [2]$$

The non-experimental design of this research implies that I cannot directly identify the counterfactual outcome but that I will have to estimate it. In particular, I will need to estimate $E(Y_0|D=1)$. In this research I will use non-parametric techniques based on the propensity score to estimate the counterfactuals. I prefer this approach to parametric approaches such as ordinary least square regressions (OLS), instrumental variables estimation procedures (IV), and Heckman's two-step model (Heckman, 1979) because these parametric techniques are heavily dependent upon the specification of the functional form (King and Zeng, 2006).

2.2 The propensity score

When I estimate the counterfactual outcomes, I want to make sure that I am comparing municipalities in the control group that are as similar as possible to those in the treatment group. In other words, I need to control for any other variables that might have affected treatment so that the difference between the treatment and control group is due to the treatment status alone and isn't influenced by any other differences between the treatment and control group. I hence rewrite equation [2] as:

$$ATT \equiv E(Y_1 - Y_0 | X, D = 1) = E(Y_1 | X, D = 1) - E(Y_0 | X, D = 1) \qquad [3]$$

where X is a vector of characteristics that predict treatment.

Conditioning on a set of covariates becomes difficult to implement when the set of covariates is large – a problem known as the 'curse of dimensionality'. Rosenbaum and Rubin (1983) overcame this problem by summarizing all the variables in X into an index function, the propensity score $p(X)$. This balancing score is defined as the conditional probability of being

treated, given the observed covariates X.

$$p(X) = Pr(D=1 \mid X) \qquad [4]$$

Rosenbaum and Rubin demonstrate that if potential outcomes are independent of treatment conditional on covariates X, they are also independent of treatment conditional on the propensity score *p(X)*:

$$Y_0, Y_1 \perp D \mid X \rightarrow Y_0, Y_1 \perp D \mid p(X) \qquad [5]$$

This result has important practical implications because it is much easier to condition on just one number (the probability of being treated, or propensity score) than on a vector of X characteristics.

When I use the propensity score to estimate ATT, I need to invoke two identifying assumptions[2]. The first assumption is known as the (weak) 'unconfoundedness' (Rosenbaum and Rubin, 1983), 'selection on observables' (Heckman and Robb, 1985) or the 'conditional independence assumption (CIA)' (Lechner, 1999). This assumption states that once I control for observable characteristics, the systematic differences in outcomes between treated and comparison municipalities are entirely attributable to treatment. In other words, the treatment is assumed to satisfy some form of exogeneity or

$$Y_0 \perp D \mid X \qquad [\text{Assumption 1}]$$

where X is the vector of observable variables that are unaffected by the treatment. Note that the CIA assumes that all relevant differences between the two groups are captured by their observables X and rules out any potential impact of unobserved explanatory characteristics.

The second assumption is related to the joint distribution of treatments and covariates. This condition is known as the (weak) 'common support condition' or 'overlap condition' and prevents a situation of perfect predictability of *D* given X.

$$P(D=1 \mid X) < 1 \qquad [\text{Assumption 2}]$$

The common support condition hence ensures that, for each treated municipality, there are control municipalities with the same X values (Heckman et al., 1999). As a result, the outcomes obtained by those municipalities from both groups that belong to this common support will be comparable.

2.3 Estimations based on the propensity score

When I use propensity-based estimators to estimate ATT, I first have to estimate the propensity score *p(X)*. That is, I have to estimate the conditional probability that a municipality receives

[2] Since I am interested in the average treatment on the treated (ATT) and not in the average treatment effect (ATE), the identifying conditions are weakened. When estimating ATE, the first assumption becomes Y_0, $Y_1 \perp D \mid X$ and the second assumption is $0 < P(D=1 \mid X) < 1$ (Heckman, Ichimura and Todd, 1998)

treatment, given the observable characteristics X. This is usually done by estimating a logit model, where the treatment status *D* is the dependent variable and the set of characteristics X is the independent variable. The choice of variables X in estimating the logit model is particularly important. These control variables X need to be observable and unaffected by the treatment, but should determine the treatment status. The set of X usually contains pretreatment variables and time-invariant characteristics. It often also includes lagged values of the outcome variable.

I specify the logit model that estimates the propensity score using the stratification approach proposed by Dehejia and Wahba (1999, 2002). I hence first estimate the propensity score using a parsimonious model for the covariates. I then divide the sample into several strata so that there is no statistically significant difference between the estimated propensity scores of the treated and the control groups within each stratum. Initially, the sample is divided into 5 strata[3]. If there remains a statistical difference between the estimated propensity scores of the treatment and control group within a stratum, I divide the stratum in half and compare the average propensity scores again. I consequently test for balance of the covariates within each stratum. That is, I use t-tests within each block to check if the mean values for each covariate are the same between the treatment and control group. If there is no balance in a certain block, I add higher-order and interaction terms in the logit model specification until such differences no longer emerge.

Once I have estimated the propensity score, I estimate ATT. In this research I apply two different techniques based on the propensity score to estimate ATT: blocking on the propensity score and reweighting based on the propensity score. I choose these two techniques because they have shown to perform well in small samples with n=100 or n=500 (Busso, McCrary and DiNardo, 2008).

2.3.1 Blocking on the propensity score

The "blocking on the propensity estimator" was first proposed by Rosenbaum and Rubin (1983) and follows immediately from the stratification approach described above. Now that the sample is divided into different strata, I compute the average difference in the outcome variable, $\hat{Y}_m$, between the treatment and control group within each stratum *m*. The blocking estimator is then the weighted average of $\hat{Y}_m$ across the strata, where the weights are the proportion of treated observations in each stratum.

2.3.2 Reweighting based on the propensity score

Reweighting on the propensity score is the second technique I consider to construct a balanced sample of treated and control units in order to estimate treatment effects. Whereas blocking on the propensity score assures that the propensity scores in the treatment and control group are balanced within each strata, reweighting on the propensity score makes the distribution of the propensity score in the entire control group similar to those in the treatment group. I am motivated to use this technique by the Monte Carlo study of Busso, McCrary and DiNardo

[3] Cochran (1968) analyzes a case with a single covariate and shows that under normality conditions 5 or 6 strata remove at least 90% of the bias associated with that covariate. Rosenbaum and Rubin (1984) state that this result also holds for the propensity score. That is, under normality conditions, five strata based on the propensity score will remove over 90 per cent of the bias in each of the covariates.

(2008). These authors find that propensity score-based reweighting estimators are unbiased in small samples and that their variance is very close to the semiparametric efficiency bound.

There are several inverse propensity score weighting estimators (IPW) described in the literature (Busso, McCrary and DiNardo, 2008). I apply the IPW proposed by Johnston and DiNardo (1996) and Imbens (2004), which is most commonly used in empirical studies. This reweighting estimator assures that the sum of the weights add up to the sample size *n*. The weighting function of this estimator is

$$\frac{\hat{p}(X_j)}{1-\hat{p}(X_j)} \Big/ \frac{1}{n_0}\sum_{k=1}^{n}\frac{(1-D_k)\hat{p}(X_k)}{1-\hat{p}(X_k)} \qquad [6]$$

where $\hat{p}(X_j)$ is the estimated propensity score, *n* is the size of the entire sample and n_0 is the size of the control group. Note that these weights should only be applied to the observations in the control group in order to make the mean of each variable X included in the propensity score approximately equal across the treatment and control groups.

3 Description of variables and data

To implement my analysis, I compiled the most detailed data available on agricultural, economic, and general population characteristics for all municipalities in São Paulo. This section first presents some statistics on sugarcane production growth in São Paulo over the past years and motivates the period of analysis. In the next part, I describe how treatment is defined in this study. I then give an overview of the control variables I selected to estimate the propensity score. Finally, I present the outcome variable of interest. I show how this outcome variable differs between the treatment and control group before performing estimations based on the propensity score.

3.1 Period of analysis

This research analyzes the impact of sugarcane production on GDP per capita growth between 2002 and 2006 in São Paulo state. I chose this period of analysis because of data availability issues: the methodology of GDP calculation in Brazil has recently been updated and as a result the only comparable time series on GDP per capita are for the period 2002-2006[4]. I however don't think that this is a critical issue since the main increase in sugarcane harvested area in São Paulo took place after 2001. Figure 1 represents the area and growth in sugarcane harvested in São Paulo between 1997 and 2006. I considered 3-year moving averages to account for the fact that agricultural data are strongly influenced by yearly fluctuations. This figure clearly demonstrates that until 2001 the area devoted to sugarcane cultivation increased slowly on an annual basis, but that from 2001 onwards the annual growth rate rose quickly. Between 2002 and 2006, the average annual increase in sugarcane harvested area in São Paulo state was 6.8%.

[4] In November 2009, IBGE will update the 1999-2001 GDP data using the new methodology. As soon as these data are available, I will extend my period of analysis to 2001-2006.

Figure 1: Sugarcane harvested area in São Paulo, 1997-2006

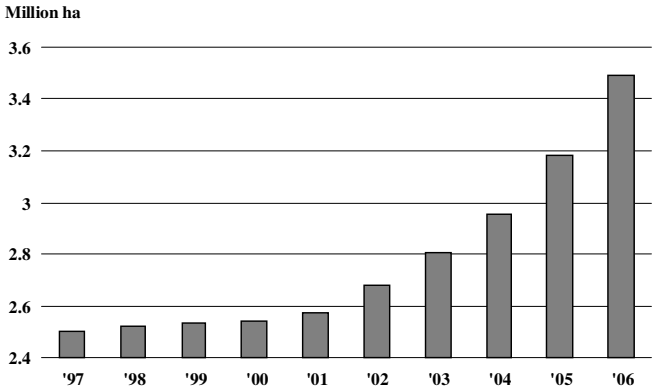

Source: IBGE

3.2 Definition of treatment

I classified municipalities into the treatment or control group by comparing their increase in sugarcane harvested area in the period 2002-2006 to the São Paulo state average of 6.8%. Municipalities where the annual mean rate of increase in sugarcane harvested area was higher than the average of the state were categorized in the treatment group. Municipalities with an annual mean rate of increase in sugarcane harvested area below 0.1% were categorized in the control group. The control group hence includes municipalities with a low or negative growth in sugarcane harvested area as well as municipalities that have never cultivated sugarcane[5]. All other municipalities, i.e. those where the annual mean rate was above 0.1% but below the state average, are removed from the analysis. I removed these municipalities to ensure that the treatment and control group were very different in terms of their growth in sugarcane harvested area.

Table 1 summarizes how many municipalities are classified in the treatment and control groups and how many municipalities have been removed from the analysis. Note that in the final analysis, the treatment and control group are composed of less municipalities due to data availability for some of the selected variables.

Table 1: Composition of treatment and control groups

	Amount of municipalities	Share of total
Treatment group	241	37.4%
Growth sugarcane harvested above 6.8%		
Control group	237	36.7%
Growth sugarcane harvested below 0.1%		
Removed from analysis	167	25.9%
Growth sugarcane harvested between 0.1% and 6.8%		
Total	**645**	**100%**

Source: IBGE

[5] Of the 237 municipalities in the control group, 56 municipalities have an annual mean rate of increase below 0.1% and for 181 municipalities no data on sugarcane harvested area was reported.

3.3 Selection of control variables

As mentioned before, the control variables used to construct the propensity score need to satisfy certain criteria. They need to be observable and unaffected by the treatment in order to satisfy the unconfoundedness condition specified in section 2.2. At the same time, they also need to determine the treatment status. The set of control variables usually contains pretreatment variables and time-invariant characteristics. It often also includes lagged values of the outcome variable. The control variables I selected for this study are listed in Table 2.

Table 2: Control variables: definitions and sources

Variable	Description	Source
area	Area municipality (km^2)	IBGE
sugharv	Sugarcane harvested, average 1990-92 (ha)	IBGE
sugharv/totharv	Share sugarcane harvested in total area of temporary crops harvested, average 1990-92 (%)	IBGE
totharv/area	Share temporary crops harvested in total area municipality, average 1990-92 (%)	IBGE
pasture/area	Share pastureland in total area municipality, 1996 (%)	IBGE
ag_rented/area	Share of municipal area that is rented out for agricultural activities, 1996 (%)	IBGE
ag_occupied/area	Share of municipal area that is occupied for agricultural activities, 1996 (%)	IBGE
ag_partner/area	Share of municipal area that is used in partnerships for agricultural activities, 1996 (%)	IBGE
ag_owned/area	Share of municipal area that is owned for agricultural activities, 1996 (%)	IBGE
rurpop/totpop	Share of rural population in total population, 1991 (%)	IBGE
gdppc80	GDP per capita, 1980 (2000 prices)	IBGE
gdppc96	GDP per capita, 1996 (2000 prices)	IBGE
suitable/area	Share of municipal area suitable for sugarcane production (%)	Gov.SP
suitable_lim/area	Share of municipal area suitable for sugarcane production under environmental limitations (%)	Gov.SP
suitable_restr/area	Share of municipal area suitable for sugarcane production under environmental restrictions (%)	Gov.SP

My data are derived from two main sources: Instituto Brasileiro de Geografia e Estatística (IBGE) and Governo de São Paulo (Gov.SP). IBGE provides most of the data used in this study. The five characteristics for 1996 (pasture/area, ag_rented/area, ag_occupied/area, ag_partner/area and ag_owned/area) are drawn from the agricultural census conducted in 1996. The other agricultural variables (sugharv, sugharv/totharv, and totharv/area) are collected on a yearly basis. I constructed a 3-year average for the period 1990-1992 to eliminate the influence of strong yearly fluctuations in agricultural production. IBGE also publishes statistics on municipal GDP per capita. Since no data on GDP is available for the beginning of the '90s due to the hyperinflation in the early 1990s, I used GDP per capita data for 1980 and 1996.

The Government of São Paulo (Gov.SP) recently published the results of its agro-environmental zoning project in São Paulo. In this project, the area in each municipality is classified according to its suitability to grow sugarcane. There are four different categories: area suitable for sugarcane production, area suitable for sugarcane production under environmental limitations, area suitable for sugarcane production under environmental restrictions, and area not suitable for sugarcane production. I only used the first three variables since the fourth one, i.e. area not suitable for sugarcane production, can be derived from the three other ones and would lead to collinearity in the logit model.

3.4 Outcome variable

The output variable of interest is mean annual growth in municipal GDP per capita between 2002 and 2006. Table 3 gives an idea of how the outcome variable differs between the treatment and control group before doing any estimations based on the propensity score. The mean annual growth in GDP per capita in the control group is on average lower than in the treatment group. I performed a two-sided t-test on the difference between GDP per capita growth between the treatment and control group and found that there is no significant difference (t-value = 1.3896).

Table 3: GDP per capita growth in treatment and control group before estimations based on the propensity score

	Observations	Mean	Std. error	Std. dev.	[95% Conf. Interval]	
Control	235	0.4980276	0.2933819	4.497459	-0.0799798	1.076035
Treatment	236	1.139974	0. 3568545	5.482103	0.4369313	1.843017
Difference		0.6419464	0.4619719			

Based on this preliminary analysis alone, one could conclude that the sugarcane expansion in São Paulo had no significant effect on GDP per capita growth. However, this result needs to be analyzed with caution. Comparing GDP growth per capita between two big groups of municipalities that only differ in their increase of sugarcane planted ignores any other factors that might have influenced GDP growth. The purpose of this study is to exactly avoid such a generalization. In the next section, I will construct estimators based on the propensity score in order to compare the outcome variable for municipalities in the treatment and control group that are similar in terms of the distribution of the observed characteristics.

4 Causal effect of sugarcane expansion on GDP per capita growth

I used Stata to obtain my estimates. Specifically, I used the Stata program *pscore* developed by Becker and Ichino (2002) to estimate the logit model based on stratification. I used the *atts* program written by these same authors to implement blocking on the propensity score and to obtain non-parametric bootstrapped standard errors for these estimators. For the reweighting on the propensity score, I coded the IPW weights in Stata and constructed bootstrapped standard errors using the technique described in Busso and Kline (2008).

4.1 Estimation of the propensity score

I used the stratification approach to construct the logit model that estimates the propensity score. I chose this approach because it is also a valid specification technique (Dehejia and Wahba (1999, 2002)). Indeed, stratification requests that I have balance in the propensity scores and in the covariates within each stratum. The model that I developed passed the specification test. This model contained linear, squared and square root versions of the control variables listed in Table 2. The full model with values for the coefficients, standard errors, z-values and confidence intervals for all the variables can be found in Appendix Table 1. The adjusted R^2 value for this model is 0.4060.

I carried out two tests to assess the goodness-of-fit of my logit model. First, I estimated the predictive power of the logit model using the area under the receiver operating characteristic curve (ROC). The area under the ROC curve is a measure of discrimination; it measures the likelihood that a treated municipality will have a higher probability of being treated than a control municipality. The area under the model's ROC was 0.89, indicating an excellent discrimination. I also used the prediction rate metric to assess whether my specified model does a good job of separating out the treatment and control municipalities. I found that 80.36% of the treatment municipalities were correctly predicted and that 85.29% of the control municipalities were correctly predicted.

Figure 2 illustrates the distribution of the estimated propensity scores in the control group (left panel) and treatment group (right panel). Recall that the propensity score in this study is defined as the conditional probability to increase sugarcane production, given a set of observable covariates. This figure clearly shows that the distributions of the estimated propensity scores for both groups are quite different. In particular, the control group has higher densities for low values of the estimated propensity scores while the treatment group has higher densities for high values of the estimated propensity scores. This indicates that the control group is composed of relatively more municipalities that are characterized by a low predicted probability to expand sugarcane production. In the treatment group, most of the municipalities are characterized by a higher probability to expand sugarcane production. The model satisfies the common support condition as the highest estimated propensity score is strictly lower than 1 (namely 0.993013).

Figure 2: Histograms of the estimated propensity scores for the control group (left panel) and treatment group (right panel)

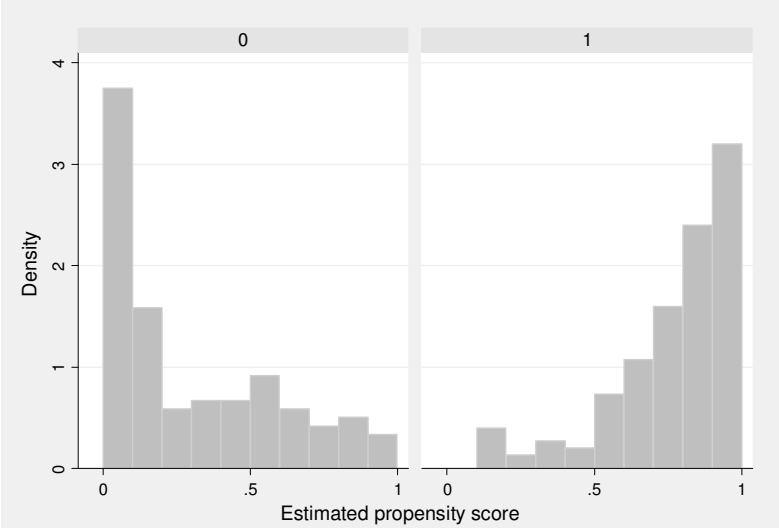

4.2 Analysis of ATT estimates

4.2.1 ATT estimates from blocking on the propensity score

After specifying the logit model using stratification, it is straightforward to estimate ATT using the blocking estimator. The ATT is obtained as the weighted average of the ATT of each stratum, where the weights are the proportion of treated observations in each stratum.

I constructed two blocking estimators that differ in the region of common support they consider. For the first blocking estimator, I imposed the common support restriction as defined by Dehejia and Wahba (1999, 2002). This approach deletes all observations in the control group with propensity score values lower than the minimum of those in the treatment group. Based on this approach, the common support was [0.136613, 0.993013]. I needed five strata to obtain balance in the propensity scores and in the mean values of the covariates within each stratum. Since this approach eliminated 54 of the 120 control observations, I also constructed a second blocking estimator for which I didn't impose a common support condition. With the stratification technique, I needed six strata to balance propensity scores and mean values of the covariates within each stratum. Since the first stratum contained 45 control observations but no treatment observations, these control observations were discarded from the analysis. As a result, the second blocking estimator had a larger region of analysis than the first one, but still didn't include the entire sample of control observations. In particular, the second region of common support became [0.105868, 0.993013]. This second blocking estimator can hence be considered as a sensitivity analysis. Indeed, one of the major concerns with imposing common support restrictions is that one might eliminate observations at the boundaries which could have an important impact on the result. The second common support includes 9 more control observations at the lower bound and therefore prevents that high quality matches are lost at the boundaries of the common support.

Figure 3 and Table 4 compare the quintile means of the estimated propensity scores for the two blocking estimators. For both estimators the estimated propensity scores are very different in the overall sample but are similar within each stratum. The means of the estimated propensity scores in the overall sample and the first quintile are slightly lower in the second blocking estimator compared to the first blocking estimator. This is because the second blocking estimator includes 9 more observations in the control group of the first quintile. These 9 extra observations are the ones with estimated propensity score values above 0.105868 but below 0.136613, and hence lower the mean value of the estimated propensity score in the first quintile and in the overall sample of the control group for the second blocking estimator.

Figure 3: Comparison of quintile means of the estimated propensity scores for both blocking estimators

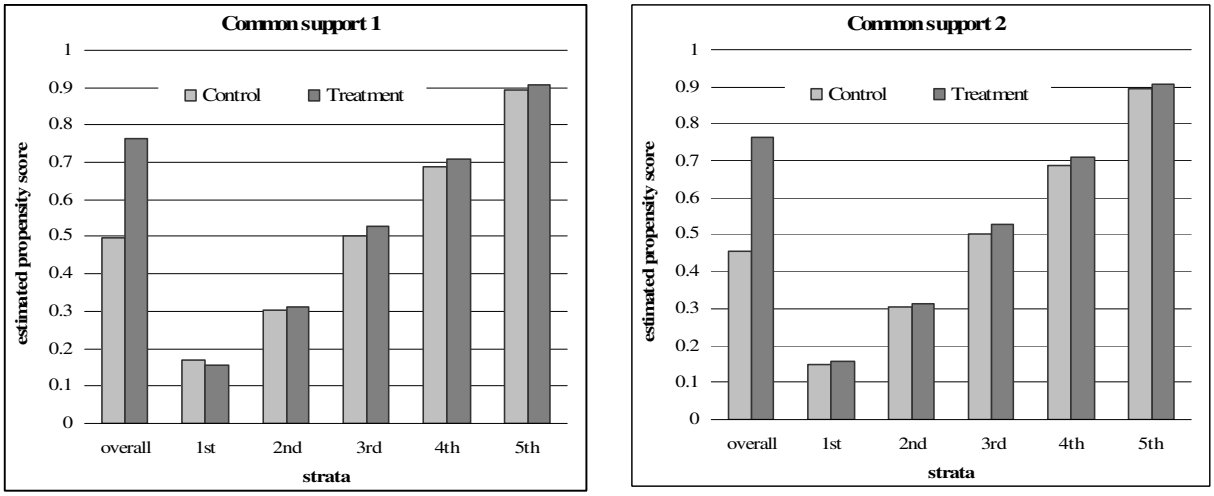

Table 4: Comparison of quintile means and standard deviations of the estimated propensity scores for both blocking estimators

		Common support 1 [0.136613, 0.993013] Propensity score			Common support 2 [0.105868, 0.993013] Propensity score		
		N	**mean**	**(SD)**	**N**	**mean**	**(SD)**
Overall	control	66	0.4996	(0.2416)	75	0.4545	(0.2576)
	treatment	150	0.7643	(0.2086)	150	0.7643	(0.2086)
1st quintile	control	10	0.1682	(0.0158)	19	0.1473	(0.0265)
	treatment	6	0.1544	(0.0178)	6	0.1544	(0.0178)
2nd quintile	control	15	0.3042	(0.0551)	15	0.3042	(0.0551)
	treatment	6	0.3142	(0.0824)	6	0.3142	(0.0824)
3rd quintile	control	19	0.5021	(0.0631)	19	0.5021	(0.0631)
	treatment	14	0.5282	(0.0407)	14	0.5282	(0.0407)
4th quintile	control	12	0.6867	(0.0302)	12	0.6867	(0.0302)
	treatment	40	0.7080	(0.0558)	40	0.7080	(0.0558)
5th quintile	control	10	0.8948	(0.0392)	10	0.8948	(0.0392)
	treatment	84	0.9061	(0.0559)	84	0.9061	(0.0559)

4.2.2 ATT estimates from reweighting on the propensity score

Reweighting the observations in the control group based on the estimated propensity scores aims at making the distribution of the estimated propensity scores of the control and treatment groups more similar. I reweighted the propensity scores using the IPW weighing function described in section 2.3.2. I considered the same two regions of common support as I used for the blocking estimators.

Figure 4 demonstrates the kernel densities of the estimated propensity scores before and after reweighting. The left hand panel is the kernel density plot of the estimated propensity scores before reweighting. This panel is the density version of the histogram in Figure 2 and shows again that the distributions of the estimated propensity scores for the treatment and control group are very different. The second and third panels are the kernel densities after reweighting on common support 1 and common support 2, respectively. Note that the density of the treatment group remained unchanged compared to the first panel since the treatment group is not reweighted. Also note that the left panel considers the entire sample of control observations, while the middle and right panel only consider those control observations on common support 1 and common support 2, respectively. This figure clearly shows that reweighting based on the estimated propensity score succeeded at making the distributions of the estimated propensity scores for the control group very similar to those of the treatment group.

Figure 4: Kernel densities of estimated propensity scores before and after reweighting

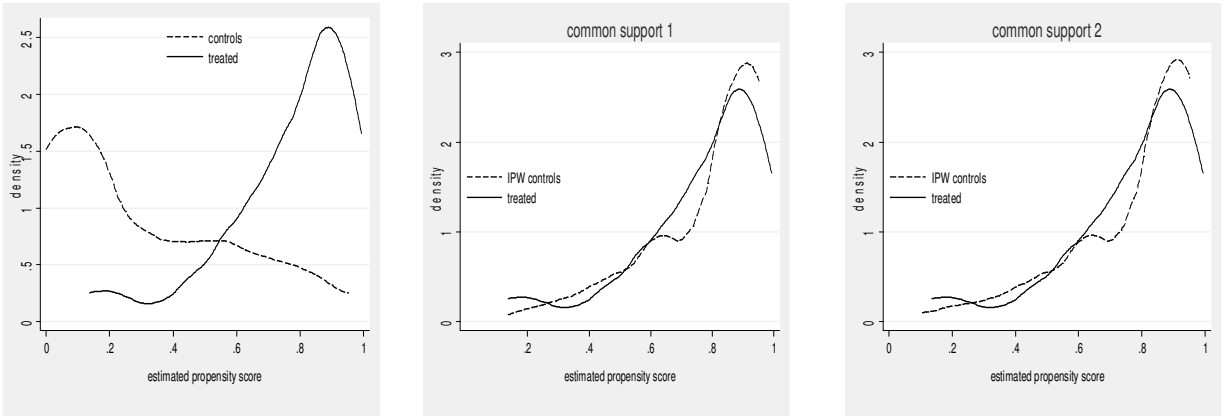

Since I do not condition on all the covariates but on the propensity score, I also checked if the reweighting procedure was able to balance the distributions of the control variables in both the treatment and control group. Before reweighting differences are expected, but after reweighting the covariates should be balanced in both groups and hence no significant differences should be found. I applied the technique described in Rosenbaum and Rubin (1985) to perform this specification test[6]. These authors used a two-sample t-test to check if there are significant differences in covariate means for both groups. Appendix Table 2 illustrates the results of the specification test. My model passes the specification test since before reweighting 8 of the 15 covariates demonstrated significant differences in their means between the treatment and control group, but after reweighting none of the covariates' means were significantly different between the treatment and control group for neither common support 1 nor common support 2.

[6] Note that stratification technique I used to estimate the propensity score is an alternative approach used to check for balance. Whereas stratification tests the mean differences in propensity scores and covariates in each stratum, the specification test described in Rosenbaum and Rubin (1985) tests the mean differences in propensity scores and covariates in the entire sample.

4.2.3 Comparison of ATT estimates obtained with different techniques

In this study, ATT measures how sugarcane expansion has impacted average annual GDP per capita growth in the municipalities that expanded sugarcane production during the period 2002-2006. I used two different types of estimators based on the propensity score to estimate the ATT. Since the propensity scores used to construct these estimators were all estimated using the same logit model specification, I can compare the results obtained from these estimators.

Table 5 presents values for ATT and summary statistics for the different estimators. I obtained bias, standard errors, t-values, mean squared errors (MSE) and 95% confidence intervals for each of these estimators using bootstrap procedures with 10,000 replications. For the blocking estimators I applied non-parametric bootstrapping, while for the reweighting estimators I constructed bootstraps as described in Busso and Kline (2008).

For the blocking estimators, the ATT is slightly higher for the estimators on common support 1 than for the estimators on common support 2. The same result holds when comparing the two reweighting estimators. The overall performance of the estimators can be compared by analyzing the MSE. The MSE between the two blocking estimators don't differ significantly, nor do the MSE between the two reweighting estimators. However, the MSE of the reweighting estimators are higher than the MSE of the blocking estimators which indicates that the blocking estimators are more effective. This result runs contrary to that of Busso, DiNardo and McCrary (2008), who found that in Monte Carle simulations with N=100 and N=500 the reweighting estimators are more effective than blocking estimators.

Table 5: ATT and summary statistics for the different estimators

	Blocking		IPW	
	common support 1	**common support 2**	**common support 1**	**common support 2**
ATT	1.0476	1.0167	1.0570	0.9926
t-test	1.352	1.302	0.747594	0.717538
bias	0.0036	-0.0064	-0.0002	0.0126
std. error	0.7747	0.7809	1.4138	1.3833
MSE	0.600	0.610	1.999	1.914
95% conf. interval	[-0.4711,2.5662]	[-0.5140,2.5473]	[-1.7144,3.8284]	[-1.7190,3.7041]
# obs. control	66	75	66	75
# obs. treatment	150	150	150	150
N	216	225	216	225
common support	[0.136613, 0.993013]	[0.105868, 0.993013]	[0.136613, 0.993013]	[0.105868, 0.993013]

Note: Values for t-test, bias, standard errors, MSE and 95% confidence intervals are obtained using bootstrap procedures with 10,000 replications

Most importantly, all estimators give the same result, namely that the ATT is statistically insignificant. This result passes the sensitivity analysis since it holds for the estimators on common support 2. Indeed, the estimators on common support 2 were constructed to check whether the observations at the boundaries of common support 1 had a significant impact on the result. This result implies that sugarcane expansion did not have a significant impact on GDP per capita growth in those municipalities that expanded sugarcane production between 2002 and 2006.

5 Conclusion

The increase in sugarcane production in Brazil is considered to be linked to GDP per capita growth in sugarcane-expanding regions. I investigated this claim by analyzing the effects of sugarcane expansion on GDP per capita growth in sugarcane-expanding municipalities in São Paulo state. This state was selected because the majority of sugarcane expansion between 2000 and 2006 occurred here. Using estimators based on the propensity score, I estimated the average treatment effect on the treated (ATT) to examine whether the sugarcane-expanding municipalities in São Paulo state experienced a higher growth in GDP per capita compared to their non-sugarcane expanding counterparts.

I classified municipalities into two groups: the treatment group, which is composed of sugarcane-expanding municipalities, and the control group, which includes municipalities that didn't expand sugarcane production. Contrary to previous studies, I controlled for a set of variables that might have caused a difference in mean annual GDP per capita growth between the treatment and control group. In particular, I used non-parametric techniques based on the propensity score to ensure that I compared municipalities in the control group that were similar to municipalities in the treatment group in terms of these control variables. In this study, the propensity score is defined as the conditional probability that a municipality expanded its sugarcane production, given a set of observable control variables. I used two techniques that are based on the propensity score: blocking on the propensity score and reweighting on the propensity score. For both techniques, I considered one estimator with a small common support and one with a larger common support.

The estimators gave similar results with respect to ATT. They all indicated that sugarcane expansion had no statistically significant impact on GDP per capita growth in sugarcane-expanding municipalities. These results challenge the findings of Walter (2008) and Sparovek et al. (2008) who established positive and statistically significant effects. Even though their studies and my research are not directly comparable because we consider different time periods and units of analysis, I think that my results are more robust. First, my counterfactual scenarios control for more factors that could have influenced GDP per capita growth. Second, my model specification passes two different balancing tests and hence ensures that the treatment and control group are similar in the mean values of their propensity scores and covariates. Third, my results also pass a sensitivity analysis by considering a larger region of common support, which includes more control observations at the boundaries.

My findings show that the sugarcane-expanding municipalities in São Paulo did not experience a larger growth in GDP per capita than their non-sugarcane expanding counterparts. This result can be explained by the fact that sugarcane has mainly replaced other crops in São Paulo. Indeed, sugarcane harvesting in São Paulo has been characterized by increasing mechanization. This mechanization has replaced laborers employed in the sugarcane sector, but might also, through crop substitution, have replaced laborers employed in other agricultural activities (Guilhoto et al., 2002). Furthermore, the expansion of sugarcane has led to a decrease in the amount of small-scale sugarcane farmers compared to an increase in large-scale sugarcane farmers (Veiga Filho and Ramos, 2006).

It will be interesting to apply this analysis to other states in Brazil, such as the states in the Central South. These states possess the agro-ecological conditions to grow sugarcane and are already increasing their sugarcane production. Furthermore, most of the planned expansion of sugarcane in Brazil will occur here. Since sugarcane will mainly replace pastureland in these states, the impacts of sugarcane expansion might be more significant.

Implications for future research are as follows. First, I will explore other techniques based on the propensity score, in particular propensity score matching and mixed methods, to estimate ATT. Second, I will repeat the analysis for the period 2001-2007 as soon as IBGE releases longer time series for GDP per capita. Third, I will extend this analysis to other regions in Brazil, such as the municipalities in the Central South. Once the results of the national agro-ecological zoning project are released, there will be data on all areas in Brazil that are suitable for sugarcane production. I will then be able to also estimate the average effect of sugarcane expansion on the control group, i.e. the average effect on the untreated, to assess the impact of the planned expansion of sugarcane in the states in the Central South of Brazil.

6 References

Becker, S.O. and A. Ichino. 2002. Estimation of average treatment effects based on propensity score. *Stata Journal* 2(4):358-377. Download software from http://www.sobecker.de/pscore.html

Blundell, R. and M. Costa Dias. 2000. Evaluation methods for non-experimental data. *Fiscal Studies* 21(4):427-468

Busso, M., J. DiNardo and J. McCrary. 2008. Finite sample properties of semiparametric estimators of average treatment effects. Working paper, University of Michigan

Busso, M. and P. Kline. 2008. Do Local Economic Development Programs Work? Evidence from the Federal Empowerment Zone Program. Cowles Foundation Discussion Paper 1638. Cowles Foundation, Yale University.

Cochran, W.G. 1968. The effectiveness of adjustment by sub-classification in removing bias in observational studies. *Biometrics* 24: 295-313

Dehejia, R. and S. 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. and S. Wahba. 2002. Propensity score matching methods for nonexperimental causal studies. *Review of Economics and Statistics* 84(1): 151-161.

Governo de São Paulo. 2008. Zoneamento agroambiental para o setor sucroalcooleiro. Online dataset. Available at: http://homologa.ambiente.sp.gov.br/etanolverde/zoneamento.asp

Guilhoto, J.J.M., A.L.M. de Barros, M.C. Marjotta-Maistro and M. Istake. 2002. O impacto da colheita da cana-de açucar sobre a geração de emprego nos setores produtores de cana-de-açucar, álcool e áçucar e nas suas macroregiones. (Versão Preliminar). ESALQ/CEPEA, Universidade de São Paulo, Piracicaba.

Heckman, J.J. 1979. Sample Selection Bias as a Specification Error. *Econometrica* 47:931-961

Heckman, J.J. and R. Robb. 1985. Alternative models for evaluating the impact of interventions. In J. Heckman and B. Singer (eds) Longitudinal Analysis of Labor Market Data (pp. 156-245). Cambridge: Cambridge University Press

Heckman, J.J. 1997. Instrumental variables: A study of implicit behavioral assumptions used in making program evaluations. *Journal of Human Resources* 32(3):441-462

Heckman, J.J., H. Ichimura and P. Todd. 1998. Matching as an econometric evaluation estimator. *Review of Economic Studies* 65:261-294

Heckman, J.J., R. LaLonde and J. Smith. 1999. The economics and econometrics of active labor market programs. In O. Ashenfelter and D. Card (eds) Handbook of Labor Economics (Vol. III, pp. 1865-2097). Amsterdam: Elsevier

Holland, P.W. 1986. Statistics and causal inference. *Journal of the American Statistical Association* 81:945-960

Imbens, G.W. 2004. Nonparametric estimation of average treatment effects under exogeneity: a review. *Review of Economics and Statistics* 86(1):4–29

Instituto Brasileiro de Geografia e Estatística (IBGE). 2009. Online database. Available at: http://www.ibge.gov.br

Jank, M.S. 2007. Criteria for Sustainable Biofuels: Views from the Brazilian Ethanol Industry. Presentation at the "Clean vehicles and fuels: European symposium and exhibition", Stockholm, Sweden, 8 November, 2007.

Johnston, J. and J.E. DiNardo. 1996. Econometric Methods. 3rd ed. McGraw-Hill, Tokyo

King, Gary and Langche Zeng. 2006. The dangers of extreme counterfactuals. *Political Analysis* 14(2):131-159

Lechner, M. 1999. Earnings and employment effects of continuous off-the-job training in East Germany after unification. *Journal of Business Economic Statistics* 17(1):74-90

Macedo, Isaias de Carvalho (editor). 2005. A energia da cana-de-açúcar: doze estudos sobre a agroindústria da cana-de-açúcar no Brasil e sua sustentabilidade. União da Agroindústria Canavieira do Estado de São Paulo (UNICA), São Paulo, Brazil.

Ramos, P. 2008. Os impactos da expansão do lavoura canaviera na estrutura fundiária e as manifestações de sua concentração no Brasil. Presentation at the "Workshop sobre os impactos da evolução do setor sucroalcooleiro". UNICAMP, 16 May 2008, Campinas, Brazil.

Rosenbaum, P.R. and D.B. Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70:41-55

Rosenbaum, P. R. and D.B. Rubin. 1984. Reducing bias in observational studies using subclassification on the propensity score. *Journal of the American Statistical Association*, 79:516-524

Rosenbaum, P. R. and D.B. Rubin. 1985. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician* 39(1):33-38

Sparovek, G. et al. 2008. Environmental, land-use and economic changes in Brazilian recent sugarcane expansion (1996-2006). (in review)

Veiga Filho, A. de A. and P. Ramos. 2006. Proálcool e evidências de concentração na produção e processamento de cana-de-açúcar. Informações Econômicas, Vol 36, n.7, July 2006, pp 48-61. São Paulo, Brazil.

Walter, A. 2008. Sustainability of Ethanol Production in Brazil. Presentation at the "Workshop internacional Brasil-Japão em biocombustível, meio ambiente e novos productos da biomassa". UNICAMP, 29 Oct – 1 Nov, 2008, Campinas, Brazil (unpublished).

7 Appendix

Appendix Table 1: Logit Model used to estimate propensity scores

Logistic Regression				Number of observations		270
				LR chi2(27)		150.59
				Prob>chi2		0.0000
Log Likelihood		-110.18273		Pseudo R2		0.4060
	Coef.	Std. Err.	z	P>z	[95% Conf.	Interval]
area	-0.0068442	0.0027931	-2.45	0.014	-0.0123186	-0.0013699
sugharv	0.0000793	0.000237	0.33	0.738	-0.0003853	0.0005439
sugharv/totharv	0.064511	0.0297832	2.17	0.030	0.0061369	0.122885
totharv/area	-0.1381926	0.0669679	-2.06	0.039	-0.2694472	-0.006938
pasture/area	0.0502768	0.0387647	1.30	0.195	-0.0257006	0.1262542
gdppc96	0.0001285	0.0001408	0.91	0.361	-0.0001474	0.0004045
ag_rented/area	0.2653065	0.3353345	0.79	0.429	-0.3919369	0.92255
ag_occupied/area	-0.1173905	0.3225961	-0.36	0.716	-0.7496673	0.5148862
ag_owned/area	-1.117622	0.5330151	-2.10	0.036	-2.162312	-0.0729314
rurpop/totpop	-0.0190729	0.0457702	-0.42	0.677	-0.1087807	0.070635
suitable_lim/area	-0.0257309	0.0258117	-1.00	0.319	-0.0763208	0.0248591
suitable_restr/area	0.0395651	0.0285334	1.39	0.166	-0.0163593	0.0954895
areasqrt	0.4174687	0.135302	3.09	0.002	0.1522816	0.6826558
totharv/areasqrt	1.458569	0.582629	2.50	0.012	0.3166369	2.600501
gdppc80sqrt	-0.0320398	0.0185945	-1.72	0.085	-0.0684843	0.0044047
pasture/areasqrt	0.5986544	0.8173727	0.73	0.464	-1.003367	2.200675
ag_partner/areasq	-1.12E-08	1.03E-08	-1.09	0.275	-3.13E-08	8.90E-09
sugharv/totharvsq	-0.0008046	0.0003235	-2.49	0.013	-0.0014387	-0.0001705
pasture/areasq	-0.000388	0.000368	-1.05	0.292	-0.0011093	0.0003333
gdppc96sq	5.32E-11	7.11E-09	0.01	0.994	-1.39E-08	1.40E-08
rurpop/totpopsq	0.0002228	0.000638	0.35	0.727	-0.0010276	0.0014732
suitable/areasq	0.0003659	0.0001391	2.63	0.009	0.0000933	0.0006385
suitable_lim/areasq	0.0004878	0.0002457	1.99	0.047	6.26E-06	0.0009692
suitable_restr/areasq	-0.0001419	0.0003581	-0.40	0.692	-0.0008439	0.00056
ag_rented/areasq	-0.0084559	0.0055979	-1.51	0.131	-0.0194276	0.0025159
ag_occupied/areasq	-0.0076147	0.0279433	-0.27	0.785	-0.0623826	0.0471532
ag_owned/areasq	0.0068732	0.0037263	1.84	0.065	-0.0004302	0.0141766
constant	33.77142	23.47553	1.44	0.150	-12.23977	79.78261

Note: the suffixes “sq” and “sqrt” stand for squared and square root, respectively

Appendix Table 2: Assessing the balance in covariates before and after reweighting based on the propensity score

	full sample		treatment		control		difference	reweighted treatment[a]		reweighted control *common support 1*		difference	reweighted control *common support 2*		difference
variable	mean	std. error	mean	std. error	mean	std. error		mean	std. error	mean	std. error		mean	std. error	
area	471.371	22.073	497.559	28.068	438.637	35.054	58.922	497.559	28.068	511.895	37.644	-14.336	510.975	35.282	-13.416
sugharv	2094.394	293.189	2182.573	308.913	1984.169	536.344	198.404	2182.573	308.913	2160.726	353.623	21.847	2167.101	340.910	15.472
sugharv/totharv	17.902	1.622	20.778	2.016	14.309	2.614	6.469**	20.778	2.016	23.241	3.021	-2.463	23.145	2.837	-2.367
totharv/area	16.172	0.892	17.701	1.019	14.261	1.540	3.440*	17.701	1.019	17.050	1.315	0.651	17.051	1.253	0.650
pasture/area	46.338	1.489	53.941	1.638	36.833	2.389	17.108***	53.941	1.638	53.041	2.600	0.900	52.936	2.444	1.005
gdppc80	6750.100	1676.776	4780.974	128.038	9211.508	3766.001	-4430.530	4780.974	128.038	5178.409	249.633	-397.435	5171.423	233.974	-390.449
gdppc96	5183.08	261.486	5160.327	220.673	5211.521	521.090	-51.194	5160.327	220.673	5404.076	270.061	-243.749	5392.420	252.953	-232.093
ag_rented/area	6.403	0.395	6.183	0.433	6.678	0.707	-0.495	6.183	0.433	5.722	0.486	0.461	5.744	0.465	0.439
ag_occupied/area	1.684	0.196	0.928	0.106	2.628	0.405	-1.700***	0.928	0.106	0.885	0.156	0.043	0.892	0.147	0.036
ag_partner/area	1.953	0.191	1.945	0.251	1.962	0.296	-0.017	1.945	0.251	2.154	0.415	-0.209	2.161	0.392	-0.216
ag_owned/area	89.961	0.537	90.944	0.548	88.732	0.987	2.212**	90.944	0.548	91.239	0.614	-0.295	91.202	0.593	-0.258
rurpop/totpop	29.604	1.036	28.121	1.177	31.458	1.800	-3.337	28.121	1.177	25.003	1.852	3.118	25.128	1.746	2.993
suitable/area	13.655	1.595	18.208	2.479	7.964	1.680	10.244***	18.208	2.479	13.049	2.798	5.159	13.004	2.616	5.204
suitable_lim/area	41.861	2.303	48.368	3.049	33.728	3.379	14.640***	48.368	3.049	52.049	4.551	-3.681	51.881	4.264	-3.513
suitable_restr/area	20.417	1.561	24.146	2.134	15.755	2.221	8.391***	24.146	2.134	25.221	3.126	-1.075	25.071	2.927	-0.925
N	270		150		120			150		66			75		

* indicates that difference between treatment and control group is statistically significant at 95% level,
** indicates that difference between treatment and control group is statistically significant at 97.5% level,
*** indicates that differences between treatment and control group is statistically significant at 99% level
t-tests are used for difference in means
[a] values for reweighted treatment are equal to the values before reweighting because only the control group was reweighted.