THE ECONOMIC IMPACTS OF SUGARCANE EXPANSION IN BRAZIL

by Annelies Deuss

This thesis/dissertation document has been electronically approved by the following individuals:

Kyle,Steven Charles (Chairperson)

Bento,Antonio Miguel R. (Minor Member)

Just,David R. (Minor Member)

THE ECONOMIC IMPACTS OF SUGARCANE EXPANSION IN BRAZIL

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Annelies Deuss

August 2010

THE ECONOMIC IMPACTS OF SUGARCANE EXPANSION IN BRAZIL

Annelies Deuss, Ph. D.

Cornell University 2010

Since 2001, Brazil has experienced a sharp increase in sugarcane production due to the upsurge in demand for sugar and ethanol, two products derived from sugarcane. While the increase in sugarcane production has led to income and employment opportunities in the sugar and ethanol sector, the benefits of sugarcane expansion could vary significantly by the region where sugarcane is cultivated. This dissertation consists of three studies that examine the economic impacts of the recent sugarcane expansion in Brazil.

Whereas previous studies only show associations between sugarcane expansion and economic indicators, this research establishes a causal relationship using estimators based on the propensity score. The propensity score is defined in this research as the probability that a municipality expands sugarcane production, given a set of observable characteristics. One of these characteristics is the suitability of a municipality to grow sugarcane. Data on suitability of land were recently published at the national level.

The first study analyzes whether municipalities in São Paulo state that expanded sugarcane production between 2002 and 2006 as a result experienced higher economic growth. The results indicate that there is no statistically significant impact of sugarcane expansion on GDP per capita growth.

The second study examines the economic growth impacts of the increased sugarcane production in the different sugarcane producing regions in Brazil. The findings show that sugarcane expansion led to GDP per capita growth in three regions: in Brazil as a whole, in the North-Northeast and in the Center-South excluding São Paulo. In addition, it is demonstrated that this latter region could benefit from future sugarcane expansions.

The final study investigates the underlying reasons for the findings in the first study. It examines the impact of sugarcane expansion in São Paulo state on growth in GDP per capita, in employment and in wages in the different sectors of the economy. The results suggest that sugarcane expansion has positive impacts on local economies in São Paulo state. Further research with updated data is needed to establish whether the positive influences at sector level affected total GDP data in lagged terms.

BIOGRAPHICAL SKETCH

Annelies Deuss was born and raised in Belgium. In 1999, she graduated from the Katholieke Universiteit Leuven, Belgium, with a Master of Science degree in Agricultural Engineering with a specialization in Agricultural Economics. In 2001, she earned a Master of Arts degree in Development, Innovation and Change from the Universitá di Bologna, Italy. She then worked for almost five years as an economist at the Food and Agriculture Organization of the United Nations. In 2005, she was accepted into the Ph.D. program in Applied Economics on a Cornell Graduate Fellowship. Upon earning her Ph.D., in August 2010, she will move to Pittsburgh, PA, to start her job as a visiting assistant professor in economics at Carnegie Mellon University.

To Majo and Sergio

ACKNOWLEDGMENTS

Above all, I would like to thank Steven Kyle, the chair of my committee, for his advice and guidance. I feel extremely fortunate that I had a mentor at Cornell University who was always available to give me support and encouragement. Our discussions were both entertaining and intellectually stimulating. I appreciate how he taught me to balance work and life and am grateful to him for funding my first trip to Brazil. I would also like to thank David Just and Antonio Bento for serving as special members on my committee and for their comments and input on earlier versions of my dissertation.

I would like to acknowledge the institutions and people who funded my studies at Cornell. In particular, I would like to thank the Department of Applied Economics and Management for the Cornell graduate fellowship during the first year; for the tuition fellowship for my entire studies; and for the graduate assistantships for several semesters. I am also grateful to Daniel Sisler's family and Cyril F. Crowe's family for their generous support.

My field work to Brazil was financed by a combination of 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. Their support is greatly appreciated.

I would like to thank the people I met in Brazil for sharing their time, insight and invaluable information. Without implicating them in the conclusions, I want to thank 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.

There are two people in the department to whom I am greatly indebted. I want to thank David Lee for giving me interesting research and teaching opportunities and for inspiring me to be a better researcher and teacher. I greatly enjoyed the conversations we shared and it was a pleasure working together. I also want to thank Linda Morehouse for helping me with practical questions. Her swift and efficient replies to every request made graduate life so much easier.

Finally, I would like to thank my friends at Cornell for making these past five years an incredible experience. I feel grateful to have met so many interesting and fun people, with whom I could discuss research and hang out. And last but not least, I am grateful to my mom and Sergio for believing in me. Without their love and support, none of this would have been possible.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

ATT	Average Treatment Effect on the Treated
ATU	Average Treatment Effect on the Untreated
BR	Brazil
CS	Center-South
CSex	Center-South excluding São Paulo state
EMBRAPA	Empresa Brasileira de Pesquisa Agropecuária
GDP	Gross Domestic Product
IBGE	Instituto Brasileiro de Geografia e Estatística
IPEA	Instituto de Pesquisa Econômica Aplicada
IPW	Inverse Propensity Score Weighting Estimators
MSE	Mean Squared Errors
NE	North-Northeast
PNUD	Programa de las Naciones Unidas para el Desarrollo
RAIS	Relação Anual de Informações Sociais
SP	São Paulo state
VA	Value Added

CHAPTER 1

INTRODUCTION

Sugarcane has been for centuries one of Brazil's main agricultural crops. Since 2001, the country has experienced a sharp increase in sugarcane production. The reason behind this increase was the upsurge in demand for sugar and ethanol, two products that are derived from sugarcane. In a country with an ideal climate and abundant amounts of suitable land available for sugarcane plantations, the impact of this recent sugarcane expansion on economic growth is assumed to be strictly positive. There are, however, indications that the benefits of sugarcane expansion could vary significantly by the region where it is cultivated. Given that Brazil plans to double the amount of land dedicated to sugarcane plantations in the next 10 years, an analysis of the regional impacts of sugarcane expansion would not only indicate where past expansions were most beneficial to the local economies, but also where future expansions should be located.

The aim of this research is to determine the impacts of sugarcane expansion in Brazil on local economies. In this introduction, we will first present a short overview of the sugarcane, sugar and ethanol sector in Brazil. We will then review the current literature that analyzes the economic impacts of sugarcane expansion in Brazil. Finally, we will give the outline of this dissertation and its contribution to current research.

1.1 The sugarcane, sugar and ethanol sector in Brazil

Sugarcane production in Brazil is concentrated in two areas: the North-Northeast and the Center-South of the country (see Appendix Figure 1 for a map of these areas). The North-Northeast was traditionally the main sugarcane growing region in the country, but was surpassed by the Center-South. The poorer soil quality and lower degree of mechanization in the North-Northeast compared to the Center-South led to lower productivity and higher costs of growing sugarcane in the North-Northeast than in the Center-South (Krivonos and Olarreaga 2006). In 2007, 16 percent of sugarcane was cultivated in the North-Northeast and 83 percent in the Center-South[1]. Most of the plantations in the Center-South are located in the state of São Paulo: this state alone was responsible for 55 percent of sugarcane production (IBGE 2010).

The recent growth in sugarcane production is mainly explained by the increase in planted area since yields have remained more or less constant for the last thirty years[2] (Brandão 2007). In fact, the area devoted to sugarcane production in Brazil has increased from 4.9 million hectares to 6.5 million hectares between 2001 and 2007, and is projected to amount to 13.9 million hectares by 2020 (Jank 2007). This increase in area harvested between 2001 and 2007 has been obtained by clearing new lands, expansion into pasture land and by replacing crop land for sugarcane plantations (Altieri 2008).

One crucial feature of the sugarcane industry is the close relationship between harvesting and processing – the raw cane has to be transported to the mill fast, since

[1] The remaining one percent is grown in other areas than the Center-South and the North-Northeast
[2] Sugarcane yields increased by only about 1 percent per year between 1970 and 2006, while the planted area expanded at an annual rate of 4 percent during that same period (Brandão 2007).

the quality of sugar deteriorates rapidly following the cutting of cane (Ueki 2007). The distance between production and the processing plant is limited to a radius of approximately 50 km (Brandão 2007). Sugarcane will then be either processed into sugar or into ethanol. Between 1980 and 2000, the share of ethanol in sugarcane production has been higher than the share of sugar. From 2000 onwards, these shares have become more or less equal (Macedo 2005).

Sugarcane is farmed on lands owned or rented by sugarcane processing millers and on lands owned by independent sugarcane farmers. Around 75 percent of sugarcane is grown by the mills, which hire seasonal workers at hourly wages, while the rest belongs to independent producers (Moraes and Pessini 2004). The 60,000 independent sugarcane producers form a very heterogeneous group with farm sizes ranging from 10 to 500 ha.

Prior to the economic liberalization in the 1990s, the Brazilian sugar and alcohol industry was highly regulated. Sugar mills and distilleries received credit guarantees and subsidized interest prices. Through the Sugar and Alcohol Institute (IAA), the government set sugar production quotas and allocated them among the sugar mills and distilleries. The IAA also fixed prices paid to sugarcane growers (Krivonos and Olarreaga 2006). Since 1999, the state considerably turned away from the sugar and alcohol sector. As a result, the sugarcane, sugar and alcohol prices became market-determined; production quotas were abolished; and ethanol producers no longer received subsidies. The only way the government still can influence prices is by changing the mix of ethanol and gasoline (Brandão 2007).

With the deregulation of the prices, the price of sugarcane became determined by a new mechanism, developed by the Consecana[3] (Sachs 2007). This new system is based on the quality of raw materials and the market price obtained for end-products (sugar and ethanol). The price of sugar is negotiated in a free and transparent market where the domestic prices depend on international prices and on the exchange rate. The price of ethanol is also determined in an open market but the role of Petrobras, the state petroleum company, and of large distributors of fuels has to be acknowledged.

1.2 Literature review

The impact of sugarcane expansion on the local economies can be examined by looking at a wide set of variables. The main variables that are discussed in the current literature –and that will be analyzed in this dissertation - are: Gross Domestic Product (GDP), employment and income.

The employment effects of the increased demand for sugar and ethanol have been extensively studied. Macedo (2005) shows that the amount of formal direct jobs in the sugarcane, sugar and ethanol sector combined rose from 643.000 to 983.000 between 2000 and 2005. Smeets et al. (2006) point out that there are also large indirect and induced employment effects: in the late 1990s these are calculated to be 940.000 and 1.800.000 jobs, respectively. Moraes (2007) however shows that despite a growth of 54.6% in the sugarcane production between 1992 and 2005, there was a reduction of 23% in the number of employees[4] in the sugarcane production sector. The increase in direct employment hence occurred especially in the sugar and ethanol sector. The

[3] The Consecana or Conselho dos Produtores de Cana-de-Açúcar, Açúcar e Álcool do Estado de São Paulo, was established in 1997 and consists of 5 representatives from the sugarcane producers and 5 representatives from the industrial sector.
[4] Formal and informal employees

main reason for the declining number of sugarcane workers is the increasing mechanization of the sugarcane harvest. Although mechanization of the sector holds many economic, social and environmental advantages, the issue emerging is that mechanization reduces the demand for workers, especially low-skilled workers.

Hoffmann and Oliveira (2008a) show that the average earnings in the sugarcane, sugar and ethanol sector in the period 2002-2006 rose by 36.0%, 1.7% and 5.3%, respectively. Even though the earnings in the sugarcane sector show the highest rate of increase, it is worthwhile noting that the earnings in this activity are still far below the earnings in the sugar and ethanol sector. The authors also demonstrate that the increase in average earnings in the sugarcane sector is strongly correlated with the rising value of the national real minimum wage. When comparing the average earnings of sugarcane workers with those of workers in other crop cultivations, Hoffmann and Oliveira (2008b) show that in Brazil as a whole, the sugarcane cutters earn more than their counterparts in most other agricultural activities. In São Paulo however, they note that two different data sources give contrasting results. According to the PNAD[5] database, the average earnings in the sugarcane sector are among the highest in the agricultural sector in São Paulo. The IEA/CATI[6] database however reveals that in São Paulo the sugarcane cutters are among the lowest earners in the agricultural sector. Moraes and Pessini (2004) analyze the relationship between the price of sugarcane and the salaries in the sugarcane and sugar sector. They demonstrate that in years where the sugarcane prices dropped, the salary reduction was

5 PNAD is the acronym for Pesquisa Nacional por Amostra de Domicílios, which is the National Household Sample Survey

6 IEA/CATI is the acronym for Instituto de Economia Agrícola/Coordenadoria de Assistência Técnica Integral

higher and that in years where these prices increased, the wage increase would be proportionally lower.

At present, there are two studies that analyze the impact of increased sugar and ethanol exports on regional income distribution and employment opportunities in the sugarcane, sugar and ethanol industries. Burnquist et al. (2004) find that a demand shock resulting from an expansion of sugar exports presented an impact of greater magnitude upon the countries' production and employment when compared to the impact from an increase in ethanol. In addition, their empirical results indicate that when the impact is generated in the North-Northeast, production and employment are more affected than when it begins at the Center-South of the region. Krivonos and Olarreaga (2006) assess the impact that changes in domestic sugar prices have on regional wages and employment depending on worker characteristics and measure the impact on household income of a 10 percent increase in world sugar prices. The authors find that workers in the sugar sector and sugar-producing regions have better employment opportunities and experience larger wage increases. More interestingly, they show that households at the top of the income distribution experience larger income gains due to higher wages, whereas households at the bottom of the distribution experience larger income gains due to movements out of unemployment.

Two studies, by Walter (2008) and by Sparovek et al. (2009), investigate the economic effects of the increased sugarcane production at municipality-level. Walter (2008) analyzes the economic and social impacts of the presence of mills and sugarcane production in various municipalities in the states of São Paulo and Alagaos for the year 2000. The author first divides the municipalities in São Paulo into two groups, municipalities with mills and municipalities without mills, and assures that the municipalities in both groups contained a population between 2.5 million to 500 million people. He then compares the 2000 values of the following indicators per group: monthly income per person, gini index, income of the 20% poorest of the population, electricity and human development index. The group of municipalities with mills scored better on all indicators compared to the group of municipalities without mills. The author then compares municipalities in São Paulo that together produce more than 90% of the total amount of sugarcane with the remaining municipalities in São Paulo, again controlling for population. Also in this situation, the municipalities with a high production of sugarcane score better on all the indicators. He finally repeats the last exercise for municipalities in Alagaos and shows that the municipalities with a high production of sugarcane have better social and economic indicators.

Sparovek et al. (2009) measure the environmental, land use and economic changes of Brazil's sugarcane expansion over the period 1996-2006. The authors divide the municipalities in Brazil into two groups (sugarcane expansion and no sugarcane expansion municipalities) according to three indicators: the presence of mills, the area cultivated with sugarcane in 2006 and the increase in harvest sugarcane area between 1997 and 2006. The authors show that in the Center-South as well as in the peripheral expansion areas of the country, 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.

1.3 Contribution of this research

The current literature holds a wide variety of studies that analyze the impacts of sugarcane expansion in Brazil on different aspects of the economy. Most of these studies however focus on one variable only, such as for example employment or income, and examine the effect of the ethanol expansion on this variable in the sugarcane, sugar and ethanol sector without looking at the impact in other sectors. The analysis is in most cases limited to the country as a whole, or at best disaggregates the effects at regional or state-level. In addition, most papers only present descriptive statistics and hence lack the power to put forward causal linkages. There are presently only two studies (Walter 2008; and Sparovek et al. 2009) that investigate the impact of sugarcane expansion on a set of economic indicators at municipal-level and beyond the sugarcane-related sectors. Although the initial concept of Walter's (2008) exercise is interesting, it lacks the robustness to draw any meaningful conclusion since the author uses only one control variable, namely population, and limits his analysis to two states in Brazil and to one point in time. Sparovek et al. (2009) have a more thorough approach since they analyze municipalities in entire Brazil for a 10-year time period and control for regional effects by comparing neighboring municipalities only. The economic impact analysis is however limited to comparing GDP and GDP growth between the two groups and no causal effects can be demonstrated.

This research will contribute to the existing literature by establishing a causal relationship between sugarcane expansion and its effects on local economies. Whereas the methodologies in the current literature only allow for associations between sugarcane expansion and different economic indicators, we use propensity score-based estimators to establish the true causality of change. We selected these types of estimators because they take into account that different municipalities have a different propensity to expand sugarcane production. We apply these estimators in three different studies, which each examine a different aspect of the impacts of sugarcane expansion. The three studies, presented in Chapters 2 through 4, were written in sequence and hence the objectives of the second and third study are a direct result of the first study's outcome.

In Chapter 2, we study the economic growth impacts of sugarcane expansion in São Paulo state. In particular, we examine whether municipalities in São Paulo state that increased their sugarcane production between 2002 and 2006 experienced as a result a higher growth in GDP per capita. We select this state for the first study because over 63 percent of the Brazilian sugarcane expansion between 2002 and 2006 occurred here. We use two different types of estimators based on the propensity score that perform well in small samples: blocking estimators and propensity score reweighting estimators. With these estimators, we estimate 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 sugarcane non-expanding counterparts. Our results are robust and show that sugarcane expansion did not have a significant impact on GDP per capita growth in those municipalities that expanded sugarcane production.

Chapter 3 extends the study conducted in Chapter 2 to the entire country and examines the same outcome variable, namely GDP per capita growth, over the period 2001-2007. Given that sugarcane production in Brazil is concentrated in two areas, namely the North-Northeast and the Center-South of the country, we estimate and compare the ATT on three regional levels: i) in Brazil as a whole, ii) between the North-Northeast and the Center-South, and iii) between São Paulo state and the Center-South excluding São Paulo state. These levels of analysis allow us not only to draw conclusions at national level, but also to compare the situation in the relatively poorer North-Northeast with the richer Center-South as well as compare the new sugarcane producing municipalities in the Center-South with those in São Paulo state. In this chapter, we also estimate for the Center-South excluding São Paulo state the average effect of sugarcane expansion on the sugarcane non-expanding municipalities, i.e. the average treatment effect on the untreated (ATU). This allows us to assess the impact of the future expansion of sugarcane production, which is planned in this region.

In Chapter 4, we redirect the analysis to the state of São Paulo. The results of Chapter 2 and Chapter 3 show that sugarcane expansion in São Paulo state didn't lead to greater growth in GDP per capita in those municipalities that increased sugarcane production. In this chapter, we study the underlying reasons behind these findings by analyzing the impact of sugarcane expansion on the different sectors of the economy during the period 2002-2006. In particular, we look at three sets of outcome variables, namely GDP per capita, employment and wages, and analyze the influence of sugarcane expansion on the growth of these variables at both the aggregate level and by sector.

It is important to note that the three chapters were written sequentially and that the choice of the region and period of analysis were influenced by the availability of the data. We decided to focus the first study, i.e. Chapter 2, on the state of São Paulo because one of the crucial covariates for the estimation of the propensity score, namely suitability of the land for sugarcane production, was only available for São Paulo state. The period of analysis in Chapter 2 extends from 2002 to 2006 and doesn't consider a longer time frame because the outcome variable, GDP per capita, was only available for those years. By the time we started working on the second study, discussed in Chapter 3, the data on suitability of the land for sugarcane production had become available for the entire country and the GDP data had been updated to cover the period 2001-2007. In the last study, we again studied the period 2002-2006 and not 2001-2007 because the GDP data by sector were only available and comparable for the former time period.

CHAPTER 2

THE ECONOMIC GROWTH IMPACTS OF SUGARCANE EXPANSION IN SÃO PAULO STATE

2.1 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 state 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. (2009) 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 2002 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 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. (2009) 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. (2009) 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 applied 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 article 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. The third section 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. The fourth section 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. The final section concludes.

2.2 The impact evaluation problem

We are 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, we 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, we need to take into account two issues. First, we need to establish the direction of causality. That is, we need to ensure that sugarcane expansion influenced GDP per capita growth and not vice versa. We address this potential endogeneity problem by considering lagged values of the control variables in the analysis. Second, we want 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 we 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 we are dealing with non-experimental data. As a result, we are 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 we 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). We overcome this problem by using estimators based on the propensity score, a non-parametric approach first proposed by Rosenbaum and Rubin (1983). In particular, we apply two different techniques: blocking on the propensity score and reweighting based on the propensity score.

2.2.1 Theoretical aspects

Our 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[7]. We then define the outcome (GDP per capita growth) for sugarcane-expanding (treated) municipalities, i.e. the municipalities for which Di=1, as Y1i and the outcome for sugarcane non-expanding (non-treated or control) municipalities, i.e. the municipalities for which Di=0, as Y0i.

$$Y_i \equiv \begin{cases} Y_{0i} & if\ D_i = 0 \\ Y_{1i} & if\ D_i = 1 \end{cases} \qquad (2.1)$$

[7] The classification into treatment or control group is described in detail in the next section.

The causal effect of treatment (sugarcane expansion) in a certain municipality is given by the difference in the potential outcomes with and without treatment, Y1i – Y0i. Since a municipality will either expand sugarcane production (Di=1) or not expand sugarcane production (Di=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 we cannot compute the individual treatment effect. We can however estimate the average effect, which compares the average outcomes of the treated and non-treated groups.

In this study, we are interested in estimating the average treatment effect on the treated (ATT). In other words, we 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):

$$(2.2) \qquad ATT \equiv E(Y_1 - Y_0 | D = 1) = E(Y_1 | D = 1) - E(Y_0 | D = 1)$$

The non-experimental design of this research implies that we cannot directly identify the counterfactual outcome but that we have to estimate it. In particular, we need to estimate E(Y0|D=1). In this research we use non-parametric techniques based on the propensity score to estimate the counterfactuals. we 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).

2.2.2 The propensity score

When we estimate the counterfactual outcomes, we want to make sure that we are comparing municipalities in the control group that are as similar as possible to those in the treatment group. In other words, we 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. Equation (2.2) can hence be rewritten as:

$$(2.3) \qquad ATT \equiv E(Y_1 - Y_0 | X, D = 1) = E(Y_1 | X, D = 1) - E(Y_0 | X, D = 1)$$

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 ρ(X). This balancing score is defined as the conditional probability of being treated, given the observed covariates X, or $\rho(X) = \text{Prob}(D=1 | X)$.

Rosenbaum and Rubin (1983) demonstrate that if potential outcomes are independent of treatment conditional on covariates X, they are also independent of treatment conditional on the propensity score ρ(X):

$$(2.4) \qquad Y_0, Y_1 \perp D \mid X \;\rightarrow\; Y_0, Y_1 \perp D \mid \rho(X)$$

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 using the propensity score to estimate ATT, two identifying assumptions8 need to be invoked. 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 we 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

$$(2.5) \qquad Y_0 \perp D \mid X$$

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.

[8] Since we are 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 $0 < Prob(D=1 \mid X) < 1$ (Heckman, Ichimura and Todd 1998).

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.

(2.6) $$Prob(D=1 \mid X) < 1$$

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

2.2.3 Estimations based on the propensity score

When we use propensity-based estimators to estimate ATT, we first have to estimate the propensity score $\rho(X)$. That is, we have to estimate the conditional probability that a municipality receives 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.

We specify the logit model that estimates the propensity score using the stratification approach proposed by Dehejia and Wahba (1999; 2002). We hence first estimate the propensity score using a parsimonious model for the covariates. We 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[9]. If there remains a statistical difference between the estimated propensity scores of the treatment and control group within a stratum, we divide the stratum in half and compare the average propensity scores again. We consequently test for balance of the covariates within each stratum. That is, we 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, we add higher-order and interaction terms in the logit model specification until such differences no longer emerge.

Once we have estimated the propensity score, we estimate ATT. In this research we apply two different techniques based on the propensity score to estimate ATT: blocking on the propensity score and reweighting based on the propensity score. We 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).

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, we compute the average difference in the outcome variable, $\bar{Y}_m$, between the treatment and control group within each stratum m. The blocking estimator is then the weighted average of $\bar{Y}_m$

[9] 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.

across the strata, where the weights are the proportion of treated observations in each stratum.

The second technique we use to construct a balanced sample of treated and control units is reweighting on the propensity score. Whereas blocking on the propensity score assures that the propensity scores in the treatment and control group are balanced within each stratum, reweighting on the propensity score makes the distribution of the propensity score in the entire control group similar to the one in the treatment group. We are motivated to use this technique by the Monte Carlo study of Busso, McCrary and DiNardo (2008). These authors find that propensity score-based reweighting estimators are unbiased in small samples and that their variance is very close to the semi-parametric efficiency bound.

There are several inverse propensity score weighting estimators (IPW) described in the literature (Busso, McCrary and DiNardo 2008). We 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

$$\text{(2.7)} \qquad \frac{\hat{\rho}(X_j)}{1-\hat{\rho}(X_j)} \Bigg/ \frac{1}{n_0}\sum_{k=1}^{n}\frac{(1-D_k)\hat{\rho}(X_k)}{1-\hat{\rho}(X_k)}$$

where $\hat{\rho}(X_j)$ 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.

2.3 Description of variables and data

To implement our analysis, we 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, we describe how treatment is defined in this study. We then give an overview of the control variables we selected to estimate the propensity score. Finally, we present the outcome variable of interest. We show how this outcome variable differs between the treatment and control group before performing estimations based on the propensity score.

2.3.1 Period of analysis

This research analyzes the impact of sugarcane production on GDP per capita growth between 2002 and 2006 in São Paulo state. We 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. We 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 2.1 represents the area and growth in sugarcane harvested in São Paulo between 1997 and 2006. We 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%.

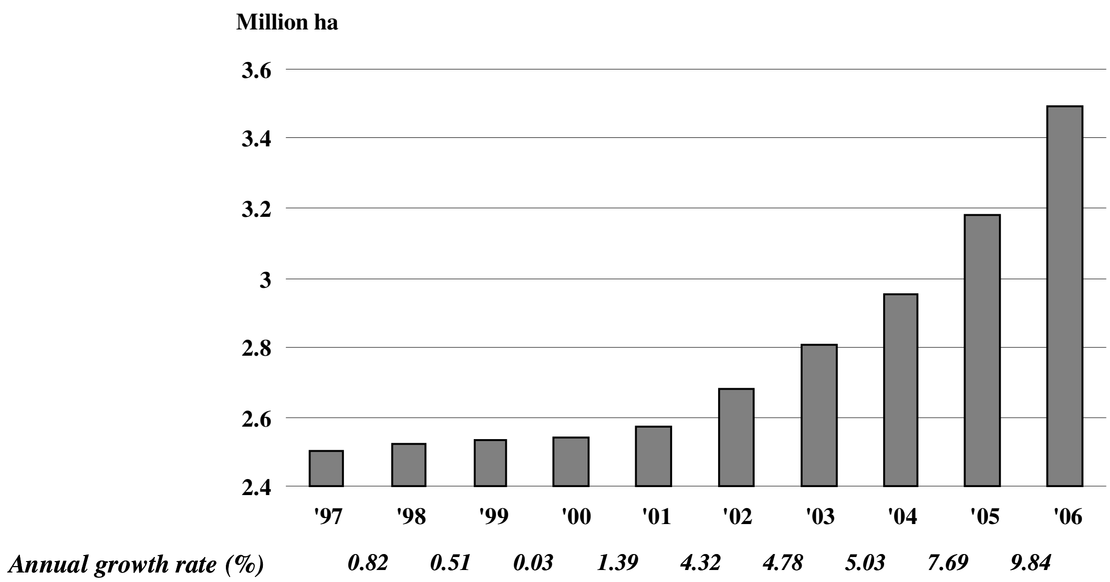

Source: IBGE

Figure 2.1: Area and growth in sugarcane harvested in São Paulo, 1997-2006

2.3.2 Definition of treatment

We 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 equal or 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 equal or below 0% 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 sugarcane10. All other municipalities, i.e. those where the annual mean rate was above 0% but below the state average, are removed from the analysis. We removed these municipalities to ensure that the treatment and control group are significantly different in terms of their growth in sugarcane harvested area.

Table 2.1 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 2.1: Composition of treatment and control groups

	Amount of municipalities	**Share of total**
Treatment group Growth sugarcane harvested ≥ 6.8%	241	37.4%
Control group Growth sugarcane harvested ≤ 0%	236	36.6%
Removed from analysis Growth sugarcane harvested between 0% and 6.8%	168	26.0%
Total	**645**	**100%**

Source: IBGE

2.3.3 Selection of control variables

As mentioned above, 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. 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

[10] Of the 236 municipalities in the control group, 55 municipalities have an annual mean rate of increase below 0% and for 181 municipalities no data on sugarcane harvested area was reported.

values of the outcome variable. The control variables we selected for this study are listed in Table 2.2.

Table 2.2: 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 limitation (%)	Gov.SP
suitable_restr/area	Share of municipal area suitable for sugarcane production under environmental restriction (%)	Gov.SP

The 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. We 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, we 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 production11. We 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.

2.3.4 Outcome variable

The output variable of interest is mean annual growth in municipal GDP per capita between 2002 and 2006. Since we are interested in growth rates, we are using GDP per capita data at constant 2000 prices, which are provided by the Instituto de Pesquisa Econômica Aplicada[12] (IPEA). Table 2.3 gives an idea of how the outcome variable differs between the treatment and control group before doing any estimation based on

---

[11] The four different categories are defined as follows. (1) areas suitable for sugarcane production: areas with favorable climatic conditions for the cultivation of sugarcane and without any specific environmental constraints; (2) areas suitable under environmental limitations: areas with favorable climate and soil for sugarcane cultivation but classified as Environmental Protection Areas (APA), or as medium priority areas for enhancing the connectivity, as directed by the BIOTA-FAPESP Project; or as critical watersheds; (3) suitable areas with environmental constraints: areas with favorable climatic conditions for the cultivation of sugarcane but classified as buffer zones of the Conservation Units of Integral Protection (UCPI), or as high priority areas for increased connectivity as indicated by the BIOTA-FAPESP Project, or as areas of high vulnerability for the groundwater in the State of São Paulo, as published by CETESB-IG-DAEE - 1997; (4) areas not suitable or inadequate areas: areas classified under the Conservation Units of Integral Protection (UCPI) at State and Federal level; areas classified as extremely important for biological conservation, indicated by the BIOTA-FAPESP Project for the creation of Conservation Units of Integral Protection (UCPI); areas classified as Zones Wildlife Areas Environmental Protection (EPA); areas with soil and climatic constraints to grow sugarcane; and areas with slopes steeper than 20%.
The land in São Paulo state is classified as follows: 26% are suitable areas, 45% are suitable areas with environmental restrictions, 28% are suitable areas with environmental restrictions, and only 1% are inadequate areas

[12] Institute of Applied Economic Research

the propensity score. The mean annual growth in GDP per capita in the control group is on average lower than in the treatment group. We 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 2.3: 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	0.498	0.293	4.497	-0.080	1.076
Treatment	236	1.140	0. 357	5.482	0.4373	1.843
Difference		0.642	0.462			

Source: IPEA

Based on this preliminary analysis alone, one could conclude that the sugarcane expansion in São Paulo state 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 in area planted with sugarcane 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, we 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.

2.4 Causal effect of sugarcane expansion on GDP per capita growth

We used Stata to obtain our estimates. Specifically, we used the Stata program pscore developed by Becker and Ichino (2002) to estimate the logit model based on stratification. We 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, we coded the IPW weights in Stata and constructed bootstrapped standard errors using the technique described in Busso and Kline (2008).

2.4.1 Estimation of the propensity score

We used the stratification approach to construct the logit model that estimates the propensity score. We chose this approach because it is also a valid specification technique (Dehejia and Wahba 1999; Dehejia and Wahba 2002). Indeed, stratification requests that we have balance in the propensity scores and in the covariates within each stratum. The model that we developed passed the specification test. This model contains linear, squared and square root versions of the control variables listed in Table 2.2. 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.

We carried out two tests to assess the goodness-of-fit of our logit model. First, we 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. We also used the prediction rate metric to assess whether our specified model does a good job of separating out the treatment and control municipalities. We found that 80.36% of the treatment municipalities were correctly predicted and that 85.29% of the control municipalities were correctly predicted.

Figure 2.2 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).

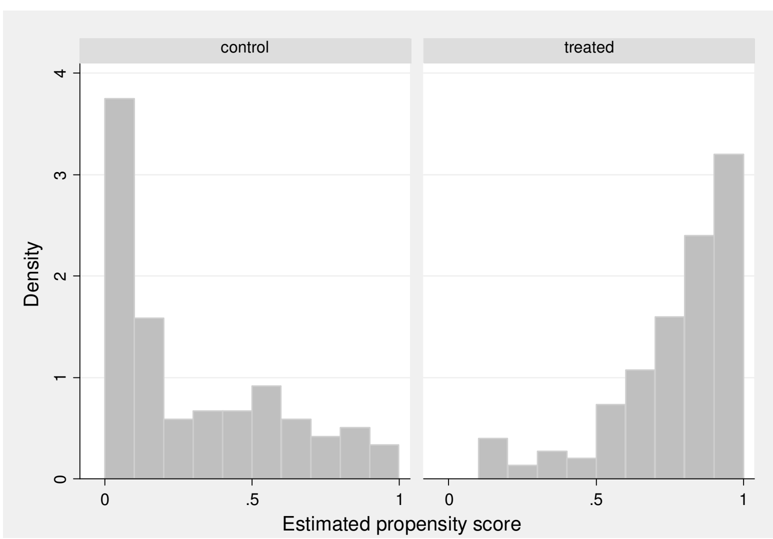

Figure 2.2: Histograms of the estimated propensity scores for the control group (left panel) and treatment group (right panel)

2.4.2 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.

We constructed two blocking estimators that differ in the region of common support they consider. For the first blocking estimator, we 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]. We 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, we also constructed a second blocking estimator for which we didn't impose a common support condition. With the stratification technique, we 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 2.3 and Table 2.4 compare the overall and 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.

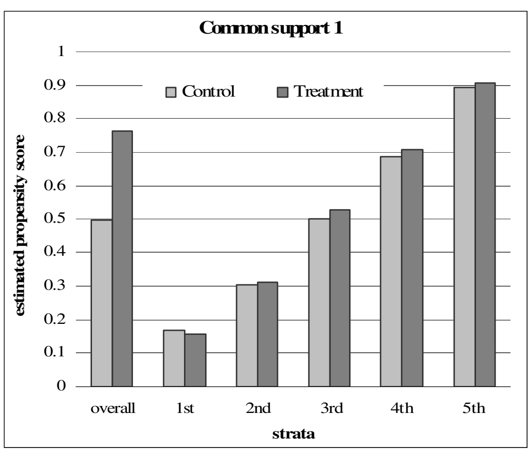

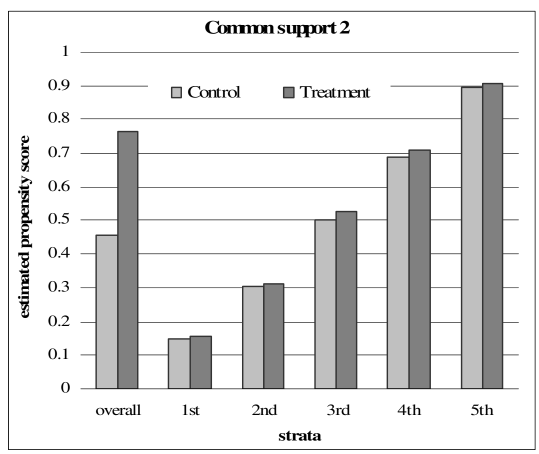

Figure 2.3: Overall and quintile means of the estimated propensity scores for blocking estimators on common support 1 (left) and common support 2 (right)

Table 2.4: Comparison of quintile means and standard deviations of the estimated propensity scores for both blocking estimators

		Common support 1 [0.136613, 0.993013]			Common support 2 [0.105868, 0.993013]		
			Propensity score			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)

2.4.3 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. We reweighted the propensity scores using the IPW weighing function described in equation (2.7). We considered the same two regions of common support as we used for the blocking estimators.

Figure 2.4 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 2.2 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.

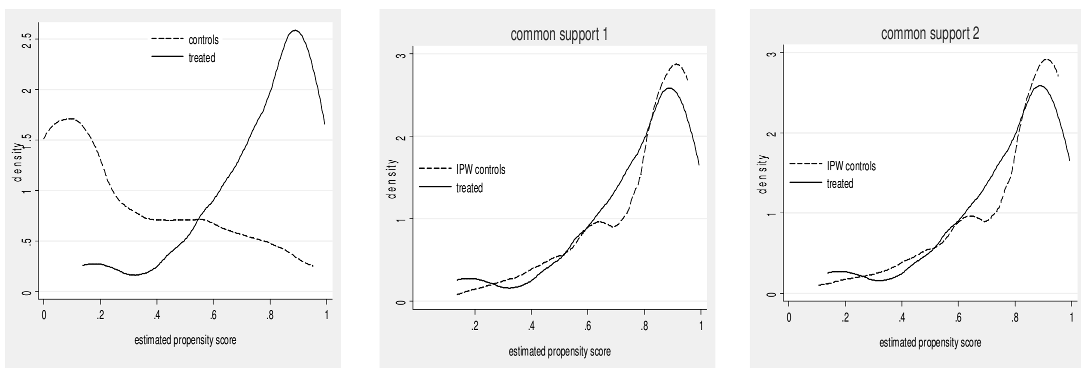

Figure 2.4: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (middle, right)

Since we do not condition on all the covariates but on the propensity score, we 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. We applied the technique described in Rosenbaum and Rubin (1985) to perform this specification test[13]. 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. Our 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.

2.4.4 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. We 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, we can compare the results obtained from these estimators.

Table 2.5 presents values for ATT and summary statistics for the different estimators. We 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 we applied non-parametric bootstrapping, while for the reweighting estimators we constructed bootstraps as described in Busso and Kline (2008).

---

[13] Note that the stratification technique we 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.

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 2.5: ATT and Summary Statistics for the Different Estimators

	Blocking common support 1	**common support 2**	**IPW 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	[-0.471,2.566]	[-0.514,2.547]	[-1.714,3.828]	[-1.719,3.704]
# obs. control	66	75	66	75
# obs. treatment	150	150	150	150
N	216	225	216	225
common support	[0.1366, 0.9930]	[0.1059, 0.9930]	[0.1366, 0.9930]	[0.1059, 0.993]

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.

2.5 Policy implications

Brazil is planning to double the amount devoted to sugarcane plantations by 2020 (Jank 2007). The result obtained in this study suggests that these future expansions should be planned carefully. In particular, the finding that sugarcane expansion in São Paulo state did not lead, on average, to economic growth has several implications for policy recommendations and further research.

First, a more detailed analysis is necessary to examine which sectors in São Paulo state benefited from sugarcane expansions and which sectors didn't. Previous studies show that sugarcane expansions lead to income and employment opportunities in the sugar and ethanol sector (Macedo 2005). However, the insignificant impact of sugarcane expansion on economic growth indicates that there have been negative effects in other sectors that offset the positive impacts in the sugar and ethanol sector. It will hence be useful to compare income and employment effects in the sugarcane, sugar, and ethanol sector with other sectors not related to sugarcane. A clear understanding of which sectors and which segments of the labor population didn't gain from sugarcane expansions will be paramount to design policies that prevent further unemployment or income losses.

Second, this type of research should also be applied to other states in Brazil. Particular attention should be given to the states in the Center-South since most of the projected expansion will take place in these states. By estimating the average treatment effect on the untreated, it will be possible to estimate how the planned expansion will influence the economies in the newly expanding municipalities. Policy makers can then use this information to recommend which areas will most probably benefit the most from sugarcane expansions.

2.6 Conclusion

The increase in sugarcane production in Brazil is considered to be linked to GDP per capita growth in sugarcane-expanding regions. We 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 most of the sugarcane expansion since 2000 occurred here. Using estimators based on the propensity score, we 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 between 2002 and 2006 compared to their sugarcane non-expanding counterparts.

We 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, we 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, we used non-parametric techniques based on the propensity score to ensure that we 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. We used two techniques that are based on the propensity score: blocking on the propensity score and reweighting on the propensity score. For both techniques, we 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. (2009) who established positive and statistically significant effects. Even though their studies and our research are not directly comparable because we consider different time periods and units of analysis, we consider our results more robust. First, our counterfactual scenarios control for more factors that could have influenced GDP per capita growth. Second, our 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, our results also pass a sensitivity analysis by considering a larger region of common support, which includes more control observations at the boundaries.

Our findings show that the sugarcane-expanding municipalities in São Paulo did not experience a larger growth in GDP per capita than their sugarcane non-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).

Implications for future research are as follows. First, it will be interesting to examine why sugarcane expansion had no significant impact on economic growth in São Paulo state by analyzing employment and income effects in different sectors of the economy. Second, it will be relevant for policy formulations to apply this analysis to other states in Brazil, such as the states in the Center-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.

CHAPTER 3

INTER-REGIONAL ANALYSIS OF THE ECONOMIC GROWTH IMPACTS OF SUGARCANE EXPANSION IN BRAZIL

3.1 Introduction

In this chapter we analyze the economic growth impacts of sugarcane expansion in the different regions in Brazil. From 2001 onwards, this country has experienced a sharp increase in sugarcane production. The reasons behind this increase were the upsurges in demand for both sugar and ethanol, two products that are derived from sugarcane. In a country with an ideal climate and abundant amounts of suitable land available for sugarcane plantations, the impact of this recent sugarcane expansion on economic growth is assumed by many to be positive, at least on a net basis. In terms of employment generation, Macedo (2005) shows that the amount of formal direct jobs in the sugarcane, sugar and ethanol sector combined rose from 643.000 to 983.000 between 2000 and 2005. Hoffmann and Oliveira (2008) demonstrate that the average earnings in the sugarcane, sugar and ethanol sector in the period 2002-2006 rose by 36.0%, 1.7% and 5.3%, respectively

There are, however, indications that the national aggregate image masks regional differences and that the benefits of sugarcane expansion could vary significantly depending on the region where sugarcane is cultivated. The two areas where sugarcane production in Brazil is concentrated are the North-Northeast and the Center-South of the country (see Appendix Figure 1 for a map of these areas). The poorer soil quality and lower degree of mechanization in the North-Northeast are responsible for the lower productivity and higher costs of growing sugarcane there than in the Center-South (Krivonos and Olarreaga 2006). Macedo (2005) demonstrates that in 2005, the mean monthly salaries in the sugarcane, sugar and ethanol sector were respectively 58.7%, 78.4% and 48.5% higher in the Center-South than in the North-Northeast.

The differences in the mean monthly salaries in the sugarcane and sugarcane-related industries between the Center-South and the North-Northeast only reveal part of the story. Even though these salaries are lower in the North-Northeast, sugarcane expansion could have a larger impact on the economic growth of local economies in this region. Burnquist et al. (2004) analyze the impacts of a demand shock resulting from an expansion of sugar and ethanol exports. Their empirical results indicate that when the shock is generated in the North-Northeast, production and employment get a greater boost than when the shock is generated in the Center-South.

A further distinction should also be made in the Center-South sugarcane producing region. In 2007, 83 percent of sugarcane was harvested in the Center-South, but the lion's share of sugarcane plantations, namely 55% of the nation's total, were found in the state of São Paulo. This state was also responsible for most of the growth in sugarcane production that occurred since 2001. Accordingly, there are many studies available that analyze the sugarcane sector in São Paulo state and its impacts on employment, income and economic growth. In Chapter 2, we show that in São Paulo state, sugarcane expanding municipalities didn't experience a significantly greater growth in GDP per capita between 2002 and 2006 than the sugarcane non-expanding municipalities.

Even though a large share of sugarcane is produced in the remaining states in the Center-South, most studies analyze either the region as a whole or focus on the state of São Paulo. Sparovek et al. (2009) compare GDP and GDP growth between neighboring sugarcane-expanding and sugarcane non-expanding municipalities in the Center-South and show that the values of these indicators are higher in the former group. Nevertheless, it is crucial to also examine the impacts of sugarcane expansion on economic growth in the Center-South region excluding São Paulo since the strong presence of sugarcane in São Paulo dominates the overall picture in this region. Another argument to study the remaining states in the Center-South as a separate region is that most of the future sugarcane expansions, which are projected to double the current amount of land cultivated with sugarcane by 2020, are planned in these states (Jank 2007).

The current literature contains a wide variety of studies that analyze the impacts of the increased sugarcane sector in Brazil on different aspects of the economy. Most of these studies, however, focus on one variable only, such as employment or income, and examine the effect of sugarcane expansion on this variable in the sugarcane, sugar and ethanol sector without looking at the impact in other sectors. The analysis is in most cases limited to the country as a whole, or at best compares the effects in the North-Northeast, Center-South and São Paulo state, but doesn't consider the region of the Center-South excluding São Paulo. In addition, most papers only present descriptive statistics and thus lack the power to analyze causal linkages. There is presently only one study (Sparovek et al. 2009) that considers an analysis beyond the sugarcane-related sectors by considering a general indicator of economic growth, i.e. GDP growth. Then again, their study cannot demonstrate a causal relation and is limited to presenting associations between sugarcane expansion and economic growth because they don't control for potential confounding factors that could have influenced GDP growth.

This study aims at filling this gap by analyzing the economic growth impacts of sugarcane expansion in Brazil and in the different sugarcane growing regions. Contrary to previous studies that examine income and employment generation in sugarcane-related sectors, we evaluate economic growth because this will give us a more complete picture of the impact on local economies. We study these impacts in Brazil as a whole, in the North-Northeast and Center-South, and make an additional distinction between São Paulo and the Center-South region excluding São Paulo.

We adopt a methodology that allows for drawing causal inferences between sugarcane expansion and economic growth. In particular, we construct two different types of counterfactual scenarios to establish the true causality of change. With the first counterfactual scenario, we analyze for the sugarcane-expanding municipalities what would have happened to economic growth if they hadn't expanded sugarcane production. The second counterfactual scenario is only applied to those municipalities in the Center-South region excluding São Paulo that didn't expand sugarcane production. We examine what would have been the effect on economic growth if they had indeed expanded sugarcane production. Analyzing what could have been the effect in these municipalities will indicate the potential impacts of future sugarcane expansions in this region.

The counterfactual scenarios are constructed using techniques based on the propensity score. We use this technique, first developed by Rosenbaum and Rubin (1983a), because it accounts for the fact that different municipalities have a different propensity to expand sugarcane. The causal effect of sugarcane expansion on economic growth is then estimated with four different estimators that are each based on the propensity score.

The rest of this chapter is organized as follows. The methodology is explained in section 2. In section 3, we describe the data and variables. The empirical results are presented in section 4. Section 5 concludes.

3.2 Methodology

In this research, we analyze whether sugarcane expansion in Brazil had a significant impact on economic growth. We address this issue in two different ways. First, we analyze whether municipalities that expanded sugarcane production experienced higher economic growth as a result. Second, we examine what would have happened to the economic growth in those municipalities that didn't expand sugarcane if they had indeed increased sugarcane production. In both cases, we compare the observed or factual outcome with the outcome that would have occurred otherwise (i.e. the counterfactual).

To formalize this problem, we use the standard framework used in evaluation analysis or the potential outcome approach (Roy 1951; Rubin 1974). The unit of analysis, $\boldsymbol{i}$, is the municipality. The treatment status is represented by a binary treatment indicator $\boldsymbol{D_i}$, which equals one if the municipality receives treatment and zero otherwise. In this study, treatment is defined in terms of sugarcane expansion: municipalities that expanded sugarcane at a certain rate will be considered treated, while municipalities that didn't expand sugarcane production are categorized as control municipalities[14]. The potential outcomes, measured in GDP per capita growth, are then denoted as $\boldsymbol{Y_{1i}}$ if the municipality expanded sugarcane production (treated) and as $\boldsymbol{Y_{0i}}$ if the municipality didn't expand sugarcane production (control).

The causal effect of treatment (sugarcane expansion) in a certain municipality is given by the difference in the potential outcomes with and without treatment, $\boldsymbol{Y_{1i} - Y_{0i}}$. Since a municipality will either expand sugarcane production ($\boldsymbol{D_i=1}$) or not expand sugarcane production ($\boldsymbol{D_i=0}$), 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 we cannot compute the individual treatment effect. We can however estimate the average effect, which compares the average outcomes of the treated and non-treated groups.

In this study, we are interested in estimating the average treatment effect on the treated (ATT) and the average treatment effect on the untreated (ATU). The ATT will evaluate what would have happened to the outcome variable of the sugarcane-expanding municipalities if they hadn't expanded sugarcane. The ATU will indicate what would have been the impact on the outcome variable of sugarcane non-expanding municipalities if they had expanded sugarcane. The ATT is then defined as the difference between expected outcome values with and without treatment for those municipalities that actually participated in the treatment (Heckman 1997):

$$ATT \equiv E(Y_1 - Y_0 | D = 1) = E(Y_1 | D = 1) - E(Y_0 | D = 1) \qquad (3.1)$$

[14] The classification into treatment or control group is described in detail in the next section

The ATU is the difference between expected outcome values with and without treatment for those municipalities classified in the control group:

$$(3.2) \qquad ATU \equiv E(Y_1 - Y_0 | D=0) = E(Y_1 | D=0) - E(Y_0 | D=0)$$

Both the ATT and ATU are composed of a factual and a counterfactual component. In the ATT, the counterfactual is $E(Y_0|D=1)$, while for ATU the counterfactual is $E(Y_1|D=0)$. Since we are dealing with a non-experimental design, we cannot observe the counterfactuals but will have to estimate them.

In this study, the counterfactuals are estimated using estimators based on the propensity score. We choose this class of estimators, first developed by Rosenbaum and Rubin (1983a), because they reduce the bias in treatment-effect estimates when the sample is not random. These estimators hence account for the fact that different municipalities have a different probability to expand sugarcane production. Indeed, one can expect that municipalities with favorable soil and climate conditions to grow sugarcane will have a higher probability to expand sugarcane than others. The propensity score is then defined as the probability that a municipality is treated (i.e. expands sugarcane production), given a set of observable control characteristics, X, or:

$$(3.3) \qquad \rho(X) = Prob(D=1|X).$$

In order to use propensity-based estimators, two assumptions need to be satisfied. The first assumption[15] or the 'conditional independence assumption (CIA)' (Lechner 1999), states that once we control for this set of observable characteristics X, the

[15] This condition is also known as the 'unconfoundedness' assumption (Rosenbaum and Rubin 1983a), 'selection on observables' (Heckman and Robb 1985)

systematic differences in outcomes between treated and comparison municipalities are entirely attributable to treatment. This assumption hence implies that observable covariates exhaustively determine selection into treatment. Since we condition on a rich set of variables and since the treatment in this study, namely sugarcane expansion, is mainly determined by the suitability of the land to grow sugarcane, a clearly observable characteristic, the CIA is considered to be satisfied. The CIA for ATT and ATU can be formalized as $Y_0 \perp D \mid X$ and $Y_1 \perp D \mid X$, respectively[16].

The second assumption is related to the joint distribution of treatments and covariates. This condition is known as the 'common support condition' or 'overlap condition' and prevents a situation of perfect predictability of D given X. As a result, the outcomes obtained by those municipalities from both groups that belong to this common support will be comparable. When estimating the ATT, the common support condition ensures that there are for each treated municipality control municipalities with the same X values (Heckman, LaLonde, and Smith 1999). Conversely, when estimating the ATU, the common support condition assures that for each control municipality a treated municipality can be found with the same X values. The common support conditions for ATT and ATU are represented as $Prob(D=1 \mid X) < 1$ and $0 < Prob(D=1 \mid X)$, respectively[17].

There are two steps involved when using the propensity score to estimate ATT and ATU. First, the propensity score needs to be estimated. Then, the different estimators based on the propensity score are constructed and the ATT and ATU are estimated.

---

[16] Note that the CIA for ATT and ATU are weakened versions of the CIA when estimating the average treatment effect or ATE. The ATE is defined as $E[Y_1 - Y_0 \mid X]$ and is thus the sub-sample-size-weighted average of ATT and ATU. The CIA for ATE is $Y_0, Y_1 \perp D \mid X$.

[17] These are again weakened versions of the common support condition for ATE, which is $0 < Prob(D=1 \mid X) < 1$.

Recall that when estimating the propensity score, one is in fact estimating the conditional probability that a municipality experiences growth in cane production, given the set of 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.

The logit model is specified using the stratification technique proposed by Dehejia and Wahba (1999; 2002). With this technique, a parsimonious model is specified to estimate the propensity score. Then, the sample is divided into several strata (or blocks) 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[18]. If there remains a statistical difference between the estimated propensity scores of the treatment and control group within a stratum, the stratum is divided in half and the average propensity scores are compared again. Consequently, the balance of the covariates within each stratum is tested. That is, using t-tests within each block it is checked whether the mean values for each covariate are the same between the treatment and control group. If there is no balance in a certain block, higher-order and interaction terms are added in the logit model specification until such differences no longer emerge.

[18] 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) show 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.

Once the propensity score is estimated, we construct four different types of estimators to estimate the ATT and ATU. In order to ensure that our results are not due to selecting one or the other of these estimators, we instead present results from all four to limit (insofar as possible) any potential bias from this methodological choice. The first two estimators are the blocking estimator and the reweighting estimator. The remaining two estimators are so-called “mixed estimators”; they are a combination of one of the above-mentioned estimators with regression. The motivation behind using several methods to compose propensity score-based estimators is twofold: i) to ensure that the results are robust and ii) to compare the relative performance of the estimators. Indeed, each method comes with its strengths and limitations and there is no consensus on which method is more effective. If the signs and values of the point estimates are similar for the different estimators, this will be an indication that the results are robust. In addition, comparing the mean squared errors (MSE) of the estimators will indicate which estimator performs better.

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

In propensity score reweighting, the estimated propensity score is used to reweight the observations in order to make the distributions of the control and treated group more similar. When estimating the ATT, the weighting estimator weights the observations in the control group, while the ATU weighting estimator weights the observations in the treated group. The weights are different for the ATT and ATU weighting estimator. In this study, the weights used are the ones proposed by Johnston and DiNardo (1996) and Imbens (2004), which are most commonly used in empirical studies. These reweighting estimators assure that the sum of the weights add up to the sample size *n*. The weighting functions of the ATT and ATU weighting estimators are represented in equations (3.4) and (3.5), respectively.

$$\frac{\hat{\rho}(X_j)}{1-\hat{\rho}(X_j)} \Bigg/ \frac{1}{n_0}\sum_{k=1}^{n}\frac{(1-D_k)\hat{\rho}(X_k)}{1-\hat{\rho}(X_k)} \tag{3.4}$$

$$\frac{1-\hat{\rho}(X_j)}{\hat{\rho}(X_j)} \Bigg/ \frac{1}{n_1}\sum_{k=1}^{n}\frac{D_k(1-\hat{\rho}(X_k))}{\hat{\rho}(X_k)} \tag{3.5}$$

In these equations, $\hat{\rho}(X_j)$ is the estimated propensity score, n is the size of the entire sample, n_0 is the size of the control group, and n_1 is the size of the treated group.

In the mixed methods the blocking estimator and respectively the reweighting estimator are combined with regression. These mixed methods are becoming increasingly popular because, although one method alone can be sufficient to obtain a consistent or even efficient estimator, combining the estimator with regression can improve precision and reduce the bias (Imbens and Wooldridge 2007). Rosenbaum and Rubin (1983b) first suggested combining the blocking estimator with regression through applying a least squares regression within the blocks. Robins et al. (1995) developed the so-called “doubly-robust” estimator, which is a regression adjustment of the reweighting estimator. The authors show that this estimator is consistent as long as either the propensity score or the regression function is specified correctly.

3.3 Data

3.3.1 Area and period of analysis

The impact of sugarcane expansion is studied on three different levels. First, the effects are analyzed at the national aggregate level (BR). Second, the impacts are examined in the two sugarcane growing regions: the North-Northeast (NE) and the Center-South (CS). Third, an additional distinction in the Center-South is made between the state of São Paulo (SP) and the remainder of the region, i.e. the Center-South excluding São Paulo (CSex). There are hence five different regions of analysis that will be considered. BR is composed of 27 states, NE consists of 10 states and CS is also composed of 10 states. Appendix Figure 1 lists the states in each of these regions and shows their geographical location.

The period of analysis considered in this study is 2001 until 2007. This period was chosen because sugarcane expansion in Brazil remained stable until 2001, but demonstrated a sharp increase from 2001 onwards. This trend is clearly illustrated in Figure 3.1, which shows the growth in planted area of sugarcane in Brazil since 1996. Note that in this study, we always use three-year moving averages for agricultural production data in order to smooth yearly fluctuations.

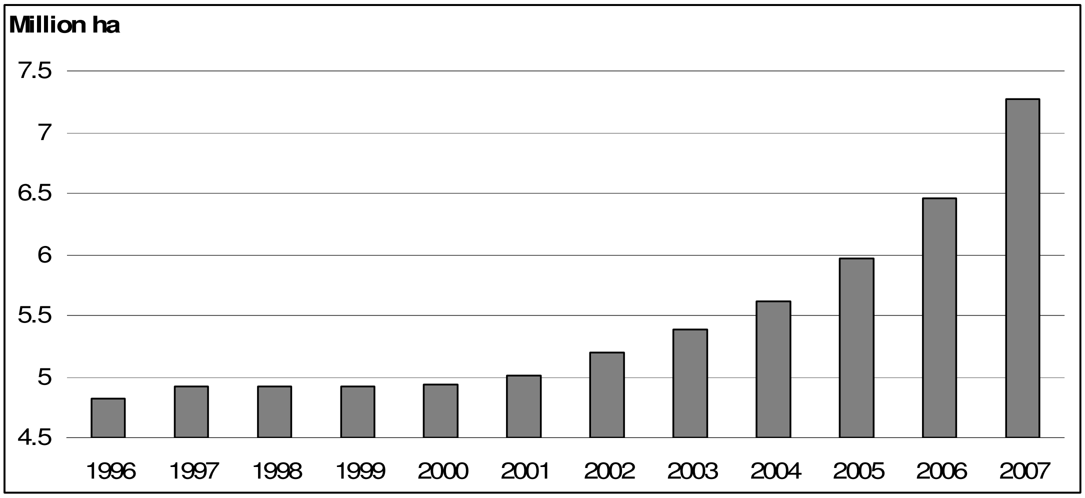

Source: IBGE

Figure 3.1: Area of sugarcane harvested in Brazil, 1996-2007

Most of the increase in sugarcane harvested area occurred in the Center-South of the country, which is also the region where most of the sugarcane plantations can be found. Table 3.1 demonstrates how much sugarcane area was harvested during the period 2001-2007 in Brazil and in the different regions of analysis. This table also summarizes each region's relative share in the national sugarcane production area. In the NE, sugarcane's production area increased by 100,000 ha between 2001 and 2007. The NE region's share in the national sugarcane area harvested however declined from 21.8% to 16.4% during that same period. The amount of land devoted to sugarcane production between 2001 and 2007 almost doubled in both subregions of the CS: in CSex it rose from 1.3 to 2.1 million ha, and in the state of São Paulo from 2.6 to almost 4 million ha. The growing importance of sugarcane production in the Center-South was evident in both CSex and SP: their shares in national sugarcane area harvested rose from 25.9% to 28.4% and from 52.1% to 55.0%, respectively

Table 3.1: Area of sugarcane harvested in different regions of analysis, 2001-2007 (1000 hectares and as share of national total)

	2001	2002	2003	2004	2005	2006	2007
Brazil	5010	5189	5382	5622	5959	6461	7270
North-Northeast (NE)	1091	1104	1119	1129	1132	1150	1190
	(21.8%)	(21.3%)	(20.8%)	(20.1%)	(19.0%)	(17.8%)	(16.4%)
Center-South (CS)	3907	4073	4251	4480	4811	5294	6061
	(78.0%)	(78.5%)	(79.0%)	(79.7%)	(80.7%)	(81.9%)	(83.4%)
Center-South excl. São Paulo (CSex)	1295	1354	1432	1517	1619	1776	2062
	(25.9%)	(26.1%)	(26.6%)	(27.0%)	(27.2%)	(27.5%)	(28.4%)
São Paulo (SP)	2612	2719	2820	2963	3193	3518	3999
	(52.1%)	(52.4%)	(52.4%)	(52.7%)	(53.6%)	(54.4%)	(55.0%)

Source: IBGE

Note: NE and CS do not add to total of Brazil because Brazil total also includes states not belonging to NE or CS

3.3.2 Definition of treatment

Each of the five regions of analysis demonstrates a different annual growth rate in the area of sugarcane harvested between 2001 and 2007. Figure 3.2 illustrates these annual growth rates and the average annual growth rate for the entire period for each of the regions of analysis. The North-Northeast (NE) had the lowest annual growth rates compared to the other regions, with an average annual growth rate from 2001 to 2007 of 1.5%. The annual growth rates in the other regions show a sharp increase over the period considered, with the highest percentage growth found in CSex (16.1% growth between 2006 and 2007). Combining this figure with Figure 3.2 hence points out that even though São Paulo is the largest contributor to growth in sugarcane expansion in absolute terms over the period 2001-07 (55% of Brazil's sugarcane is harvested in SP), sugarcane expanded at the highest rate in the other states of the Center-South region (CSex).

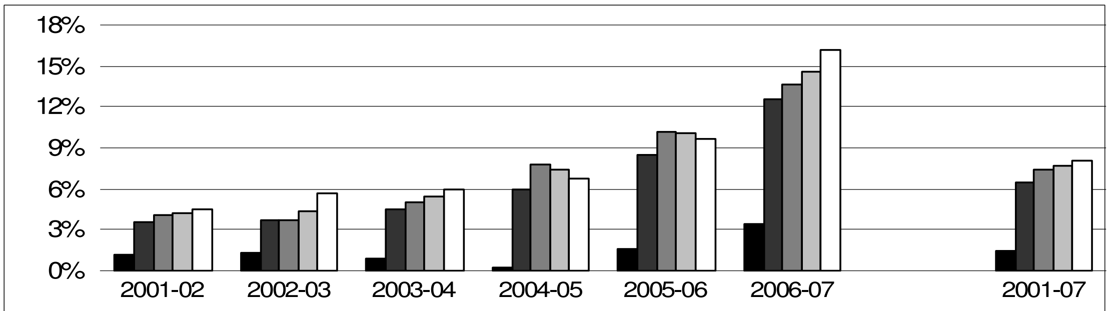

	2001-02	2002-03	2003-04	2004-05	2005-06	2006-07	2001-07
NE	1.2%	1.3%	0.9%	0.3%	1.6%	3.4%	1.5%
Brazil	3.6%	3.7%	4.5%	6.0%	8.4%	12.5%	6.5%
SP	4.1%	3.7%	5.1%	7.8%	10.2%	13.7%	7.4%
CS	4.3%	4.4%	5.4%	7.4%	10.0%	14.5%	7.7%
Csex	4.6%	5.7%	5.9%	6.7%	9.7%	16.1%	8.1%

Figure 3.2: Annual growth rates in sugarcane harvested area by region of analysis, 2001-2007

The average annual growth rates for each region over the period 2001-2007 were used to classify the municipalities into the treatment group. Municipalities with an average annual growth rate equal to or above the regional average were categorized in the treatment group. Municipalities with no or negative sugarcane expansion between 2001 and 2007 were categorized in the control group. Municipalities for which the average annual growth rate was above zero but strictly below the regional average were excluded from the analysis. This implies that depending on the region of analysis that is considered, the cutoff value that determines whether municipalities are classified in the treatment group varies. It is expected that excluding the municipalities in the middle of the spectrum of no-growth to high-growth will sharpen the ability of the statistical analysis to distinguish the differences between the high and low growth areas.

Table 3.2 shows for each region of analysis how many municipalities were classified in the control and treatment group and how many municipalities were excluded from the analysis. The last column in this table corresponds to the last column in Figure 3.2. This column lists the average annual growth rates in sugarcane area harvested. When calculating the ATT we will use these cutoff values to classify municipalities into the treatment group.

Table 3.2: Composition of control and treatment group by region of analysis

	control group Growth sugarcane harvested ≤ 0% (amount of municipalities)	**treatment group** Growth sugarcane harvested ≥ cutoff	**excluded from analysis** Growth sugarcane harvested strictly between 0 and cutoff	cutoff value for treatment (%)
Brazil	3158	1428	975	6.5%
NE	1336	539	56	1.5%
CS	1634	910	775	7.7%
CSex	1421	635	618	8.1%
SP	213	255	177	7.4%

Source: IBGE

Note: NE and CS do not add to total of Brazil because Brazil total also includes states not belonging to NE or CS

3.3.3 Control variables

We use a rich set of covariates at the municipal level to construct the propensity-score. As mentioned in the methodology section above, these covariates should be selected carefully: they should influence the treatment status but not be influenced by the treatment status. Table 3.3 lists the selected covariates with their respective descriptions and sources. Most of the variables are collected by the Instituto Brasileiro de Geografia e Estatística (IBGE), the remaining variables are published by the

Programa de las Naciones Unidas para el Desarrollo (PNUD) and the Empresa Brasileira de Pesquisa Agropecuária (EMBRAPA)[19]. We use both variables that are time-invariant and variables that vary over time.

Table 3.3: Control variables: abbreviation, description and source

Variable	Description	Source
area	Area municipality (km^2)	IBGE
high	Share of municipal area with high suitability for sugarcane production (%)	EMBRAPA
med	Share of municipal area with medium suitability for sugarcane production (%)	EMBRAPA
low	Share of municipal area with low suitability for sugarcane production (%)	EMBRAPA
sugharv	Sugarcane harvested, average 1990-92 (ha)	IBGE
sug_totharv	Share sugarcane harvested in total area of temporary crops harvested, average 1990-92 (%)	IBGE
totharv	Share temporary crops harvested in total area municipality, average 1990-92 (%)	IBGE
pasture	Share pastureland in total area municipality, 1996 (%)	IBGE
rented	Share of municipal area that is rented out for agricultural activities, 1996 (%)	IBGE
occupied	Share of municipal area that is occupied for agricultural activities, 1996 (%)	IBGE
partner	Share of municipal area that is used in partnerships for agricultural activities, 1996 (%)	IBGE
owned	Share of municipal area that is owned for agricultural activities, 1996 (%)	IBGE
popdens	Population density, 1991 (persons/(km^2)	IBGE
rurpop	Share of rural population in total population, 1991 (%)	IBGE
metrop	Metropolitan city (dummy)	IBGE
idhm	Municipal human development index, 1991	PNUD
gdppc80	GDP per capita, 1980 (2000 prices)	IBGE
gdppc96	GDP per capita, 1996 (2000 prices)	IBGE

The time-invariant characteristics are the area of the municipality and the share of municipal area suitable for sugarcane production. Indeed, one of the main determinants of sugarcane growth is the suitability of the land to grow sugarcane. In September 2009, EMBRAPA published the results of the agro-ecological zoning project (Manzatto et al. 2009). In this project, the area in each municipality was classified according to its suitability to grow sugarcane[20]. The main indicators used to

[19] The Instituto Brasileiro de Geografia e Estatística (IBGE) is the Brazilian Institute of Geography and Statistics, the Programa de las Naciones Unidas para el Desarrollo (PNUD) is the United Nations Development Programme and the Empresa Brasileira de Pesquisa Agropecuária (EMBRAPA) is the Brazilian Agricultural Research Corporation.
[20] The project excluded the following areas: 1- lands with an inclination higher than 12%; 2- areas with native vegetation; 3- the Amazon rainforest and the swamplands; 4- areas under environmental

define this suitability were the vulnerability of the land, the climate risk, the potential for sustainable agricultural production and the current environmental legislation. There were three different categories: high, medium and low suitability.

The variables that vary over time include agricultural, population and socio-economic characteristics. To assure that the variables that vary over time aren't influenced by the treatment status, we use lagged values of these variables. Wherever possible, we constructed a 3-year average to eliminate the influence of strong yearly fluctuations in agricultural production. We hence use a 3-year average for the period 1990-1992 for the following agricultural variables: sugarcane harvested area, share of sugarcane harvested area in total area of temporary crops harvested, and share of temporary crops harvested in the total area of the municipality. The other agricultural variables were collected during the agricultural census of 1996. These variables are: share of total pastureland in the total area of the municipality, and share of municipal area that is rented, occupied, used in partnerships or owned for agricultural activities. The population characteristics are population density and share of rural population in total population. The socio-economic characteristics include the municipal human development index in 1991 and GDP per capita in 1980 and 1996[21].

3.3.4 Outcome variable

The outcome variable is economic growth measured as average annual GDP per capita growth between 2001 and 2007. Table 3.4 presents the average value of the outcome

---

protection; 5- indigenous lands; 6- other remaining forests; 7- dunes; 8- mangroves; 9- rock formations; 10- reforested areas; and 11- urban areas and mineral areas. Note that these data are available for 20 out of the 27 states of Brazil. The treatment cutoff value for BR remains the same if only these 20 states are considered, namely 6.5%.

[21] No data on GDP is available for the beginning of the '90s due to the hyperinflation in the early 1990s.

variable for the treatment and control group by region of analysis before doing estimations based on the propensity score. The difference in the average outcome between treatment and control group is also displayed as well as the t-statistic of this difference. From this table, one would conclude that in BR, CS and SP the average annual GDP per capita growth was higher in the control group than in the treatment group, while in NE and CSex the converse was true. However, the difference is only statistically significant in the regions CSex and SP. The reason we don't pick up any statistical significant result in the region CS most probably is because: i) the values of the average growth rates are much higher in CSex than in SP and ii) in CSex the treatment has a higher average annual GDP per capita growth, while in the SP the control group has a higher growth rate than the treatment group. Aggregating the outcome values of CSex and SP in the region CS hence washes away these relative differences.

Table 3.4: GDP per capita growth in treatment and control group by region of analysis before estimations based on the propensity score

Region	Group	Obs	Mean	Std.Err.	Std.Dev.	[95% Conf.Interval]		t-stat
BR	Control	3079	4.884	0.116	6.409	4.657	5.110	
	Treatment	1375	4.838	0.166	6.156	4.513	5.164	
	Difference		0.045	0.202		-0.351	0.442	0.224
NE	Control	1319	4.838	0.121	4.387	4.601	5.075	
	Treatment	531	5.063	0.232	5.348	4.608	5.519	
	Difference		-0.225	0.262		-0.739	0.288	-0.861
CS	Control	1572	4.926	0.195	7.736	4.543	5.308	
	Treatment	864	4.459	0.224	6.596	4.018	4.899	
	Difference		0.467	0.297		-0.116	1.050	1.571
CSex	Control	1360	5.204	0.210	7.737	4.792	5.615	
	Treatment	597	6.066	0.257	6.284	5.561	6.572	
	Difference		-0.863	0.332		-1.514	-0.211	-2.599
SP	Control	212	3.140	0.515	7.502	2.124	4.156	

Table 3.4 (Continued)

Treatment	247	0.380	0.363	5.711	-0.336	1.095	
Difference		2.760	0.631		1.521	4.000	4.378

The summary statistics in Table 3.4 should be interpreted with great caution. GDP per capita growth rates are compared between two groups of municipalities that differ only in terms of sugarcane expansion. There are however many variables besides sugarcane expansion that influence GDP per capita growth and that could explain the difference in the value of the outcome variable between treatment and control group. The purpose of this study is exactly to take these other variables in consideration by performing estimations based on the propensity score.

3.4 Results

All estimates are obtained with STATA. We use the program *pscore*, developed by Becker and Ichino (2002), to estimate the propensity score with the stratification technique. We use the program *atts*, written by these same authors, to estimate the blocking estimator and obtain bootstrapped standard-errors, bias and confidence intervals. For the reweighting and mixed estimators, we construct bootstrapped standard errors using the technique described in Busso and Kline (2008).

3.4.1 Average treatment effect on the treated (ATT)

3.4.1.1 Estimation of the propensity score

The propensity score is estimated using the stratification technique. This technique was chosen because it assures that there is balance in the propensity scores and in the covariates within each stratum. The logit model used to estimate the propensity score is different for each region of analysis. This means that for each region of analysis, the propensity score is estimated using a different combination of linear, quadratic, square root and/or interaction terms of the covariates listed in Table 3.3. Appendix Table 3 through Appendix Table 7 display the selected full models for each region of analysis with values for the coefficients, standard errors, z-values and confidence intervals for all the covariates.

Table 3.5 summarizes, by region of analysis, the value of the pseudo R^2, the number of blocks that were needed to obtain balance, the amount of treated and control observations, and the region of common support. Note that the amount of control and treated observations considered in the final analysis is lower than in the original sample (compare Table 3.5 with Table 3.4). This is because no propensity score could be estimated for municipalities for which there were no data on one or more of the covariates. Consequently, these municipalities were excluded from the analysis. When analyzing the ATT, the region of common support is obtained as defined by Dehejia and Wahba (1999; 2002). To guarantee that each treated municipality can be compared with a control municipality with similar characteristics, these authors suggest that the common support should be imposed by eliminating all those observations in the control group that have an estimated propensity score lower than the lowest estimated propensity score in the treatment group.

Table 3.5: Summary of propensity score specification by region of analysis

region	pseudo R^2	blocks	No. of control	No. of treated	region of common support
BR	0.1448	8	1,738	890	[.02574491, .99761196]
NE	0.1509	6	738	373	[.06954797, .99995372]
CS	0.1679	6	888	578	[.04828420, .99934311]
CSex	0.1696	7	737	388	[.04985790, .99999674]
SP	0.3598	5	108	178	[.02523855, .99863938]

Figure 3.3 represents the histograms of the estimated propensity scores for both the control and treated group by region of analysis. The histograms show that the distribution of the estimated propensity scores for the control group differ from those for the treated group. In particular, for all regions of analysis, the control group has a relatively higher proportion of observations that have a low propensity score than the treated group. Recall that in this study the propensity score is defined as the probability that a municipality expands sugarcane production, given a set of observable characteristics. The histograms hence show that there are proportionally more municipalities that have a low probability to expand sugarcane production in the control group than in the treated group.

The key point of analyzing these histograms[22] is to make sure that there is sufficient overlap between the two groups, i.e. ensure that the common support condition is satisfied. Whereas the two groups obviously differ in their distribution of estimated propensity scores, these figures clearly show that the support of the estimated propensity scores nearly covers the entire unit interval. Therefore, there will clearly be municipalities in each group that can be matched with municipalities in the other group that are similar in all respects except for the extent to which they experienced growth in sugarcane production.

[22] In these histograms, no region of common support is imposed and hence no control observations are excluded.

BR

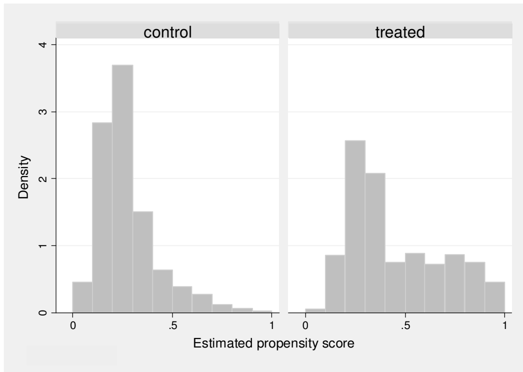

NE

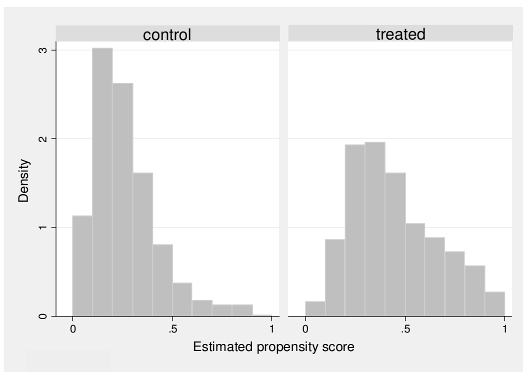

CS

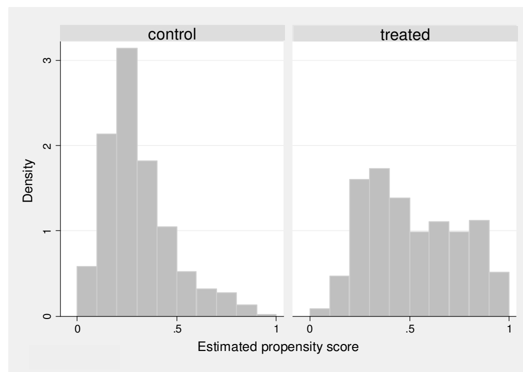

CSex

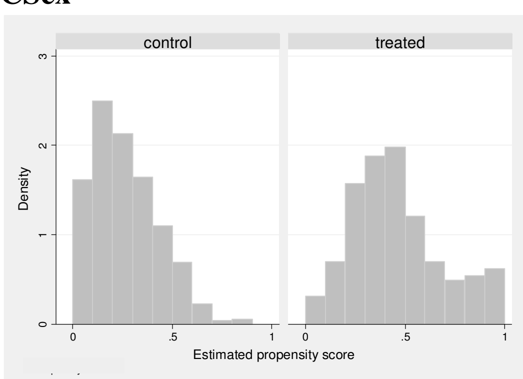

SP

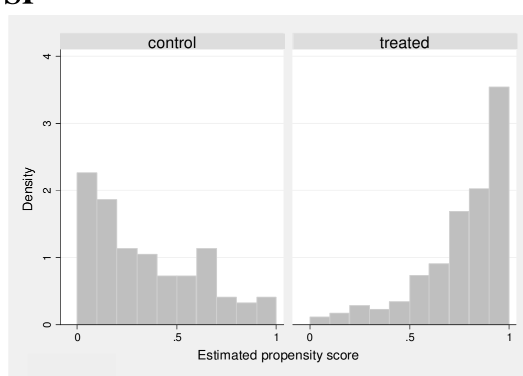

Figure 3.3: Histograms of estimated propensity scores for control group (left) and treated group (right) by region of analysis

3.4.1.2 Covariate balancing

Before estimating the ATT, we check whether the propensity score specification was able to balance the distribution of the relevant covariates in the treatment and control group. In this study, we use four techniques based on the estimated propensity score to combine the treated and control observations: two are based on blocking estimators and the remaining two are based on reweighting estimators. Since the blocking estimator follows immediately from the stratification technique, balance in the covariates is guaranteed. Indeed, the stratification technique is a valid balancing test for blocking estimators since the technique itself guarantees that balance in the covariates is obtained within each stratum.

For the reweighting estimators, we apply the technique suggested by Rosenbaum and Rubin (1985) to check for covariate balancing. In particular, we carry out two-sample t-tests: one before and one after reweighting based on the propensity score. Before reweighting, differences in the covariate means between the treatment and control groups are expected. After reweighting, there should be no statistically significant differences in the covariate means between the treated and control groups. Appendix Table 8 through Appendix Table 12 show that the selected propensity score specifications for all regions of analysis were able to balance the covariates. Appendix Table 8 shows that in BR, 14 out of 17 covariates weren't balanced before reweighting, while all covariates were balanced after reweighting. The following annex tables illustrate that similar results are obtained for the other regions: there were 14, 13, 7, and 7 covariates that didn't balance for respectively NE, CS, CSex and SP. After reweighting, all the covariates balance.

3.4.1.3 Estimation of the average treatment effect on the treated (ATT)

We use four different techniques to estimate the ATT, two of them are non-parametric techniques based on the propensity score and the other two are “mixed methods” or a combination of one of these propensity score-based techniques with regression. In particular, the four techniques used are: i) blocking on the propensity score; ii) propensity score reweighting; iii) “mixed blocking” or blocking and regression; and iv) “mixed reweighting” or weighting and regression.

Blocking on the propensity score follows immediately from the stratification technique. Once the sample is divided into different strata, the blocking estimator is the weighted average of the difference in outcomes in the treatment and control group, weighted by the proportion of treated observations in each stratum.

Propensity score reweighting aims at eliminating biases associated with differences in observed covariates. In this study, we use the weights as specified in equation 3.4 to reweight the observations in the control group. Appendix Figure 2 through Appendix Figure 6 illustrate, by region of analysis, how reweighting the observations in the control group makes the kernel densities of the estimated propensity scores of the treatment and control group more similar. These figures consider the region of common support and hence eliminate control observations not belonging to the region of common support. Each of these annex figures is composed of two panels: the left hand panel represents the kernel densities of the estimated propensity scores for the treatment and control group before reweighting, while the right hand panel represents the kernel densities of the estimated propensity scores for the treatment and the reweighted control group. It is clear in these figures that the kernel densities of the estimated propensity scores for the treated and control group are very different in the left hand panels, but have become much more similar in the right hand panels.

The mixed methods combine the above-mentioned techniques with regression. We select the subset of covariates to be included in the regression model using backward stepwise selection with inclusion and removal criteria set at 0.05 and 0.051, respectively[23]. The same subset is used for both the mixed blocking and mixed reweighting estimators. Note that the subset of covariates is different for each of the regions of analysis.

Table 3.6 presents, by region of analysis, the ATT estimates and summary statistics for each type of estimator. Bias, standard errors, t-values, and mean squared errors (MSE) for each of the estimators are obtained using bootstrap procedures with 10,000 replications. We applied non-parametric bootstrapping for the blocking estimators, while for the reweighting estimators and mixed estimators, we constructed bootstraps as described in Busso and Kline (2008).

In BR, the ATT estimates range between 0.452 and 0.537 percent. This implies that in Brazil municipalities that expanded sugarcane production between 2001 and 2007 experienced an average annual GDP per capita growth that was around 0.5% higher than in their sugarcane non-expanding counterparts. Over the entire 2001-2007 period,

[23] Changing the values of the inclusion and removal criteria resulted in the inclusion or exclusion of one or two other covariates. However, the final ATT estimates didn't change. We also used other techniques to select the subset of covariates to be included in the regression model: we carried out forward selection and backward elimination at different inclusion and removal criteria, and included all the covariates in the model that had a t-statistic above a certain cutoff value. In most instances, the same subset of covariates was selected, independent of the technique. If an alternative technique did suggest including or excluding one or more of the covariates, we found that the final ATT estimate gave similar results.

this annual average growth differential would result in an overall difference between the groups of around 3 percent.

When separating the ATT by the two sugarcane-producing regions, NE and CS, the ATT estimates are again positive, but only statistically significant in NE. The effects in NE are even stronger than in BR: ATT estimates take on values between 0.818 and 0.959 percent. This suggests that sugarcane expanding municipalities in NE on average experienced a 0.9 percent higher annual economic growth than sugarcane non-expanding municipalities. This would translate into an overall difference of 5.5 percent over the entire period 2001-2007.

The statistically insignificant effect of sugarcane expansion on economic growth in CS masks an important intra-regional difference. When detaching SP from the region and comparing the ATT estimates between SP and CSex, an interesting result appears. While there is no statistically significant effect in SP, the ATT estimates are statistically significant in CSex. These states experience a positive impact of sugarcane expansion on economic growth with values varying between 0.494 and 0.550 percent.

Table 3.6: ATT and summary statistics

region		blocking	reweighting	mixed blocking	mixed reweighting
BR	ATT	0.500	0.537	0.452	0.491
	std.err.	(0.231)**	(0.263)**	(0.239)*	(0.217)**
	Bias	0.001	-0.004	-0.005	-0.009
	MSE	0.053	0.069	0.057	0.047
NE	ATT	0.818	0.959	0.933	0.923
	std.err.	(0.335)**	(0.440)**	(0.330)***	(0.370)**
	Bias	-0.006	-0.008	0.007	-0.006
	MSE	0.112	0.194	0.109	0.137

Table 3.6 (Continued)

CS	ATT	0.020	0.099	0.136	0.143
	std.err.	(0.323)	(0.312)	(0.314)	(0.292)
	Bias	-0.002	0.018	0.006	-0.006
	MSE	0.104	0.098	0.099	0.085
CSex	ATT	0.498	0.494	0.550	0.508
	std.err.	(0.258)*	(0.255)*	(0.320)*	(0.262)*
	Bias	0.000	-0.005	0.027	0.007
	MSE	0.066	0.065	0.103	0.069
SP	ATT	0.250	0.666	0.407	0.617
	std.err.	(1.270)	(1.325)	(1.010)	(1.062)
	Bias	0.012	-0.078	-0.007	-0.099
	MSE	1.612	1.761	1.021	1.137

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

The motivation to use not only one technique but four different techniques to estimate the ATT is twofold: i) to see which technique performs better than the others, and ii) to check whether these techniques give similar results. Since we use in each region the same propensity score specification for all four estimators, we can use the MSE to compare the relative effectiveness of the estimators. It is clear from Table 3.6 that there is no single estimator that performs uniformly better in all regions. Only in SP, the mixed estimators outperform the other two estimators. Interestingly, all estimators give similar ATT estimates for the statistically significant results. This gives the estimates extra credibility since the choice of the type of estimator should only affect the ATT estimates marginally. In the next section, we perform extra robustness checks to further support these outcomes.

3.4.1.4 Robustness checks

Alternative treatment definitions

Section 3.3.2 describes how the classification of municipalities into the treated and control groups was based on their growth in sugarcane harvested over the period 2001-2007. For each region of analysis, the mean average growth rate in sugarcane harvested was calculated. Municipalities with a growth rate equal or greater than the regional average were classified in the treatment group, while municipalities with a zero or negative growth rate were classified in the control group. The remaining municipalities, i.e. with a positive growth rate below the regional average, were excluded from the analysis. Table 3.2 shows how these cutoff points vary by region of analysis. In this section, we investigate the sensitivity of the ATT estimates to these cutoff points. In particular, we consider three alternative thresholds and classify municipalities into the treated group if their growth in sugarcane harvested during 2001-2007 was at least 1 percent, 5 percent and 10 percent, respectively.

Table 3.7 through Table 3.11 report the ATT estimates and standard errors by region of analysis for the three alternative treatment thresholds. The amount of municipalities classified in the treated group decreases as the value of the treatment threshold increases, i.e. with a higher treatment cutoff, more municipalities get excluded from the analysis and less fall into the treated group. Note that the number of observations in the control group is not constant in each region of analysis. This is because the specification of the propensity score is slightly different for the different treatment definitions. Indeed, as the threshold value changes, the total sample size changes and hence the propensity score that satisfies the stratification technique specification changes as well.

Table 3.7: ATT for varying treatment thresholds – region BR

	blocking	reweighting	mixed blocking	mixed reweighting	*no. treated*	*no. control*
Treatment: sugarcane expansion ≥ 1%					*1373*	*1746*
ATT	0.493	0.554	0.394	0.423		
std.err.	(0.202)**	(0.236)**	(0.219)*	(0.210)**		
Treatment: sugarcane expansion ≥ 5%					*1003*	*1736*
ATT	0.458	0.541	0.369	0.489		
std.err.	(0.218)**	(0.244)**	(0.217)*	(0.205)**		
Treatment: sugarcane expansion ≥ 10%					*680*	*1708*
ATT	0.297	0.320	0.226	0.317		
std.err.	(0.254)	(0.259)	(0.251)	(0.233)		

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

In BR, the original cutoff value for classification into treatment was 6.5%. This implies that the sample size is now larger for the 1% and 5% thresholds than in the original analysis, while for the 10% threshold, the sample size is smaller. Comparing Table 3.7 with Table 3.6 shows that the ATT estimates and statistical significance levels are comparable between the 1% and 5% cutoffs and the original cutoff of 6.5%. However, when focusing on those municipalities in Brazil that experienced high sugarcane growth, i.e. at least 10%, the ATT estimates are smaller and the results are no longer statistically significant. This can be explained by the fact that the largest growths in sugarcane expansion occurred in the state of São Paulo. Hence, the treated sample considered with the 10% threshold definition will contain mostly municipalities from SP, for which is shown in the previous section that they have experienced no impact of sugarcane expansion on economic growth.

Table 3.8: ATT for varying treatment thresholds – region NE

	blocking	reweighting	mixed blocking	mixed reweighting	*no. treated*	*no. control*
Treatment: sugarcane expansion ≥ 1%					*382*	*743*
ATT	0.777	0.967	0.879	0.936		
std.err.	(0.333)**	(0.411)**	(0.321)***	(0.348)***		
Treatment: sugarcane expansion ≥ 5%					*262*	*746*
ATT	0.823	0.763	0.796	0.849		
std.err.	(0.367)**	(0.416)*	(0.408)*	(0.387)**		
Treatment: sugarcane expansion ≥ 10%					*168*	*759*
ATT	0.926	0.672	1.105	0.826		
std.err.	(0.431)**	(0.491)	(0.515)**	(0.449)*		

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

The average annual growth in sugarcane production in the region NE was 1.5% during the period 2001-2007. There is hence, as expected, no big difference in ATT estimates and significance levels between the original cutoff and the 1% cutoff (Table 3.8). Also at the 5% and 10% threshold levels, the ATT estimates remain comparable with the original estimates. Depending on the estimator and threshold level, ATT estimates average around 0.8 and 0.9 percent and remain statistically significant. Only in the case of the reweighting estimator at the 10% cutoff level, the ATT estimate is not statistically significant. This table shows that municipalities in NE that increased sugarcane expansion beyond 1% experienced as a result a higher economic growth of around 0.8 to 0.9 percent than sugarcane non-expanding municipalities.

Table 3.9: ATT for varying treatment thresholds – region CS

	blocking	reweighting	mixed blocking	mixed reweighting	*no. treated*	*no. control*
Treatment: sugarcane expansion ≥ 1%					*950*	*895*
ATT	-0.027	0.016	0.307	0.084		
std.err.	(0.313)	(0.294)	(0.310)	(0.272)		
Treatment: sugarcane expansion ≥ 5%					*709*	*893*
ATT	-0.012	0.110	0.103	0.140		
std.err.	(0.285)	(0.297)	(0.255)	(0.259)		
Treatment: sugarcane expansion ≥ 10%					*486*	*878*
ATT	0.006	0.042	0.093	0.137		
std.err.	(0.326)	(0.332)	(0.318)	(0.283)		

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Table 3.9 demonstrates that the ATT estimates in CS take on slightly different values at different thresholds. None of these values are statistically significant. This result coincides with the results obtained at the original cutoff of 7.7%, namely that in the region CS as a whole there is no significant impact of sugarcane expansion on economic growth.

Table 3.10: ATT for varying treatment thresholds – region CSex

	blocking	reweighting	mixed blocking	mixed reweighting	*no. treated*	*no. control*
Treatment: sugarcane expansion ≥ 1%					*679*	*779*
ATT	0.264	0.324	0.254	0.377		
std.err.	(0.265)	(0.244)	(0.260)	(0.256)		
Treatment: sugarcane expansion ≥ 5%					*471*	*804*
ATT	0.506	0.500	0.476	0.506		
std.err.	(0.241)**	(0.245)**	(0.223)**	(0.236)**		
Treatment: sugarcane expansion ≥ 10%					*330*	*734*
ATT	0.483	0.464	0.537	0.469		
std.err.	(0.287)*	(0.252)*	(0.250)**	(0.254)*		

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

The original cutoff value in CSex was the highest among the regions of analysis, 8.1%. Table 3.10 shows that municipalities in CSex need to have more than a modest growth in sugarcane production in order to experience impacts on economic growth. Indeed, at a threshold of 1%, the ATT estimates are not statistically significant. At the higher threshold levels of 5% and 10%, the ATT estimates have similar values as the ones at the 8.1% threshold, and are all statistically significant.

Table 3.11: ATT for varying treatment thresholds – region SP

	blocking	reweighting	mixed blocking	mixed reweighting	*no. treated*	*no. control*
Treatment: sugarcane expansion ≥ 1%					*271*	*116*
ATT	0.171	1.071	0.537	0.913		
std.err.	(0.834)	(1.094)	(1.366)	(0.922)		
Treatment: sugarcane expansion ≥ 5%					*208*	*108*
ATT	0.404	0.763	1.258	0.920		
std.err.	(1.154)	(1.203)	(1.278)	(1.089)		
Treatment: sugarcane expansion ≥ 10%					*156*	*109*
ATT	0.750	0.604	0.471	0.369		
std.err.	(1.318)	(1.347)	(1.057)	(0.924)		

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

The alternative treatment thresholds confirm the original results in SP, namely that sugarcane expanding municipalities did not experience higher economic growth as a result of sugarcane expansion. Under the original treatment threshold of 7.4%, and with the other three thresholds in Table 3.11, none of the ATT estimates are statistically significant.

Alternative regions of common support

The region of common support is imposed according to the technique described by Dehejia and Wahba (1999, 2002). When estimating the ATT, these authors discard all the observations in the control group that have an estimated propensity score below the lowest estimated propensity score in the treated group. Todd (2008), however, notes that this rule might be too stringent as potentially good matches just outside the common support might be lost. We hence perform a second robustness test to ascertain that the original common support didn't exclude important observations at the boundaries. In particular, we select around 20 percent of the originally discarded control observations that have the highest estimated propensity scores and include them into the analysis.

Table 3.12 reports the ATT estimates and summary statistics when considering this slightly larger region of common support. When comparing the results in this table with those in Table 3.6, it is clear that the ATT estimates only change marginally and that the levels of statistical significance are comparable to the ones obtained with the original region of common support. In particular, the ATT estimates are statistically significant in BR, NE and CSex, with point estimates around 0.5 percent, 0.9 percent and 0.5 percent, respectively. In CS and SP on the other hand, there is no statistically significant impact of sugarcane expansion on economic growth. These robustness checks hence reinforce the main findings.

Table 3.12: ATT for alternative regions of common support

region		blocking	reweighting	mixed blocking	mixed reweighting	no. treated	no. control
BR	ATT	0.498	0.538	0.448	0.492	890	1744
	std.err.	(0.235)**	(0.264)**	(0.241)*	(0.217)**		
	Bias	0.000	-0.002	-0.010	-0.009		
	MSE	0.055	0.069	0.058	0.047		
NE	ATT	0.818	0.967	0.936	0.936	373	745
	std.err.	(0.337)**	(0.419)**	(0.333)***	(0.358)***		
	Bias	-0.008	-0.007	0.002	-0.007		
	MSE	0.114	0.176	0.111	0.128		
CS	ATT	0.024	0.116	0.141	0.170	578	894
	std.err.	(0.324)	(0.325)	(0.317)	(0.289)		
	Bias	0.002	-0.002	0.005	-0.011		
	MSE	0.105	0.105	0.101	0.084		
CSex	ATT	0.508	0.518	0.558	0.534	388	748
	std.err.	(0.259)**	(0.246)**	(0.308)*	(0.259)**		
	Bias	0.001	-0.003	0.023	0.002		
	MSE	0.067	0.061	0.096	0.067		
SP	ATT	0.257	0.859	0.436	0.921	178	111
	std.err.	(1.267)	(1.319)	(0.995)	(1.069)		
	Bias	-0.006	-0.085	0.014	-0.154		
	MSE	1.607	1.747	0.990	1.167		

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

3.4.2 Average treatment effect on the untreated (ATU)

For the region CSex, we also estimate the average treatment effect on the untreated. That is, we estimate what would have happened to economic growth in those municipalities that didn't expand sugarcane expansion during the period 2001-2007 if they had increased sugarcane production. The motivation behind this analysis is that most of the future sugarcane expansion in Brazil will occur in the region CSex. A retrospective examination of the potential impacts of sugarcane expansion could provide useful insights into the future potential impacts.

3.4.2.1 Estimation of the propensity score

We consider three different scenarios when estimating the ATU in CSex. Namely, we investigate how an increase in average annual sugarcane production at 1%, 5% or 10% would have impacted economic growth in those municipalities that didn't expand sugarcane production. We use the same four types of propensity-score based estimators we used to estimate the ATT: the blocking estimator, the reweighting estimator, the mixed blocking estimator and the mixed reweighting estimator. For each scenario and for each type of estimator, we use the same logit model to estimate the propensity score. Appendix Table 13 through Appendix Table 15 displays the full models for each of the three scenarios with values for the coefficients, standard errors, z-values and confidence intervals for all the variables. These logit models are specified using the stratification technique and hence for each of the three scenarios the estimated propensity scores and the mean covariates are balanced in each of the blocks.

Table 3.13 compares the pseudo R2, the amount of blocks, the amount of observations in the control and treated group, and the region of common support for each of the scenarios. For the 1% and 5% scenarios, 7 blocks were needed to obtain balance while in the 10% scenario 6 blocks were needed. The number of control observations remains relatively stable across the scenarios because the same control group is used for all three scenarios, namely municipalities for which the growth in sugarcane expansion over the period 2001-2007 is zero or negative. The number of treated observations, on the other hand, decreases as the growth cutoff value increases. In other words, for the first scenario (1% growth), the observations in the control group are compared with all those municipalities that expanded sugarcane at least 1 percent per year. For the third scenario (10% growth), the observations in the control group are compared to a smaller treated group, which is composed of those municipalities that increased sugarcane production at least 10 percent per year.

Table 3.13: Summary of propensity score specification by scenario – ATU estimation

Scenario	pseudo R^2	blocks	no. control	no. treated	region of common support
Scenario 1: sugarcane expansion ≥ 1%	0.1433	7	791	632	[.00005155, .91765401]
Scenario 2: sugarcane expansion ≥ 5%	0.1612	7	790	466	[.00000002, .90133682]
Scenario 3: sugarcane expansion ≥ 10%	0.1954	6	785	307	[.00000011, .90836808]

The region of common support is constructed differently when estimating the ATU than when estimating the ATT. When estimating the ATU, all observations in the control group are used while the observations in the treated group that fall outside the region of common support are excluded. Figure 3.4 represents the histogram of the estimated propensity scores for each of the three scenarios. These figures clearly demonstrate that there is sufficient overlap between the control and treatment group and hence show that the region of common support can be imposed.

Scenario 1: at least 1% growth

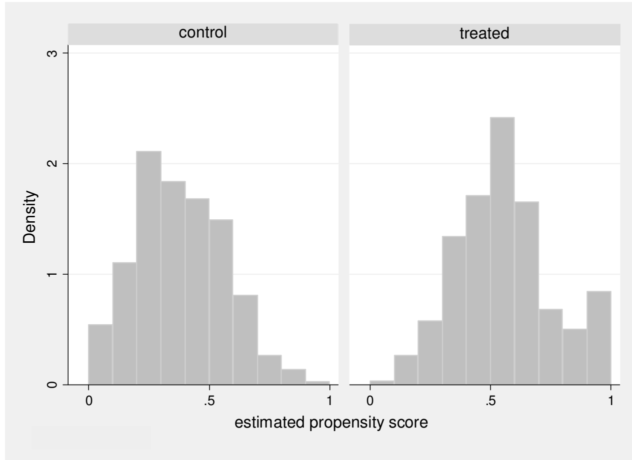

Scenario 2: at least 5% growth

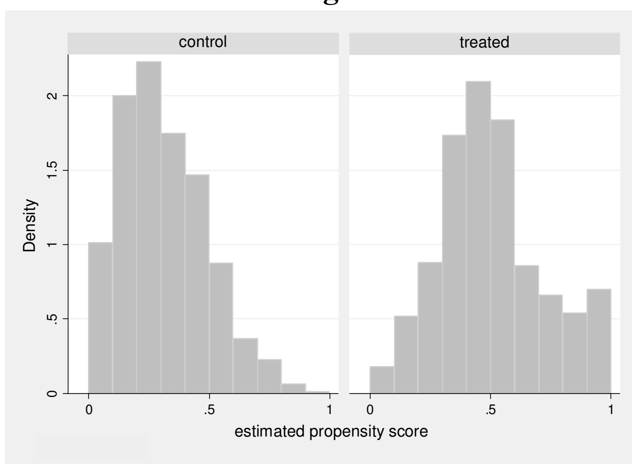

Scenario 3: at least 10% growth

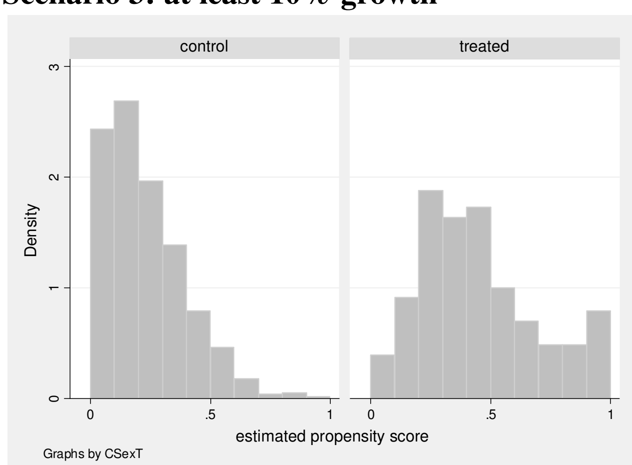

Figure 3.4: Histograms of estimated propensity scores for control group (left) and treated group (right) by scenario – ATU estimation

3.4.2.2 Covariate balancing

The stratification technique is a valid balancing technique when estimating ATU using the blocking estimator (Dehejia and Wahba 1999). In order to check for balance with the reweighting estimator, we compare the mean covariates between treated and control group before and after reweighting. Appendix Table 16 through Appendix Table 18 demonstrate that before reweighting, respectively 7, 9, and 9 out of the 15 covariates didn't balance for respectively the 1%, 5% and 10% scenario. After reweighting almost all covariates balance between the treated and control group.

Appendix Figure 7 through Appendix Figure 9 show for each of the scenarios how reweighting manages to make the distribution of the estimated propensity scores between the treated and control group more similar. The left hand panels in these figures show the distributions before reweighting and the right hand panels show the distributions after reweighting the observations in the treated group. In each scenario the group of control observations remains stable and hence the density distributions are the same. The group of treated observations changes for each scenario and hence the distributions differ by scenario. These figures demonstrate that for each scenario, reweighting made the distributions of the estimated propensity scores of treated and control group comparable. Note that these figures consider the region of common support as specified above and hence those observations in the treated group that don't belong to the common support are excluded.

3.4.2.3 Estimation of the average treatment effect on the untreated (ATU)

Table 3.14 represents the ATU estimates for each of the three scenarios and for each type of estimator. For scenario 1 (at least 1% growth in sugarcane production), the ATU estimates are positive and significant for each of the estimators and range from 0.609 to 0.733 percent. The ATU estimates are slightly lower for scenario 2 and for the blocking-based estimators the significance levels have decreased. In the third scenario, the ATU estimates are no longer statistically significant for most of the estimators. The relative performance of the estimators is analyzed by examining the MSE. The lowest MSE are obtained with blocking-based estimators. However, the difference in MSE is too small to draw conclusions on which estimator performs best in this analysis.

Table 3.14: ATU and summary statistics for CSex

	blocking	**reweighting**	**mixed blocking**	**mixed reweighting**
Scenario 1: sugarcane expansion ≥ 1%				
ATU	0.712	0.610	0.733	0.609
std.err.	(0.324)**	(0.334)*	(0.334)**	(0.328)*
bias	0.002	-0.024	-0.068	-0.036
MSE	0.105	0.112	0.116	0.109
Scenario 2: sugarcane expansion ≥ 5%				
ATU	0.597	0.588	0.611	0.606
std.err.	(0.334)*	(0.348)*	(0.320)*	(0.332)*
bias	-0.009	0.008	-0.003	-0.015
MSE	0.112	0.121	0.103	0.111
Scenario 3: sugarcane expansion ≥ 10%				
ATU	0.676	0.651	0.639	0.632
std.err.	(0.404)*	(0.415)	(0.415)	(0.403)
bias	0.010	0.002	-0.011	-0.005
MSE	0.163	0.172	0.172	0.163

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

The results in Table 3.14 suggest that if those municipalities that didn't expand sugarcane production over the period 2001-2007, had in fact expanded sugarcane production at annual rates of 1% or 5%, they would have experienced as a result a higher growth in GDP per capita of around 0.6 to 0.7, or 0.6 percent, respectively. However, if these municipalities had increased sugarcane production at an even higher rate, i.e. 10%, then the impact on economic growth would have become less straightforward because one out of the four estimators shows a significant effect.

3.4.2.4 Robustness checks

We perform two sets of robustness tests to check the results above. The first robustness checks enlarge the region of common support for each scenario to ensure that potentially important observations at the boundaries aren't excluded from the analysis. In particular, we include around 20 percent of the treated observations that were excluded in the original analysis and that had the highest estimated propensity scores. Table 3.15 reports the resulting ATU estimates and other summary statistics when considering a slightly larger region of common support for each of the scenarios. The ATU point estimates are very similar to the ones found in the original analysis. The significance levels in Table 3.15 are also comparable to those found in Table 3.14. The only difference with the original region of common support is that the blocking estimator for scenario 3 (at least 10% sugarcane expansion) is no longer significant with the larger region of common support. Increasing the region of common support hence reinforces the results obtained in Table 3.14, i.e. that sugarcane expansion rates of 1% or 5% would have caused higher economic growth in those municipalities that didn't expand sugarcane production, while a 10% expansion rate would have had no effect.

Table 3.15: ATU for alternative regions of common support

	blocking	reweighting	mixed blocking	mixed reweighting	no. treated	no. control
Scenario 1: sugarcane expansion ≥ 1%					*640*	*791*
ATU	0.710	0.629	0.726	0.628		
std.err.	(0.332)**	(0.325)*	(0.320)**	(0.319)**		
bias	-0.008	-0.002	-0.043	-0.011		
MSE	0.110	0.105	0.104	0.102		
Scenario 2: sugarcane expansion ≥ 5%					*473*	*790*
ATU	0.597	0.624	0.614	0.629		
std.err.	(0.331)*	(0.337)*	(0.319)*	(0.325)*		
bias	0.001	0.020	0.004	0.000		
MSE	0.110	0.114	0.101	0.106		
Scenario 3: sugarcane expansion ≥ 10%					*312*	*785*
ATU	0.676	0.660	0.639	0.642		
std.err.	(0.406)	(0.403)	(0.399)	(0.396)		
bias	0.008	-0.011	-0.010	-0.016		
MSE	0.166	0.163	0.159	0.157		

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

The second robustness check is to vary the cutoff points of the different scenarios. We consider four alternative thresholds for the growth in sugarcane production: growth that is i) strictly positive, ii) at least 2.5%, iii) at least 7.5%, and iv) at least 15%. The first alternative threshold will indicate whether a minimal increase in sugarcane production was already sufficient to induce an increase in economic growth. The three other thresholds are meant to solidify the original results. In particular, it is expected that at the 2.5% threshold, the results will fall in between those in scenario 1 (at least 1% growth) and scenario 2 (at least 5% growth). The results of the 7.5% threshold will indicate whether the 10% threshold is a turning point where no longer positive effects of sugarcane expansion can be seen or whether this effect already occurs at lower rates of sugarcane expansion. At the 15% threshold, which lies well above the threshold of scenario 3 (at least 10% growth), a finding of no statistically significant estimates is expected to reinforce the results for scenario 3. Note that these robustness checks are not intended to identify a clear threshold that guarantees a positive and significant impact of sugarcane expansion on economic growth but are conducted to reinforce the original results.

Table 3.16 represents the results of the second set of robustness checks. For the first alternative threshold, i.e. sugarcane expansion above zero percent, the ATU estimates are positive, but not statistically significant. At sugarcane expansion rates of at least 2.5% and 7.5%, the ATU estimates are around 0.6 and 0.7 percent and statistically significant. For the last alternative threshold of 15%, the ATU estimates are higher than in the previous situations, but no longer statistically significant.

Table 3.16: ATU for alternative thresholds of sugarcane expansion

	blocking	reweighting	mixed blocking	mixed reweighting	no. treated	no. control
i) Sugarcane expansion > 0%					*687*	*791*
ATU	0.468	0.413	0.445	0.443		
std.err.	(0.302)	(0.297)	(0.294)	(0.288)		
bias	-0.006	0.006	-0.021	-0.007		

Table 3.16 (Continued)

MSE	0.091	0.088	0.087	0.083		
ii) Sugarcane expansion ≥ 2.5%					*567*	*791*
ATU	0.759	0.630	0.700	0.676		
std.err.	(0.411)*	(0.368)*	(0.373)*	(0.357)*		
bias	-0.010	-0.009	-0.038	-0.024		
MSE	0.169	0.136	0.141	0.128		
iii) Sugarcane expansion ≥ 7.5%					*378*	*790*
ATU	0.698	0.732	0.641	0.688		
std.err.	(0.376)*	(0.401)*	(0.364)*	(0.380)*		
bias	-0.009	-0.017	-0.037	-0.019		
MSE	0.141	0.161	0.134	0.145		
iv) Sugarcane expansion ≥ 15%					*221*	*776*
ATU	0.882	1.227	0.943	1.187		
std.err.	(0.611)	(1.049)	(0.573)	(0.963)		
bias	0.021	0.036	-0.015	0.056		
MSE	0.374	1.101	0.328	0.930		

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

The second set of robustness checks hence suggests that expanding sugarcane production at low rates in those municipalities that didn’t expand sugarcane production has no significant impact on economic growth. Increasing sugarcane expansion above 1% per year but below 10% on the other hand would have resulted in higher GDP per capita growth.

3.5 Conclusion

In this chapter we analyze the causal relationship between sugarcane expansion and economic growth in Brazil and in the main sugarcane-producing regions in Brazil during the period 2001-2007. We examine this relationship by answering two types of counterfactual questions. First, we examine what would have happened to economic growth in sugarcane-expanding municipalities if they hadn't expanded sugarcane production. Second, we analyze what would have been the impact on economic growth in municipalities that didn't expand sugarcane if they had expanded sugarcane production.

We address these counterfactual questions using the evaluation methodology. Accordingly, we classify municipalities into treated and control groups based on their degree of sugarcane expansion and then compare the outcome variable, i.e. economic growth measured as GDP per capita growth, between these groups. The first counterfactual question is then answered by estimating the average treatment effect on the treated, or ATT. The second counterfactual question is assessed by estimating the average treatment effect on the untreated, or ATU. Both ATT and ATU are estimated using four estimators based on the propensity score. These estimators are the blocking estimator, the reweighting estimator, and the so-called "mixed estimators" that combine regression with one these two estimators.

We estimate the ATT in five different regions: Brazil (BR), the North-Northeast (NE), the Center-South (CS), São Paulo state (SP) and in the region comprised of the states in the Center-South excluding São Paulo state (CSex). The classification of municipalities into the treatment group varies by region of analysis. A municipality is categorized into the treatment group if it has an average annual growth in sugarcane production between 2001 and 2007 that is at least as high as the average annual growth in the region. In one of the robustness checks we vary this cutoff value for classification into the treated group. Municipalities with no or a negative average annual growth in sugarcane production are classified in the control group.

In BR, the ATT estimates indicate that those municipalities that expanded sugarcane production experienced a 0.5 percent higher average annual economic growth than their sugarcane non-expanding counterparts. These estimates are even higher in the NE, where the economic growth difference would have been 0.9 percent. There is no significant effect found in the CS. The results in NE and CS align with the conclusions of Burnquist et al. (2004), who show that the NE benefits more from a demand shock in sugar and ethanol exports than the CS.

The regional picture in CS, however, conceals an important intra-regional result. When estimating the ATT for the CSex, the estimates become significant and attain values around 0.5 percent. In SP the estimates remain statistically insignificant. This last result coincides with the results in Chapter 2 that sugarcane-expanding municipalities in São Paulo state didn’t experience higher economic growth than the sugarcane non-expanding municipalities.

The different sets of robustness checks reinforce the abovementioned results. Furthermore, one of the robustness checks reveals two additional results. In this robustness check, the threshold for classification into treatment is changed to fixed cutoff values at 1%, 5%, and 10%. In NE, the ATT estimates show that sugarcane expansion at any of the specified rates leads to a significant increase in economic growth. In CSex, however, the ATT estimates become insignificant when the group of treated municipalities is composed of all municipalities that experienced at least 1% growth. These estimates indicate that in CSex at least a growth of 5% in sugarcane production is needed to establish a significant causal relationship between sugarcane expansion and economic growth.

The ATU is estimated for CSex only because most of the future sugarcane expansion is planned in this region. Here we look at the control municipalities and examine what would have happened to their economic growth if they had expanded sugarcane production at respectively 1%, 5% or 10%. The ATU estimates are positive and significant at the 1% and 5% level, but no longer significant at the 10% level. The robustness checks confirm these results and show similar ATU estimates of 0.6 and 0.7 percent at 2.5% and 7.5% levels. Interestingly, the estimates become insignificant when the cutoff is lowered so that all municipalities with a strictly positive growth in sugarcane expansion are classified into the treatment group.

These results have several policy implications. First, even though NE is characterized by lower productivity and higher production costs compared to CS, this region still benefited from sugarcane expansion. This suggests that sugarcane, which has a long tradition in this region, is still a lucrative business which raises the economic growth in these municipalities. Second, the lack of any significant impact of sugarcane production on economic growth in São Paulo state in both this chapter and in chapter 2, suggests that future expansions most probably shouldn't be planned in this state. The reasons behind this result are not obvious from this study and deserve more attention in other empirical studies. Finally, the ATT and ATU estimates in CSex and the robustness checks of these estimates show that sugarcane expanding municipalities need to expand at above a certain threshold, i.e. higher than 1%, in order to reap economic growth benefits. Municipalities in CSex that didn't expand sugarcane would have experienced significant positive economic growth effects if they expanded sugarcane production at at least 1%, but not at too high rates because then the effects become no longer significant. The projected expansions in CSex thus have the potential to generate economic growth for the region. Caution should be applied though, so that sugarcane expansion isn't too large in those municipalities that originally didn't increase sugarcane production because it might disrupt current economic and agricultural activities that are beneficial to the region.

CHAPTER 4

THE SUGAR AND ETHANOL BOOM IN SÃO PAULO STATE AND ITS EFFECTS ON THE MAIN SECTORS OF THE ECONOMY

4.1 Introduction

The state of São Paulo in Brazil has experienced a sharp growth in sugarcane production since 2002. The drivers behind this increase were the rise in demand for both sugar and ethanol, two products that are derived from sugarcane. Consequently, it was assumed that the municipalities in São Paulo state that expanded sugarcane production would experience higher economic growth due to the creation of employment opportunities in the sugarcane, sugar and ethanol sectors. In Chapter 2, however, we show that the sugarcane expansion between 2002 and 2006 had no significant effect on growth in GDP per capita.

The aim of this study is to examine the underlying reasons behind these findings by focusing on sugarcane's impact on the different sectors of the economy. Whereas the outcome variable GDP per capita, which was used in both Chapter 2 and Chapter 3, is a comprehensive indicator of sugarcane's direct effect on the local economies, the results obtained for São Paulo state in these two chapters deserve further analysis. In this chapter, we hence focus on variables that give more insight into the different sectors of the economy and examine variables that contribute directly or indirectly to the original variable, i.e. GDP per capita growth. We hypothesize that the growth in sugarcane production affected those sectors of the economy that only have a relatively small importance on the total economy. An additional hypothesis that we explore is that sugarcane expansion had significant positive effects on some sectors of the economy but that these positive effects were offset by negative effects in other sectors.

To analyze these hypotheses, we consider different sets of outcome variables that can be disaggregated by economic sector. A first set of outcome variables is GDP by sector. The Brazilian Institute of Geography and Statistics recently updated its methodology to compute GDP (IBGE 2008). As a result, total GDP is now defined as the sum of four different components: agriculture Value Added (VA), industry VA, services VA and taxes. Accordingly, our first hypothesis translates into analyzing whether sugarcane expansion led to growth in the agriculture and industry components of total GDP, and whether the shares of these components in total GDP were too small to show significant results in the aggregate figure.

Sugarcane's impact on agriculture VA is not straightforward. Especially in terms of employment generation, it is not clear whether sugarcane expansion created additional jobs. On the one hand, it is evident that the extension of sugarcane plantations has brought forth more land to be cultivated and more sugarcane to be cut. As mentioned in Chapter 2, the amount of land devoted to sugarcane plantations in São Paulo state increased from 2.7 to 3.5 million hectares between 2002 and 2006. On the other hand however, the increased mechanization of sugarcane harvesting implies that less labor is required for these activities. Fredo et al. (2007) report that following the new mechanization law[24] in São Paulo state, 20 percent of sugarcane was harvested mechanically in 2002, while by 2006 this share increased to 30 percent. Furthermore, sugarcane expansion might have caused unemployment of agricultural workers engaged in other agricultural activities. Indeed, sugarcane expansion in São Paulo state

[24] Lei 11.241/2002

did not only occur on pastureland, but also on cropland that was cultivated with labor-intensive crops (Rudorff 2010).

The impact on the industrial sector is more likely to be positive. Sugarcane's end products, sugar and ethanol, are both produced in industrial mills. Between 2002 and 2006, the production of these two products increased substantially. São Paulo produced almost 20 percent more sugar (in tonnes) in 2006 compared to 2002, while ethanol production (in liters) increased 30 percent during that same time period (UNICA 2010). This upsurge in sugar and ethanol production is expected to translate into higher industry VA. The employment effects in the industry sector are assumed to be strictly positive since there are no substitution effects at play as in the agricultural sector.

Besides employment creation in the industry sector, there might be also positive employment effects in the construction and trade sector. The enlarged sugar and ethanol production led to the construction of over 30 mills in São Paulo state. In the harvest year 2000/2001, there were 133 mills in São Paulo state (Amaral and Neves 2003). This number rose to 169 mills by December 2006 (Ueki 2007). The trade sector also experienced an impressive boost. Between 2002 and 2006, ethanol exports from São Paulo grew from 112 billion US$ (FOB) to 1,210 billion US$ (FOB), while sugarcane exports more than tripled during the same time period, i.e. from 690 million to 2.6 billion US$ (FOB) (SECEX/MDIC 2010).

It is also important to analyze wages in the sugarcane-related sectors and compare them with those in other sectors. In fact, higher wages in sugarcane-related activities combined with growth in employment result in greater purchasing power. This in turn translates into increased demand for consumption goods and services and new employment opportunities in these sectors. Smeets et al. (2006), for example, estimate that the indirect and induced employment effects of sugarcane expansion in the late 1990s were around 940.000 and 1.800.000 jobs, respectively.

In this research, we investigate the different channels through which sugarcane expansion impacted the local economies of sugarcane-expanding municipalities. We examine the causal relation between sugarcane expansion and three sets of outcome variables, namely GDP, employment and wages. By looking at these outcome variables at the aggregate and sector level, we take into consideration that sugarcane expansion might have led to positive impacts in one sector (most likely the sectors directly related to sugarcane) but might have had negative repercussions on other sectors.

The methodology we use is similar to the one in Chapter 2. That is, we use propensity score-based estimators to assess the impact of sugarcane expansion on the set of outcome variables. We choose this methodology because it accounts for the fact that not all municipalities have the same propensity to expand sugarcane production. Comparing the outcome variables between municipalities that have the same propensity to expand sugarcane production but have a different actual growth in sugarcane production, allows us to draw causal relations between sugarcane expansion and the outcome variables.

In addition, we also examine the relative importance of the different sectors in the economy and explore how the employment and wages in the sugarcane, sugar and ethanol sector compare to those in the agricultural and industry sector. This additional analysis will help us interpret the estimates obtained with the propensity score-based estimators.

The rest of the chapter is organized as follows. In the next section we review the methodology on propensity-score based estimators. In section 3, we describe the data and variables. The empirical results are presented in section 4. We first analyze the relative importance of the sugarcane-related activities and of the agricultural and industrial sector in the economy. Consequently, we estimate the average treatment effect on the treated using estimators based on the propensity score. Section 5 concludes.

4.2 Methodology

We analyze whether municipalities in São Paulo state that expanded sugarcane production as a result experienced higher growth in GDP, income and employment generation in the different sectors of the economy. To identify the causal relation between sugarcane expansion and these outcome variables, we assess what the situation would have been like if no sugarcane expansion had taken place. In other words, we seek to answer the question: "What would have happened to the values of these outcome variables if the sugarcane-expanding municipalities had not expanded sugarcane?".

Since we are comparing a factual outcome with a counterfactual outcome, we formalize this problem with the standard framework used in evaluation analysis or the potential outcome approach (Roy 1951; Rubin 1974). The unit of analysis, $\boldsymbol{i}$, is the municipality. The treatment status is represented by a binary treatment indicator $\boldsymbol{D_i}$, which equals one if the municipality receives treatment and zero otherwise. In this study, treatment is defined in terms of sugarcane expansion: municipalities that expanded sugarcane at a certain rate will be considered treated, while municipalities that didn't expand sugarcane production are categorized as control municipalities[25]. The potential outcomes are then denoted as $\boldsymbol{Y_{1i}}$ if the municipality expanded sugarcane production (treated) and as $\boldsymbol{Y_{0i}}$ if the municipality didn't expand sugarcane production (control).

The causal effect of treatment (sugarcane expansion) in a certain municipality is given by the difference in the potential outcomes with and without treatment, $\boldsymbol{Y_{1i} - Y_{0i}}$. Since a municipality will either expand sugarcane production ($\boldsymbol{D_i=1}$) or not expand sugarcane production ($\boldsymbol{D_i=0}$), 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 we cannot compute the individual treatment effect. We can however estimate the average effect, which compares the average outcomes of the treated and non-treated groups.

In this study, we are interested in estimating the average treatment effect on the treated (ATT). The ATT will evaluate what would have happened to the outcome variables 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):

$$ATT \equiv E(Y_1 - Y_0 | D = 1) = E(Y_1 | D = 1) - E(Y_0 | D = 1) \tag{4.1}$$

[25] The classification into treatment or control group is described in detail in the next section

The first component in expression (1) is the factual component, while the second part, $E(Y_0|D=1)$, is the counterfactual. Since we are dealing with a non-experimental design, we cannot observe the counterfactuals but will have to estimate them.

In this study, the counterfactuals are estimated using estimators based on the propensity score. We choose this class of estimators, first developed by Rosenbaum and Rubin (1983a), because they reduce the bias in treatment-effect estimates when the sample is not random. These estimators hence account for the fact that different municipalities have a different probability to expand sugarcane production. Indeed, one can expect that municipalities with favorable soil and climate conditions to grow sugarcane will have a higher probability to expand sugarcane than others. The propensity score is then defined as the probability that a municipality is treated (i.e. expands sugarcane production), given a set of observable control characteristics, X, or:

$$\rho(X) = Prob(D=1|X). \tag{4.2}$$

In order to use propensity-based estimators, two assumptions need to be satisfied. The first assumption[26] or the ‘conditional independence assumption (CIA)’ (Lechner 1999), states that once we control for this set of observable characteristics X, the systematic differences in outcomes between treated and control municipalities are entirely attributable to treatment. This assumption hence implies that observable covariates exhaustively determine selection into treatment. Since we condition on a rich set of variables and since the treatment in this study, namely sugarcane expansion, is mainly determined by the suitability of the land to grow sugarcane, a clearly

[26] This condition is also known as the ‘unconfoundedness’ assumption (Rosenbaum and Rubin 1983a), ‘selection on observables’ (Heckman and Robb 1985)

observable characteristic, the CIA is considered to be satisfied. The CIA for ATT can be formalized as $\boldsymbol{Y_0 \perp D \mid X}$[27].

The second assumption is related to the joint distribution of treatments and covariates. This condition is known as the 'common support condition' or 'overlap condition' and prevents a situation of perfect predictability of $\boldsymbol{D}$ given X. As a result, the outcomes obtained by those municipalities from both groups that belong to this common support will be comparable. When estimating the ATT, the common support condition ensures that there are for each treated municipality control municipalities with the same X values (Heckman, LaLonde, and Smith 1999). The common support condition for ATT is represented as $\boldsymbol{Prob(D=1 \mid X) < 1}$[28].

There are two steps involved when using the propensity score to estimate ATT. First, the propensity score needs to be estimated. Then, the different estimators based on the propensity score are constructed and the ATT is estimated. Recall that when estimating the propensity score, one is in fact estimating the conditional probability that a municipality experiences growth in cane production, given the set of observable characteristics X. This is usually done by estimating a logit model, where the treatment status $\boldsymbol{D}$ 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

[27] Note that the CIA for ATT is a weakened version of the CIA when estimating the average treatment effect or ATE. The ATE is defined as $\boldsymbol{E[Y_1-Y_0 \mid X]}$ and the CIA for ATE is $\boldsymbol{Y_0, Y_1 \perp D \mid X}$.

[28] This is again a weakened version of the common support condition for ATE, which is $\boldsymbol{0 < Prob(D=1 \mid X) < 1}$.

contains pretreatment variables and time-invariant characteristics. It often also includes lagged values of the outcome variable.

The logit model is specified using the stratification technique proposed by Dehejia and Wahba (1999; 2002). With this technique, a parsimonious model is specified to estimate the propensity score. Then, the sample is divided into several strata (or blocks) 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[29]. If there remains a statistical difference between the estimated propensity scores of the treatment and control group within a stratum, the stratum is divided in half and the average propensity scores are compared again. Consequently, the balance of the covariates within each stratum is tested. That is, using t-tests within each block it is checked whether the mean values for each covariate are the same between the treatment and control group. If there is no balance in a certain block, higher-order and interaction terms are added in the logit model specification until such differences no longer emerge.

Once the propensity score is estimated, four different types of estimators are constructed to estimate the ATT. The first two estimators are the blocking estimator and the reweighting estimator. The remaining two estimators are so-called "mixed estimators"; they are a combination of one of the above-mentioned estimators with regression. In Chapter 2, we only used the blocking and reweighting estimators because Busso, McCrary and DiNardo (2008) show that these estimators perform best

[29] 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.

in small samples with n=100 or n=500. In this chapter, we decided to also construct two mixed estimators.

The motivation behind using several methods to compose propensity score-based estimators is twofold: i) to ensure that the results are robust and ii) to compare the relative performance of the estimators. Indeed, each method comes with its strengths and limitations and there is no consensus on which method is more effective. If the signs and values of the point estimates are similar for the different estimators, this will be an indication that the results are robust. In addition, comparing the mean squared errors (MSE) of the estimators will indicate which estimator performs better.

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

In propensity score reweighting, the estimated propensity score is used to reweight the observations in the control group in order to make the distributions of the control and treated group more similar. In this study, we use the weights proposed by Johnston and DiNardo (1996) and Imbens (2004), which are most commonly used in empirical studies. These reweighting estimators assure that the sum of the weights add up to the sample size *n*. The weighting function of the ATT weighting estimator is

$$\text{(4.3)} \qquad \frac{\hat{\rho}(X_j)}{1-\hat{\rho}(X_j)} \Bigg/ \frac{1}{n_0}\sum_{k=1}^{n}\frac{(1-D_k)\hat{\rho}(X_k)}{1-\hat{\rho}(X_k)}$$

In this equation, $\hat{\rho}(X_j)$ is the estimated propensity score, n is the size of the entire sample, and n_0 is the size of the control group.

In the mixed methods the blocking estimator and respectively the reweighting estimator are combined with regression. These mixed methods are becoming increasingly popular because, although one method alone can be sufficient to obtain a consistent or even efficient estimator, combining the estimator with regression can improve precision and reduce the bias (Imbens and Wooldridge 2007). Rosenbaum and Rubin (1983b) first suggested combining the blocking estimator with regression through applying a least squares regression within the blocks. Robins et al. (1995) developed the so-called "doubly-robust" estimator, which is a regression adjustment of the reweighting estimator. The authors show that this estimator is consistent as long as either the propensity score or the regression function is specified correctly.

4.3 Data

This study is motivated by the results in Chapter 2 and uses a similar methodological approach. The period of analysis hence extends from 2002 until 2006. The definition for classification of municipalities into treatment and control groups is again based on the average annual growth in sugarcane harvested between 2002 and 2006. Municipalities with a growth equal to or above the state's annual average of 6.8% are classified in the treated group. Municipalities with no growth or a negative growth are categorized in the control group. Consequently, municipalities that experienced a positive growth below 6.8% are excluded from the analysis. Table 4.1 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 robustness checks we will vary the cutoff point for classification into treatment.

Table 4.1: Composition of treated and control groups

	Amount of municipalities	Share of total
Treatment group Growth sugarcane harvested ≥ 6.8%	241	37.4%
Control group Growth sugarcane harvested ≤ 0%	236	36.6%
Removed from analysis Growth sugarcane harvested between 0% and 6.8%	168	26.0%
Total	**645**	**100%**

Source: IBGE

The control variables used to construct the propensity score remain the same as in Chapter 2. Table 4.2 lists these variables and their sources. The Instituto Brasileiro de Geografia e Estatística (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, sugar harvest, sugar harvest as a share of total harvest, and total harvest divided by area (sugharv, sugharv/totharv, and totharv/area) are collected on a yearly basis. We 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 population data and on municipal GDP per capita. Since no data on GDP is available for the beginning of the '90s due to the hyperinflation in that period, we used GDP per capita data for 1980 and 1996.

Table 4.2: 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 limitation (%)	Gov.SP
suitable_restr/area	Share of municipal area suitable for sugarcane production under environmental restriction (%)	Gov.SP

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 for growing 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[30]. We only used the first three variables since the fourth one, i.e. area not

[30] The four different categories are defined as follows. (1) areas suitable for sugarcane production: areas with favorable climatic conditions for the cultivation of sugarcane and without any specific environmental constraints; (2) areas suitable under environmental limitations: areas with favorable climate and soil for sugarcane cultivation but classified as Environmental Protection Areas (APA), or as medium priority areas for enhancing the connectivity, as directed by the BIOTA-FAPESP Project; or as critical watersheds; (3) suitable areas with environmental constraints: areas with favorable climatic conditions for the cultivation of sugarcane but classified as buffer zones of the Conservation Units of Integral Protection (UCPI), or as high priority areas for increased connectivity as indicated by the BIOTA-FAPESP Project, or as areas of high vulnerability for the groundwater in the State of São Paulo, as published by CETESB-IG-DAEE - 1997; (4) areas not suitable or inadequate areas: areas classified under the Conservation Units of Integral Protection (UCPI) at State and Federal level; areas classified as extremely important for biological conservation, indicated by the BIOTA-FAPESP Project for the creation of Conservation Units of Integral Protection (UCPI); areas classified as Zones Wildlife Areas Environmental Protection (EPA); areas with soil and climatic constraints to grow sugarcane; and areas with slopes steeper than 20%.

suitable for sugarcane production, can be derived from the other three and would lead to collinearity in the logit model.

There are three sets of outcome variables. The first set concerns GDP per capita data broken down by its several components: agriculture (value added), industry (value added), service (value added) and taxes. Since we are interested in growth rates, we are using GDP per capita data at constant 2000 prices, which are provided by the Instituto de Pesquisa Econômica Aplicada[31] (IPEA).

The second and third sets are composed of data on employment and wages, respectively. Both sets are disaggregated by the main sectors in the economy, namely agriculture, industry, services, civil construction and trade. These data are obtained from the Relação Anual de Informações Sociais (RAIS), which is an administrative dataset collected on an annual basis by the Ministry of Employment and Labor. The main limitation of this dataset is that it only reports formal labor[32]. Since the workers' mean wages are expressed in nominal values, we deflate these data using the National Consumers' Price Index for São Paulo (IPEA).

Table 4.3 gives an idea of the annual average growth of these outcome variables for the treatment and control group over the period 2002-2006. This table also displays t-statistics that determine whether the growth between the treatment and control group is significantly different. The control and treated groups are as defined in Table 4.1

The land in São Paulo state is classified as follows: 26% are suitable areas, 45% are suitable areas with environmental restrictions, 28% are suitable areas with environmental restrictions, and only 1% are inadequate areas

[31] Institute of Applied Economic Research

[32] The National Household Sample Research (or Pesquisa Nacional por Amostra de Domicílios, PNAD), collected by IBGE, also contains data on informal labor. However, the PNAD doesn't have data at the municipal level.

and hence are composed of 237 and 241 municipalities, respectively. Based on this table alone, one would conclude that the treated group experienced a significantly higher growth in industry GDP per capita, total employment, mean industry wages and mean trade wages than the control group. The control group, on the other hand, displays a significantly higher growth in employment in agriculture than the treated group.

Table 4.3: Average annual growth (%) of the outcome variables over the period 2002-2006 for treated and control group

	GDP per capita					
	total	agriculture	industry	services	taxes	
control	0.50%	-2.29%	1.83%	0.90%	1.73%	
treated	1.14%	-0.60%	5.27%	0.46%	2.24%	
t-test	*-1.390*	*-1.547*	*-3.513*	*1.062*	*-0.684*	
	Employment					
	total	agriculture	industry	services	construction	trade
control	5.51%	21.19%	24.06%	5.72%	34.04%	12.28%
treated	7.26%	8.10%	37.40%	7.29%	49.61%	14.89%
t-test	*-1.882*	*2.167*	*-1.102*	*-1.314*	*-0.546*	*-1.357*
	Wages					
	total	agriculture	industry	services	construction	trade
control	2.63%	4.46%	3.05%	2.51%	4.24%	1.81%
treated	2.75%	5.16%	5.25%	1.66%	9.09%	2.50%
t-test	*-0.215*	*-1.079*	*-2.504*	*1.472*	*-1.618*	*-1.846*

Source: Own calculations based on data from IPEA and RAIS

It is important to note that this table is only an illustration of the average situation between the control and treated group and that no causal relation can be drawn from these tables. The reason is that these tables just compare the averages of two groups and don't take into consideration other factors that might have influenced the values of these outcomes besides sugarcane expansion. The purpose of this study is to establish a causal relation between sugarcane expansion and these outcomes variables using a propensity score-based methodology.

4.4 Results

This section is composed of two parts. In the first part, we present descriptive statistics that demonstrate the relative importance of the sugarcane-related activities and of the agricultural and industrial sector in the economy. In the second part, we estimate the Average Treatment Effect on the Treated (ATT) using propensity-based estimators and then rely on the findings of the first part to interpret these estimates. The analysis in both parts is based on the treatment and control groups that are obtained after the estimation of the propensity score. Since this study uses the same data and model specification for the estimation of the propensity score as in Chapter 2, we refer the reader to this chapter for the model specification as well as the covariate balancing tests and histograms. Table 4.4 summarizes the key statistics of this logit model. Note that in the final analysis, both the treatment and control group are composed of fewer municipalities due to data availability issues for the control variables and because some of the municipalities in the control group were eliminated since they did not belong to the "region of common support"[33].

[33] The region of common support is obtained as defined by Dehejia and Wahba (1999; 2002). These authors suggest that the common support should be imposed by eliminating all those observations in the control group that have an estimated propensity score lower than the lowest estimated propensity score in the treatment group.

Table 4.4: Summary of propensity score specification and size of treated and control group within region of common support

pseudo R^2	blocks	No. of control	No. of treated	region of common support
0.4060	5	65	150	[0.136613, 0.993013]

4.4.1 Relative importance of the sugarcane-related activities and the agricultural and industrial sector in the economy

In order to interpret the differences in growth of the outcome variables between the municipalities where sugarcane production grew strongly and those where it did not grow at all, we need to understand the importance of sugarcane for the treated municipalities and how the two groups differ in the composition of their municipal level GDP. All the municipalities in the treated group experienced, by definition, an increase of sugarcane production of at least 6.8 percent per year. Consequently, the expansion of the sector has influenced employment figures in this group of municipalities. Figure 4.1 displays the employment in the sugarcane sector in the treated group as a share of employment in the agricultural sector and in absolute numbers. Between 2002 and 2005, around 15 percent of the workers employed in agriculture in the treated group were employed in the sugarcane sector. By 2006, this share reached almost 25 percent and was equivalent to 25,000 people working in the sugarcane sector.

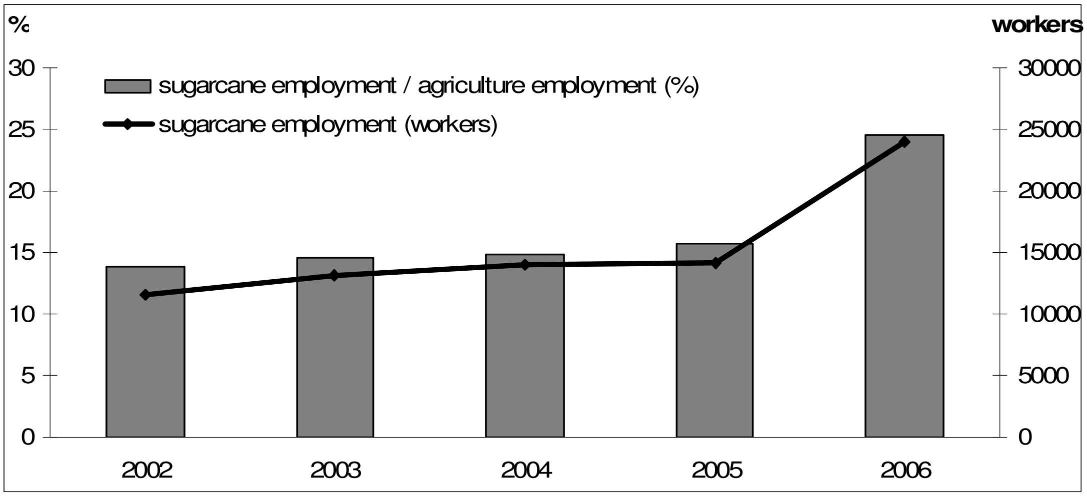

Source: RAIS

Figure 4.1: Employment in the sugarcane sector in the treated group, as share of employment in the agriculture sector and in absolute numbers, 2002-2006

Figure 4.2 demonstrates that employment in the sugar and ethanol sector also rose during the period 2002-2006. Whereas employment in the sugar sector demonstrates a clear upward trend, the employment in the ethanol sector stagnated during some of the years. The employment in these two sectors as a share of total employment in the industry almost doubled between 2002 and 2006. In 2006, over 27 thousand people were employed in the sugar and ethanol sector combined, which accounted for almost 12 percent of total employment in the industry sector.

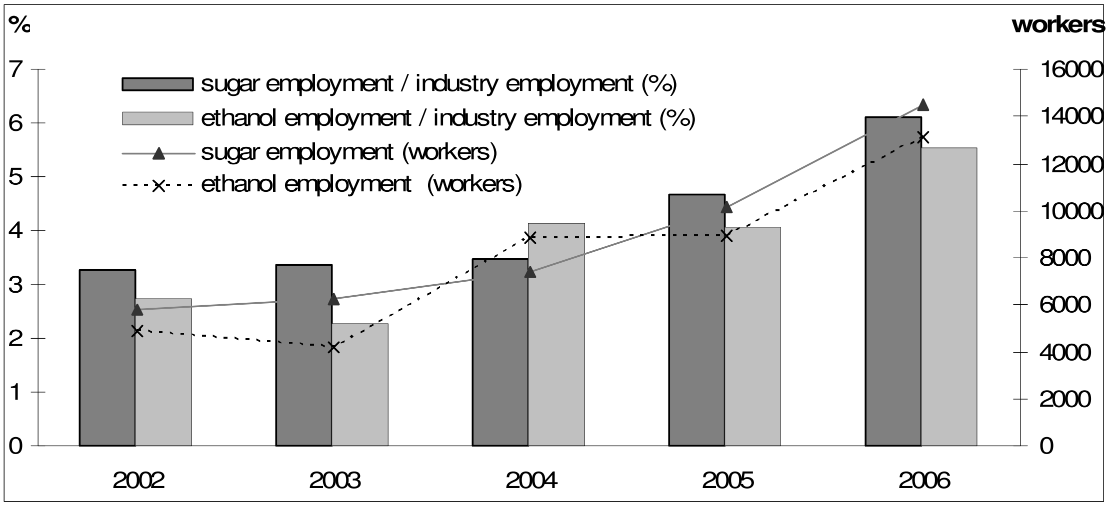

Source: RAIS

Figure 4.2: Employment in the sugar and ethanol sector in the treated group, as share of employment in the industry sector and in absolute numbers, 2002-2006

One of the outcome variables is growth in average monthly income in the different sectors. In order to ascertain sugarcane's role in these outcome variables, we compare the mean monthly incomes in the sugarcane-related sectors with the overall average of the associated sectors. Figure 4.3 compares the mean monthly incomes per worker in the sugarcane sector with those in the agriculture sector, while Figure 4.4 compares these incomes between the sugar, ethanol and industry sectors. These figures both clearly demonstrate that jobs in the sugarcane-related sectors on average pay more than jobs in the associated sectors. In 2006, the mean monthly income of people employed in the sugarcane sector was on average 15 percent more than the average in the agricultural sector. This gap was even more pronounced between the sugar and ethanol sector and the industry sector. In 2006, the mean monthly income per worker in the sugar and ethanol sector was respectively 32 percent and 42 percent higher than the industry sector average.

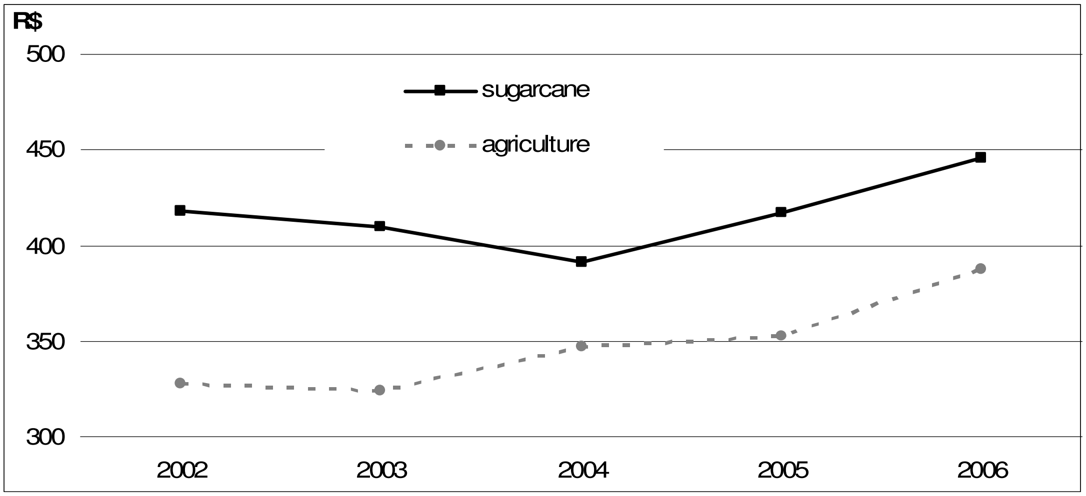

Source: Own calculations based on RAIS and IPEA datasets

Figure 4.3: Mean monthly income per worker in the sugarcane and agriculture sector; municipalities in treated group; 2002-2006 (R$, constant 2000 values)

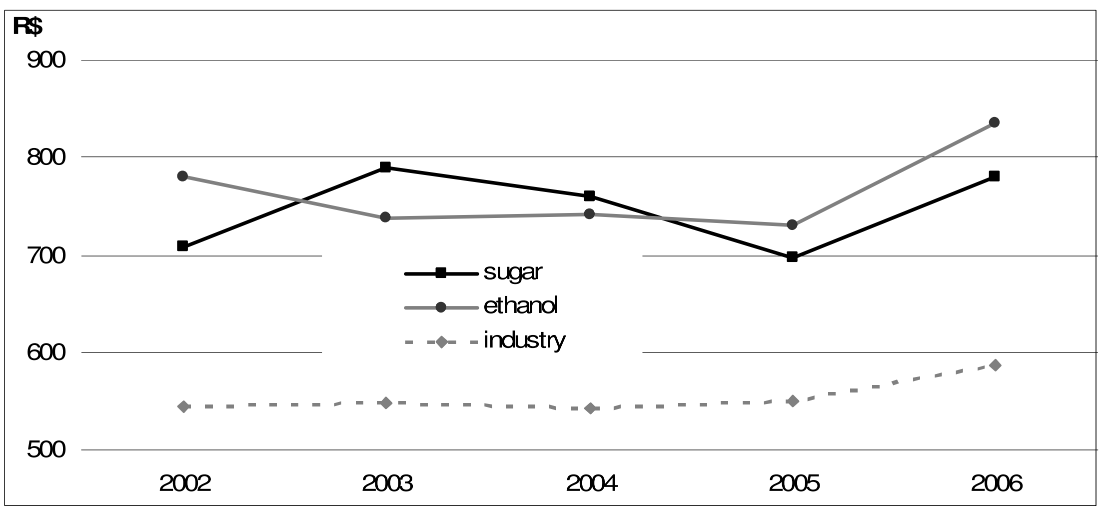

Source: Own calculations based on RAIS and IPEA datasets

Figure 4.4: Mean monthly income per worker in the sugar, ethanol and industry sector; municipalities in treated group; 2002-2006 (R$, constant 2000 values)

One of the possible reasons that sugarcane-expanding municipalities didn't see a greater growth in GDP per capita compared to the sugarcane non-expanding municipalities is that most of the growth occurred in agriculture Value Added (VA) but that this is only a small component of overall GDP. Another explanation is that it occurred in industry VA but that this sector grew at similar rates in the control and treated groups. In order to interpret and analyze the difference in growth in the several components of GDP, we need to examine the relative importance of the different sectors for both the control and treated groups. Table 4.5 compares the relative shares of the different components of total GDP between the treated and control group. This table shows that agriculture and services play a relatively larger role in the treated group than in the control group, while industry and taxes account for a relatively larger share of total GDP in the control group. In both the treated and control group, the service sector is the largest contributor to total GDP, accounting for around 60 and 50 percent of total GDP in the treated and control group, respectively.

Table 4.5: Relative share of different components in total GDP (%), control and treated groups, 2002-2006

		2002	2003	2004	2005	2006
Agriculture	control	5.3	5.0	4.5	4.0	3.7
	treated	8.9	8.7	7.9	6.9	7.7
Industry	control	29.3	30.3	32.7	32.4	32.7
	treated	22.3	22.7	24.8	24.3	24.3
Services	control	52.2	52.1	49.3	50.6	50.4
	treated	59.0	58.9	57.6	59.2	58.3
Taxes	control	13.2	12.6	13.5	13.0	13.2
	treated	9.7	9.7	9.7	9.7	9.6

Source: IBGE

4.4.2 Average treatment effect on the treated (ATT)

All estimates are obtained with STATA. We use the program *atts*, developed by Becker and Ichino (2002), to estimate the blocking estimator and obtain bootstrapped standard-errors, bias and confidence intervals. For the reweighting and mixed estimators, we construct bootstrapped standard errors using the technique described in Busso and Kline (2008). The covariates that were included in the mixed models' regression part are selected using backward stepwise selection. The same subset of regression covariates is used for both the mixed blocking and mixed reweighting estimators.

4.4.2.1 ATT estimation results

Table 4.6 represents the ATT estimates and summary statistics for the first set of outcome variables, namely growth in GDP per capita. The table first displays the estimates for total GDP per capita and then reports these estimates broken down by the different components (agriculture, industry, services, and taxes). The ATT estimates for growth in total GDP per capita confirm the results Chapter 2, which showed that the blocking and reweighting estimators displayed no statistically significant effect. The current results for the blocking and reweighting estimators are very similar to those in Chapter 2, and the mixed estimators reinforce these results.

The ATT estimates are statistically significant for the agriculture and industry sectors, with greater significance levels for the latter. In terms of agriculture VA, the treated group experienced an average annual growth that was around 3 to 4 percent (depending on the estimator) higher than in the control group. For the industrial sector, this difference in growth was around 3 percent. These two results hence suggest that the sectors that are directly related to sugarcane, namely agriculture and industry, experienced a higher growth in the treated group due to the sugarcane production increase.

Even though the ATT estimates for the two other components, namely services and taxes, are positive, they are not statistically significant. This might explain why there is no statistically significant difference in growth in total GDP per capita. Indeed, when looking at Table 4.5, which illustrates the relative importance of the different components of total GDP, we see that the agriculture and industry sector make up around 30 percent of total GDP, while services and taxes are responsible for the remaining 70 percent.

When analyzing the performances of the four different estimators, two issues arise. First, all estimators report rather similar ATT estimates, hence reinforcing the results. Second, the mixed blocking estimator has the highest mean squared errors (MSE) for most of the estimations, which indicates that this estimator is the least efficient. The MSE of the mixed reweighting estimators, on the other hand, are comparable to those of the blocking and reweighting estimators. This finding shows that in addition to the blocking and reweighting estimators, which were selected because Busso, McCrary and DiNardo (2008) demonstrate that they perform best in small samples, also the mixed reweighting estimator can now be considered as an effective estimator for small sample analyses.

Table 4.6: ATT and summary statistics for GDP per capita growth (total and by sector)

		blocking	reweighting	mixed blocking	mixed reweighting
Total	ATT	1.048	1.178	1.226	1.222
	std. err.	(0.773)	(0.831)	(8.846)	(0.845)
	bias	0.000	-0.022	0.237	-0.004
	MSE	0.597	0.691	78.301	0.714
Agriculture	ATT	3.335	3.592	3.511	3.823
	std. err.	(1.867)*	(2.000)*	(2.332)	(2.030)*
	bias	0.016	-0.037	-0.160	-0.070
	MSE	3.486	4.002	5.463	4.126
Industry	ATT	2.840	2.865	2.839	2.867
	std. err.	(1.228)**	(1.291)**	(1.336)**	(1.297)**
	bias	-0.002	-0.004	-0.044	-0.015
	MSE	1.508	1.667	1.788	1.683
Services	ATT	0.031	0.218	0.431	0.286
	std. err.	(0.504)	(0.513)	(5.433)	(0.543)
	bias	0.006	-0.028	-0.083	-0.012
	MSE	0.254	0.264	29.528	0.295
Taxes	ATT	0.247	0.355	0.687	0.439
	std. err.	(1.220)	(1.122)	(32.577)	(1.040)
	bias	0.023	-0.035	0.302	-0.026
	MSE	1.489	1.260	1061.328	1.082

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Table 4.7 reports the ATT estimates for growth in employment, both in aggregate terms and for each of the main sectors of the economy. Similar to the analysis of growth in GDP per capita, the mixed blocking estimator displays the highest MSE for all estimations. Moreover, none of the mixed blocking ATT estimates are statistically significant, while the remaining three estimators report similar significance levels for the different outcome variables. Since the mixed blocking estimators are clearly the least efficient estimators, we base our analysis of the estimates and summary statistics in Table 4.7 on the blocking, reweighting and mixed reweighting estimators.

Growth in total employment is significantly higher in the treated group than in the control group, with estimates indicating a difference of 3 percent per annum. This difference is not driven by employment growth in the agricultural sector. On the contrary, even though the estimates are not statistically significant, the treated group experienced a negative growth compared to the control group. The main contributors to the growth in total employment are employment increases in the industry, services and trade sector, which all demonstrate positive and statistically significant ATT estimates.

In the industry sector, the blocking and reweighting estimators show that growth in employment in the industry sector is around 20 percent higher in the treated group than in the control group. The mixed reweighting estimator displays lower ATT estimates of 14 percent which are statistically significant at the 11 percent level (t-stat of 1.605). These estimates suggest that the boom in the sugar and ethanol production in São Paulo state led to higher growth in employment in the industry sector. This is in concordance with Figure 4.2 presented in the previous section. Indeed, this figure demonstrates that the share of employment in the sugar and ethanol sector in total industry employment almost doubled between 2002 and 2006, from 6 percent to 12 percent.

The sugarcane-expanding municipalities also experienced higher growth in employment in the trade sector than their sugarcane non-expanding counterparts. As mentioned in the introduction, both sugar and ethanol are export products: sugar is mainly exported to the international market while ethanol is transported to the national and international markets. The increase in sugarcane production was driven by the upsurge in demand for these two products on the domestic and international markets. As such, it is not surprising that the treated group has experienced a significantly higher growth in employment in the trade sector.

Table 4.7: ATT and summary statistics for employment growth (total and by sector)

		blocking	reweighting	mixed blocking	mixed reweighting
Total	ATT	3.311	3.281	3.355	2.701
	std. err.	(1.339)**	(1.377)**	(2.292)	(1.555)*
	bias	-0.008	-0.025	-0.030	0.019
	MSE	1.793	1.896	5.254	2.417
Agriculture	ATT	-3.582	-0.962	-3.982	-1.735
	std. err.	(8.359)	(5.609)	(8.099)	(5.758)
	bias	0.076	-0.119	-0.271	-0.098
	MSE	69.886	31.473	65.673	33.159
Industry	ATT	21.837	20.768	13.072	14.290
	std. err.	(7.909)***	(8.133)**	(359.763)	(8.902)
	bias	0.001	-0.208	-6.498	-0.335
	MSE	62.545	66.196	1.29E+05	79.355
Services	ATT	3.249	3.550	2.684	2.964
	std. err.	(1.545)**	(1.481)**	(1.942)	(1.753)*
	bias	-0.012	-0.051	0.347	-0.046
	MSE	2.389	2.197	3.891	3.076
Construction	ATT	34.935	45.926	41.666	49.222
	std. err.	(41.408)	(42.433)	(54.193)	(45.621)
	bias	0.339	-2.431	0.381	-2.654
	MSE	1714.724	1806.429	2937.056	2088.335
Trade	ATT	3.958	4.297	2.658	3.053
	std. err.	(1.882)**	(1.694)**	(4.214)	(1.505)**
	bias	0.021	-0.051	-0.042	-0.103
	MSE	3.541	2.872	17.763	2.275

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

As demonstrated in Table 4.8, there is no statistically significant effect of sugarcane expansion on the growth in overall mean wages in the treated group. On a disaggregated level, we do however see that mean monthly wages in the agriculture, industry, construction and trade sector have experienced growth due to the increase in sugarcane production. This is shown by the ATT estimates in Table 4.8 where the estimated values for agriculture are below 1 percent and significant at the 5 percent level for two of the four estimators. The ATT estimates are significant for all estimators in the industry sector and reach almost 3 percent. Figure 4.3 and Figure 4.4 show that both in the agriculture and industry sector, the sugarcane-related mean wages were well above the average in the respective sector. In addition, mean monthly wages in the sugarcane, sugar and ethanol sector demonstrated a steady growth between 2002 and 2006, which contributes to the findings in Table 4.8. Sugarcane expansion also resulted in higher growth in wages in the construction and trade sector. Since there are no disaggregated data available on wages in construction and trade related to sugarcane, sugar and ethanol activities, it is harder to interpret the exact reason for these increases in growth.

Table 4.8: ATT and summary statistics for growth in wages (total and by sector)

		blocking	reweighting	mixed blocking	mixed reweighting
Total	ATT	-0.112	-0.004	-0.596	-0.417
	std. err.	(0.691)	(0.717)	(2.506)	(0.925)
	bias	0.009	-0.053	0.197	-0.066
	MSE	0.477	0.517	6.317	0.860

Table 4.8 (Continued)

Agriculture	ATT	0.921	0.927	0.696	0.679
	std. err.	(0.451)**	(0.446)**	(0.490)	(0.458)
	bias	0.008	-0.008	-0.008	0.006
	MSE	0.203	0.199	0.240	0.210
Industry	ATT	2.689	2.508	2.871	2.479
	std. err.	(1.089)**	(1.023)**	(1.147)**	(1.028)**
	bias	-0.008	-0.004	0.087	0.000
	MSE	1.187	1.046	1.323	1.058
Services	ATT	-0.802	-0.754	-1.205	-1.165
	std. err.	(0.740)	(0.765)	(1.393)	(0.914)
	bias	0.011	-0.066	0.118	-0.078
	MSE	0.548	0.590	1.954	0.841
Construction	ATT	8.629	8.476	10.332	7.222
	std. err.	(4.264)**	(3.903)**	(5.924)*	(3.241)**
	bias	0.085	0.013	0.305	0.034
	MSE	18.187	15.234	35.184	10.508
Trade	ATT	0.925	1.108	1.470	1.312
	std. err.	(0.559)*	(0.528)**	(2.482)	(0.537)**
	bias	0.008	-0.016	-0.084	-0.052
	MSE	0.312	0.279	6.166	0.291

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

4.4.2.2 Robustness checks

Two different sets of robustness checks are performed to test the findings in Table 4.6, Table 4.7, and Table 4.8. In the first set, we vary the region of common support so as to include more control observations in the analysis. In the second set, we use alternative definitions for classification into the treated group.

Alternative regions of common support

Under the original analysis, the region of common support is defined as proposed by by Dehejia and Wahba (1999; 2002). That is, the common support's lower bound is the lowest estimated propensity score found in the treated group and all observations in the control group that have estimated propensity scores below this value are eliminated. Todd (2008) notes that this rule might be too stringent as potentially good matches just outside the common support might be lost. Therefore, we define an alternative region of common support that includes 20 percent of those control variables with the highest estimated propensity scores[34]. Consequently, the control group is now composed of 75 observations, compared to 66 observations in the original analysis. Table 4.9 reports the ATT estimates for GDP per capita growth with the alternative region of common support. When comparing this table with Table 4.6, we see that the same sectors display statistically significant ATT estimates, namely agriculture and industry. Furthermore, the ATT point estimates are comparable between the two tables.

Table 4.9: ATT for alternative region of common support - GDP per capita growth

		blocking	reweighting	mixed blocking	mixed reweighting
Total	ATT	1.017	1.160	1.164	1.204
	std. err.	(0.762)	(0.819)	(2.472)	(0.835)
Agriculture	ATT	3.239	3.538	3.094	3.579
	std. err.	(1.852)*	(1.996)*	(4.838)	(2.003)*
Industry	ATT	2.803	2.856	2.722	2.720
	std. err.	(1.169)**	(1.279)**	(1.327)**	(1.229)**
Services	ATT	0.012	0.208	0.413	0.264

[34] The alternative region of common support is exactly the same as the one called "region of common support 2" in Chapter 2.

Table 4.9 (Continued)

	std. err.	(0.489)	(0.508)	(1.608)	(0.531)
Taxes	ATT	0.207	0.349	0.567	0.107
	std. err.	(1.164)	(1.113)	(1.304)	(0.986)

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Table 4.10 displays the ATT estimates for employment growth under the broader region of common support. This table demonstrates that the results in Table 4.7 are robust to an alternative region of common support since in both tables the same outcome variables display statistically significant ATT estimates. These outcome variables are total employment, employment in industry, in services and in trade. The values of the statistically significant ATT estimates in Table 4.10 are similar to those in Table 4.7.

Table 4.10: ATT for alternative region of common support – employment growth

		blocking	reweighting	mixed blocking	mixed reweighting
Total	ATT	3.444	3.307	3.463	2.671
	std. err.	(1.328)**	(1.374)**	(4.699)	(1.601)*
Agriculture	ATT	-3.805	-0.989	-4.739	-2.053
	std. err.	(8.458)	(5.666)	(8.654)	(5.983)
Industry	ATT	21.501	20.658	11.322	14.012
	std. err.	(7.951)***	(8.203)**	(25.047)	(9.092)
Services	ATT	3.566	3.591	3.266	2.662
	std. err.	(1.501)**	(1.469)**	(1.918)*	(1.699)
Construction	ATT	32.186	45.154	39.861	48.272
	std. err.	(42.064)	(42.480)	(51.141)	(45.334)
Trade	ATT	4.010	4.302	2.802	3.103
	std. err.	(1.847)**	(1.679)**	(2.308)	(1.462)**

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Also for the set of outcome variables related to growth in monthly mean wages, the results are robust to the alternative region of common support. Indeed, comparing Table 4.11 with Table 4.8 shows that the ATT estimates for growth in wages in the agriculture, industry, construction and trade sectors are statistically significant in both tables and attain similar point values.

Table 4.11: ATT for alternative region of common support – growth in wages

		blocking	reweighting	mixed blocking	mixed reweighting
Total	ATT	-0.117	-0.006	-0.469	-0.387
	std. err.	(0.688)	(0.716)	(0.952)	(0.930)
Agriculture	ATT	0.938	0.924	0.702	0.688
	std. err.	(0.455)**	(0.445)**	(0.490)	(0.451)
Industry	ATT	2.643	2.491	2.803	2.465
	std. err.	(1.099)**	(1.011)**	(1.106)**	(1.018)**
Services	ATT	-0.851	-0.758	-0.951	-1.115
	std. err.	(0.727)	(0.761)	(7.303)	(0.909)
Construction	ATT	8.620	8.488	10.489	9.530
	std. err.	(4.128)**	(3.943)**	(4.849)**	(4.023)**
Trade	ATT	0.919	1.098	1.414	1.313
	std. err.	(0.562)	(0.531)**	(0.714)**	(0.549)**

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Alternative treatment definitions

In the second set of robustness checks, we vary the threshold for inclusion into the treated group. Recall that municipalities were classified into the treated group if their average annual growth in sugarcane harvested area over the period 2002-2006 was higher than the average of São Paulo state, which was 6.8%. In this section, we consider two alternative thresholds, one that is lower and one that is higher than the São Paulo state average. The first alternative treatment definition classifies municipalities with at an expansion in sugarcane harvested area of at least 5% in the treated group. Under the second alternative treatment definition, the threshold is moved up to 10%.

When the threshold is changed, the group of municipalities considered in the analysis also changes. Table 4.12 summarizes how many municipalities were considered in the treated and control group under the original treated definition and under the two alternative definitions. This table also reports the pseudo R^2, number of blocks and regions of common support for the different definitions. The first row represents these values for the original analysis and is a repetition of the values mentioned in Table 4.4. In the case of the 5% threshold, the treated group is larger than in the original analysis; it includes 164 municipalities (compared to 150 under the 6.8% threshold). With the 10% threshold, the treated group consists of fewer municipalities, namely 123 municipalities. The group of control variables also changes under the two alternative scenarios. Although the estimation of the propensity score is done with the same covariates as specified in the original model[35], the sample has changed. As a result, the

[35] The estimation of the propensity score is also here based on the stratification technique. For the different treatment definitions, we used the same covariates as in the original model and obtained balance in both the covariates and estimates propensity scores in each block.

region of common support is different and includes a different set of control municipalities.

Table 4.12: Summary of propensity score specification and size of treated and control group within region of common support, under original treatment threshold and two alternative treatment thresholds

Treatment threshold:	No. of treated	No. of control	pseudo R^2	blocks	region of common support
≥ 6.8%	150	65	0.4060	5	[.136613, .993013]
≥ 5%	164	69	0.4125	5	[.138965, .993705]
≥ 10%	123	74	0.3837	5	[.095992, .985298]

Table 4.13 reports, under the two alternative treatment definitions, ATT estimates and standard errors for growth in GDP per capita and in its different components. When comparing this table with Table 4.6, we see that the same sectors demonstrate statistically significant results, namely agriculture and industry. Furthermore, the values of the ATT estimates are similar to the ones obtained in the original analysis. This shows that the results in the original analysis are robust to changes in the definition of treatment.

Table 4.13: ATT for varying treatment thresholds – GDP per capita growth

		blocking	reweighting	mixed blocking	mixed reweighting
Treatment: sugarcane expansion ≥ 5%					
Total	ATT	1.015	1.293	1.137	1.353
	std. err.	(0.751)	(0.902)	(29.725)	(0.934)
Agriculture	ATT	3.120	3.516	3.362	3.630
	std. err.	(1.904)	(2.089)*	(2.727)	(2.120)*
Industry	ATT	2.620	2.829	2.517	2.693
	std. err.	(1.185)**	(1.350)**	(4.409)	(1.283)**
Services	ATT	0.040	0.279	0.261	0.329
	std. err.	(0.490)	(0.538)	(7.124)	(0.549)
Taxes	ATT	0.519	0.614	0.647	0.300

Table 4.13 (Continued)

	std. err.	(1.275)	(1.166)	(43.955)	(1.089)
Treatment: sugarcane expansion ≥ 10%					
Total	ATT	1.087	1.212	1.267	1.174
	std. err.	(0.779)	(0.805)	(2.537)	(0.790)
Agriculture	ATT	2.861	3.271	2.759	3.447
	std. err.	(1.756)	(1.772)*	(2.334)	(1.804)*
Industry	ATT	3.204	2.953	3.134	2.883
	std. err.	(1.249)**	(1.321)**	(1.534)**	(1.254)**
Services	ATT	0.016	0.269	0.475	0.337
	std. err.	(0.528)	(0.533)	(1.217)	(0.553)
Taxes	ATT	0.164	0.504	0.669	0.218
	std. err.	(1.219)	(1.090)	(12.944)	(0.995)

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Table 4.14 presents the ATT estimates for varying treatment thresholds for growth in employment. When comparing these estimates with the ones in Table 4.7, we can conclude that the estimates for employment growth are also robust for the different definitions for treatment. Indeed, the findings show that sugarcane expansion led to growth in total employment and in employment in industry, services and trade under the two alternative definitions. The ATT point estimates are also comparable to the ones in Table 4.7. However, at the 10% treatment threshold, the ATT estimates for industry and services are only statistically significant for the reweighting estimators.

Table 4.14: ATT for varying treatment thresholds – employment growth

		blocking	reweighting	mixed blocking	mixed reweighting
Treatment: sugarcane expansion ≥ 5%					
Total	ATT	2.877	3.271	2.592	2.677
	std. err.	(1.352)**	(1.375)**	(29.232)	(1.549)*
Agriculture	ATT	-3.719	-0.320	-5.073	-1.788
	std. err.	(8.301)	(5.224)	(8.615)	(5.529)

Table 4.14 (Continued)

Industry	ATT	18.727	18.458	10.982	13.231
	std. err.	(7.476)**	(7.630)**	(214.485)	(8.083)
Services	ATT	2.950	3.519	2.711	2.810
	std. err.	(1.588)*	(1.480)**	(7.358)	(1.524)*
Construction	ATT	28.185	45.554	41.978	50.695
	std. err.	(35.481)	(40.188)	(49.690)	(44.247)
Trade	ATT	3.968	4.383	2.243	3.003
	std. err.	(1.815)**	(1.655)***	(8.108)	(1.457)**
Treatment: sugarcane expansion ≥ 10%					
Total	ATT	3.071	3.266	2.672	2.517
	std. err.	(1.798)*	(1.457)**	(2.309)	(1.642)
Agriculture	ATT	-3.479	-2.794	-4.083	-3.747
	std. err.	(7.471)	(6.164)	(7.791)	(6.454)
Industry	ATT	18.993	23.159	7.162	13.477
	std. err.	(12.663)	(10.041)**	(74.528)	(9.364)
Services	ATT	2.463	3.199	1.648	2.770
	std. err.	(1.988)	(1.529)**	(2.416)	(1.723)
Construction	ATT	-21.690	-4.956	15.818	1.562
	std. err.	(22.992)	(21.509)	(181.652)	(23.711)
Trade	ATT	5.296	4.769	3.918	3.575
	std. err.	(2.046)***	(1.811)***	(4.798)	(1.576)**

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

The ATT estimates and standard errors for growth in wages (total and by sector) for alternative treatment definitions are represented in Table 4.15. Under the original threshold (Table 4.8) the ATT estimates were statistically significant for agriculture, industry, construction and trade. The same output variables display statistically significant ATT estimates for the 5% threshold. At the 10% threshold, however, the ATT estimates for construction are no longer statistically significant. The ATT point estimates in Table 4.15 are comparable to those in Table 4.8.

Table 4.15: ATT for varying treatment thresholds – growth in wages

		blocking	reweighting	mixed blocking	mixed reweighting
Treatment: sugarcane expansion ≥ 5%					
Total	ATT	-0.254	-0.092	-0.913	-0.316
	std. err.	(0.804)	(0.660)	(2.249)	(0.870)
Agriculture	ATT	0.929	1.073	0.683	0.828
	std. err.	(0.453)**	(0.423)**	(0.704)	(0.451)*
Industry	ATT	2.545	2.457	2.654	2.322
	std. err.	(1.041)**	(0.983)**	(1.063)**	(1.011)**
Services	ATT	-0.985	-0.637	-1.402	-1.006
	std. err.	(0.807)	(0.708)	(4.567)	(0.844)
Construction	ATT	8.091	7.917	9.288	7.528
	std. err.	(3.767)**	(3.570)**	(10.197)	(3.218)**
Trade	ATT	0.926	1.156	1.309	1.410
	std. err.	(0.534)*	(0.540)**	(0.759)*	(0.563)**
Treatment: sugarcane expansion ≥ 10%					
Total	ATT	-0.894	-0.253	-1.502	-0.657
	std. err.	(1.221)	(0.813)	(1.527)	(1.041)
Agriculture	ATT	0.989	1.026	0.875	0.833
	std. err.	(0.443)**	(0.442)**	(0.925)	(0.462)*
Industry	ATT	2.681	2.777	2.709	2.792
	std. err.	(1.094)**	(1.101)**	(1.108)**	(1.114)**
Services	ATT	-1.537	-0.965	-2.115	-1.454
	std. err.	(1.197)	(0.842)	(1.912)	(1.033)
Construction	ATT	4.552	4.206	6.343	3.497
	std. err.	(4.337)	(3.958)	(27.193)	(3.686)
Trade	ATT	0.666	0.937	1.129	1.154
	std. err.	(0.611)	(0.569)*	(0.649)*	(0.556)**

Note: values for bias, standard errors, t-values, and MSE are obtained using bootstrap procedures with 10,000 replications. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

4.5 Conclusion

In Chapter 2 we find that sugarcane expansion in São Paulo state didn't lead to greater growth in GDP per capita in those municipalities that increased sugarcane production. In this chapter, we study the underlying reasons behind these findings by analyzing the impact of sugarcane expansion on the different sectors of the economy. In particular, we look at three sets of outcome variables, namely GDP per capita, employment and wages, and analyze the influence of sugarcane expansion on the growth of these variables at both the aggregate level and by sector.

Since this research is motivated by the results in Chapter 2, we use a similar methodology. This means that we consider the same time period, namely 2002-2006, and use the same definition to classify municipalities into the treated and control group. We also rely on the same set of observable covariates and logit model to estimate the propensity score. Recall that the propensity score is defined as the probability that a municipality expands sugarcane production, given the set of observable covariates. Consequently, our analysis and estimations are applied to the same groups of treated and control municipalities.

The analysis is divided into two parts: a descriptive part and an analytical part. The descriptive part is designed to help interpret the estimates obtained in the second part; it demonstrates the importance of sugarcane-related activities in the treated group and compares the composition of GDP between the treated and control group. In this part, we show that employment in the sugarcane, respectively the sugar and ethanol, sector has not only increased in absolute terms but also as a share of agricultural, respectively industrial, employment. We also show that the mean monthly wages in the sugarcane sector are above the average agricultural wages and that the mean monthly wages in the sugar and ethanol sector are higher than the average monthly wages in industry. Interestingly, wages in both the sugarcane-related sectors and in the agricultural and industry sectors increased in constant terms over the period of analysis and display parallel growth paths. When analyzing the different components of GDP, we observe that the service sector is the largest contributor to total GDP in both the treated and control group. Agriculture contributes on average more to total GDP in the treated group than in the control group, while industry has a higher share in total GDP in the control group.

In the second part of the analysis, we use propensity score-based estimators to estimate the impact of sugarcane expansion on the three sets of outcome variables. In particular, we estimate the average treatment effect on the treated (ATT) for growth in GDP per capita, in employment and in wages, at the aggregate level and at sector levels. We use four different propensity score- based estimators to estimate the ATT: the blocking estimator, the reweighting estimator, the mixed blocking estimator and the mixed reweighting estimator. The former two estimators were selected because Busso, McCrary and DiNardo (2008) show that they perform best in small samples. The latter two estimators were selected as robustness checks to reinforce the results and to find out whether there are other estimators that perform well in small samples. In most of the estimations, the mixed blocking estimator is the least efficient, with MSE values well above those of the other estimators. We hence base most of our analysis on the estimates obtained with the remaining three estimators. Interestingly, the mixed reweighting estimators display MSE values that are comparable to those of the blocking and reweighting estimators. This finding demonstrates that also mixed reweighting estimators can be used in small sample propensity score-based estimations.

The first set of ATT estimates are related to growth in GDP per capita in aggregated terms and at sector level. We confirm the findings of Chapter 2 and note that there is no statistically significant growth in total GDP per capita in the sugarcane-expanding municipalities. Also in the services Value Added (VA) and taxes, the ATT estimates are not statistically significant. In terms of agriculture and industry VA, however, we find that sugarcane expansion had a positive and significant effect. In particular, municipalities that expanded sugarcane production experienced as a result a 3 percent higher growth in those two sectors. The fact that these positive estimates didn't influence total GDP estimates can be explained by the analysis in the first part, which shows that agriculture and industry combined only constitute 30 percent of total GDP in the treated group.

In terms of total employment, the ATT estimates show that the treated municipalities didn't experience a higher growth in total employment due to their sugarcane expansion. In the agricultural sector, the ATT estimates were negative, but not significant. This suggests that even though mechanization has become more prevalent in sugarcane cultivation, it hasn't led to a significant decrease in employment growth in the agricultural sector of the sugarcane-expanding municipalities. The ATT estimates were also insignificant, albeit positive, for growth in employment in the construction sector, indicating that the construction of new mills didn't lead to significant higher employment growth in the treated group.

In three sectors, the ATT estimates related to employment growth were positive and significant, namely industry, services and trade. Employment in industry grew at annual rates that were 20 percent higher due to sugarcane expansion. This is in line with the growth in employment in the ethanol and sugar industry, which expanded at an average annual rate of 29 percent between 2002 and 2006. The ATT estimates related to employment growth in the trade and services sectors averaged to 3 percent for both sectors. The significant effect in the trade sector clearly traces back to the importance of sugar and ethanol on the export market. The positive impact of sugarcane expansion on employment in the service sector reinforces the findings by Smeets et al. (2006), who point out that there are large indirect and induced employment effects associated with sugarcane expansion. Indeed, sugarcane expansion has led to direct employment growth effect in those sectors directly related to sugarcane (agriculture, industry, trade). The fact that there has also been an increase in employment in the services sector suggests that the larger workforce in sugarcane-related activities has stimulated demand for services and hence contributed indirectly to employment growth in the services sector as well.

The impact of sugarcane expansion on wages is positive and significant in the agriculture, industry, construction and trade sector. The ATT estimates attain values around 8 percent in the construction sector, around 3 percent in the industry sector and around 1 percent in the agriculture and trade sectors. The steady growths of constant wages in the sugarcane, sugar and ethanol sector, as illustrated in the first section of the analysis, corroborate these findings. In the agriculture sector, the findings are supported by those of Hoffmann and de Oliveira (2006), who show that wages in the sugarcane sector are higher than the wages of any other agricultural activity.

All ATT estimates are subjected to two different robustness tests. In the first test, we broaden the region of common support so that more control observations are included in the analysis. The second test considers two alternative definitions for classification into treatment. We show that the ATT estimates are robust to these tests and as such reinforce our results.

The results in this study suggest that sugarcane expansion has positive impacts on local economies. Indeed, all ATT estimates that are statistically significant are positive. These positive effects extend to growth in agricultural and industrial VA; growth in total employment as well as employment in the industry, services and trade sector; and growth in wages in agriculture, industry, construction and trade. Interestingly, the ATT estimates for growth in total GDP per capita and total wages are not significant. This might be due to the fact that for the effects to filter through to the entire economy, we need to consider a longer time frame. It will hence be interesting to update these findings in the future with data for 2007 through 2010, which are years characterized by even more extensive expansion of sugarcane production.

CHAPTER 5

CONCLUSION

This dissertation consists of three studies, which each examine a different aspect of the economic impacts of sugarcane expansion in Brazil. The three studies are described in Chapter 2 through Chapter 4. The study in Chapter 2 was conducted first and motivated the research topics of the other two studies. In particular, in Chapter 2 we analyze whether sugarcane expansion in São Paulo state resulted in higher economic growth. In Chapter 3, we evaluate the same question for different regions in Brazil. In this Chapter, we also estimate what would have happened to economic growth in sugarcane non-expanding municipalities if they had expanded sugarcane production. Finally, in Chapter 4, we further examine the results found in Chapter 2 and Chapter 3 by analyzing the effects of the sugar and ethanol boom in São Paulo state on the main sectors of the economy.

We use the same methodological design in each study. To estimate the causal effects of sugarcane expansion on local economies, we rely on estimators that are based on the propensity score. We choose this set of estimators because they account for the fact that different municipalities have a different propensity to expand sugarcane production. Indeed, in order to establish a causal relationship between sugarcane expansion and economic variables, we need to construct appropriate counterfactuals. By defining counterfactuals in terms of the propensity score, we assure that we are comparing municipalities that are as similar as possible in all observable aspects, measured by the propensity score, except for the extent to which they experienced growth in sugarcane production. The propensity score in this dissertation is then defined as the probability that a municipality expands sugarcane production, given a set of observable characteristics. One of the most important characteristics for which we control when we construct the propensity score is the suitability of the land to grow sugarcane. The Government of São Paulo and, more recently, EMBRAPA have published the results of their agro-ecological zoning project. In this project, the area of each municipality is classified according to its suitability to grow sugarcane. We also control for possible endogeneity problems by including lagged values of the outcome variables in the propensity score.

Given that over two thirds of Brazilian sugarcane is grown in the state of São Paulo and that most of the recent expansion took place in this state, we expected that sugarcane expansion would have had a positive effect on economic growth. The results of Chapter 2 and Chapter 3, however, indicate that sugarcane expansion had no significant impact on economic growth in the sugarcane-expanding municipalities in São Paulo state. In Chapter 4, we examine the underlying reasons for these findings by looking at growth in GDP per capita, in employment and in wages in the different sectors of the economy. We formulate two initial hypotheses that could explain why we didn't see any statistically significant effect of sugarcane expansion on total GDP per capita growth: i) sugarcane expansion influenced GDP per capita growth in some of the sectors, namely agriculture and industry, but these two sectors only constitute a small part of total GDP, ii) the positive impacts of sugarcane expansion in one sector were offset by negative impacts on other sectors. The estimates we obtain show us that the first hypothesis holds true, while the second one doesn't hold. In particular, sugarcane expanding municipalities in São Paulo state did experience a statistically significant growth in GDP per capita in terms of agriculture and industry VA, but these two sectors together only contribute to 30% of total GDP in these municipalities.

The second hypothesis, however, doesn’t hold because all the statistically significant estimates are positive and the few estimates that are negative aren’t statistically significant.

In addition to testing the above two hypotheses, we also analyze growth in employment and wages in the different sectors in order to get a more detailed picture. We show that employment in the sugarcane expanding municipalities grew significantly in aggregate terms, and in the industry, services and trade sectors. We also find that average wages in the agriculture, industry, services and trade sector all demonstrate a significant positive growth as a result of sugarcane expansion. The fact that all statistically significant estimates are positive combined with the fact that we see positive growth rates in employment and wages in many of the sectors suggest that the final findings for São Paulo state in Chapter 2 and Chapter 3 might need to be updated. That is, the results in Chapter 4 strongly indicate that we should have found greater economic growth in terms of total GDP per capita in the sugarcane expanding municipalities. We hypothesize that there is some delayed reaction and hence one of the first future steps of research will be to update the estimates for a longer period. It will be particularly interesting to update these findings with data for 2007 through 2010, which are years characterized by even more extensive expansion of sugarcane production.

In Chapter 3 we analyze the economic growth effects of sugarcane expansion in five different regions: Brazil, the North-Northeast, the Center-South, São Paulo state, and in the region comprised of the states in the Center-south excluding São Paulo state. The ATT estimates are positive and significant for Brazil, the North-Northeast and the Center-South excluding São Paulo. In Brazil and the Center-South excluding São

Paulo, the ATT estimates are around 0.5 percent, indicating that those municipalities that expanded sugarcane production in these regions experienced as a result an average annual growth in GDP per capita of 0.5 percent. This translates into an accumulated growth difference of 3 percent for the entire period 2001-2007. The effect is even larger in the North-Northeast, where the ATT estimates for average annual growth in GDP per capita reach 0.9 percent.

The robustness checks reveal that in the North-Northeast, expansion of sugarcane production at any rate has positive and significant impacts on economic growth. In the Center-South excluding São Paulo, the estimates become statistically insignificant at low rates of sugarcane expansion. These findings show that, even though sugarcane production in the North-Northeast is less productive than in the rest of the country, it should still be promoted because it has positive effects on the local economies. In the Center-South excluding São Paulo however, sugarcane expansion should reach a minimum threshold in order to influence economic growth.

In Chapter 3 we also estimated the average treatment effect on the untreated (ATU) for sugarcane non-expanding municipalities in the Center-South excluding São Paulo. We find that at expansion rates between 1 percent and 7.5 percent, the ATU estimates reach statistically significant values of 0.7 percent. When the threshold of expansion rates also includes very low rates or is limited to rates above 10 percent, the ATU estimates become statistically insignificant. These results show that the sugarcane non-expanding municipalities in this region would have experienced an average annual growth of 0.7 percent in GDP per capita if they had increased sugarcane production between 1 percent and 10 percent.

As mentioned in the introduction, Brazil is planning to double the land devoted to sugarcane plantations between 2010 and 2020. Most of the future expansions are planned in the states in the Center-South excluding São Paulo. The findings in this paper show that these states indeed would benefit in terms of economic growth from sugarcane expansion as long as it stays below a certain level. In addition, the results demonstrate that sugarcane expansion should also be planned in the North-Northeast because even low sugarcane expansion rates in this region result in positive economic growth effects.

APPENDIX

Appendix Table 1: Logit model used to estimate propensity scores (Chapter 1 and Chapter 3)

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 (Chapter 1 and Chapter 3)

	treatment	control	difference	reweighted control	difference	reweighted control	difference
				common support 1		*common support 2*	
variable	mean	mean		mean		mean	
area	497.559	438.637	58.922	511.895	-14.336	510.975	-13.416
sugharv	2182.573	1984.169	198.404	2160.726	21.847	2167.101	15.472
sugharv/totharv	20.778	14.309	6.469**	23.241	-2.463	23.145	-2.367
totharv/area	17.701	14.261	3.440*	17.050	0.651	17.051	0.650
pasture/area	53.941	36.833	17.108***	53.041	0.900	52.936	1.005
gdppc80	4780.974	9211.508	-4430.530	5178.409	-397.435	5171.423	-390.449
gdppc96	5160.327	5211.521	-51.194	5404.076	-243.749	5392.420	-232.093
ag_rented/area	6.183	6.678	-0.495	5.722	0.461	5.744	0.439
ag_occupied/area	0.928	2.628	-1.700***	0.885	0.043	0.892	0.036
ag_partner/area	1.945	1.962	-0.017	2.154	-0.209	2.161	-0.216
ag_owned/area	90.944	88.732	2.212**	91.239	-0.295	91.202	-0.258
rurpop/totpop	28.121	31.458	-3.337	25.003	3.118	25.128	2.993
suitable/area	18.208	7.964	10.244***	13.049	5.159	13.004	5.204
suitable_lim/area	48.368	33.728	14.640***	52.049	-3.681	51.881	-3.513
suitable_restr/area	24.146	15.755	8.391***	25.221	-1.075	25.071	-0.925
N	150	120		66		75	

Note: t-tests are used for difference in means. Single, double, or triple asterisks (*, **, ***) indicate that difference between treatment and control group is statistically significant at 5%, 2.5%, and 1% levels, respectively.

Appendix Table 3: Logit Model used to estimate propensity scores for BR

Logistic Regression			Number of observations			2656
			LR chi2(51)			490.5
			Prob>chi2			0
Log Likelihood	-1448.55		Pseudo R2			0.1448

	Coef.	Std. Err.	z	P>z	[95% Conf. Interval]	
idhm	10.3062	10.7115	0.960	0.336	-10.6881	31.3004
sug_totharv	0.0644	0.0086	7.480	0.000	0.0475	0.0812
popdens	-0.0042	0.0015	-2.780	0.005	-0.0071	-0.0012
gdppc80	0.0009	0.0003	2.870	0.004	0.0003	0.0015
gdppc96	-0.0017	0.0006	-2.600	0.009	-0.0030	-0.0004
rented	-0.0705	0.0762	-0.920	0.355	-0.2198	0.0789
occupied	-0.0068	0.0721	-0.090	0.924	-0.1482	0.1345
owned	-2.0264	1.3601	-1.490	0.136	-4.6921	0.6393
rurpop	0.0085	0.0141	0.600	0.548	-0.0192	0.0361
areasq	0.0000	0.0000	0.740	0.456	0.0000	0.0000
idhmsq	-9.7716	9.8983	-0.990	0.324	-29.1718	9.6287
sugharvsq	0.0000	0.0000	-3.040	0.002	0.0000	0.0000
sug_totharvsq	-0.0006	0.0001	-5.060	0.000	-0.0009	-0.0004

Appendix Table 3 (Continued)

gdppc80sq	0.0000	0.0000	0.510	0.607	0.0000	0.0000
gdppc96sq	0.0000	0.0000	2.050	0.041	0.0000	0.0000
rentedsq	-0.0002	0.0013	-0.190	0.851	-0.0028	0.0023
occupiedsq	-0.0013	0.0012	-1.120	0.263	-0.0036	0.0010
ownedsq	0.0052	0.0030	1.710	0.088	-0.0008	0.0111
highsq	-0.0001	0.0002	-0.380	0.707	-0.0004	0.0002
lowsq	-0.0002	0.0007	-0.260	0.792	-0.0016	0.0012
gdppc96sqrt	0.0635	0.0243	2.610	0.009	0.0158	0.1112
rentedsqrt	0.2015	0.2232	0.900	0.367	-0.2359	0.6389
occupiedsqrt	0.0402	0.2194	0.180	0.855	-0.3899	0.4702
ownedsqrt	20.6055	15.6477	1.320	0.188	-10.0635	51.2745
rurpopsqrt	-0.1916	0.1903	-1.010	0.314	-0.5646	0.1814
sug_totharv%totharv	0.0001	0.0002	0.930	0.354	-0.0002	0.0004
popdens%metrop	-0.0026	0.0014	-1.790	0.074	-0.0054	0.0003
high%totharv	-0.0001	0.0001	-1.070	0.286	-0.0004	0.0001
high%sug_totharv	0.0002	0.0002	1.020	0.310	-0.0002	0.0007
high%popdens	0.0000	0.0001	-0.310	0.757	-0.0001	0.0001
med%sug_totharv	-0.0008	0.0002	-3.700	0.000	-0.0012	-0.0004
med%sugharv	0.0000	0.0000	2.640	0.008	0.0000	0.0000
med%popdens	0.0002	0.0001	2.840	0.004	0.0000	0.0003
low%totharv	0.0004	0.0010	0.400	0.686	-0.0016	0.0024
low%sug_totharv	-0.0002	0.0004	-0.390	0.693	-0.0010	0.0007
low%sugharv	0.0000	0.0000	2.010	0.044	0.0000	0.0000
low%pasture	0.0020	0.0007	3.040	0.002	0.0007	0.0034
low%popdens	0.0002	0.0002	0.980	0.325	-0.0002	0.0007
rented%high	0.0008	0.0006	1.320	0.187	-0.0004	0.0021
rented%med	0.0022	0.0006	3.790	0.000	0.0011	0.0033
rented%low	0.0029	0.0040	0.730	0.468	-0.0049	0.0106
occupied%high	-0.0010	0.0020	-0.510	0.611	-0.0049	0.0029
occupied%low	-0.0006	0.0027	-0.220	0.828	-0.0059	0.0047
occupied%pasture	-0.0008	0.0005	-1.550	0.120	-0.0019	0.0002
owned%high	0.0003	0.0001	2.450	0.014	0.0001	0.0005
owned%low	-0.0015	0.0005	-2.780	0.005	-0.0026	-0.0004
gdppc80%popdens	0.0000	0.0000	1.830	0.067	0.0000	0.0000
gdppc80%idhm	-0.0014	0.0005	-3.040	0.002	-0.0023	-0.0005
gdppc80%rurpop	0.0000	0.0000	0.350	0.727	0.0000	0.0000
gdppc96%popdens	0.0000	0.0000	-0.100	0.918	0.0000	0.0000
gdppc96%idhm	0.0017	0.0007	2.350	0.019	0.0003	0.0032
constant	-59.5104	50.5004	-1.180	0.239	-158.4892	39.4685

Note: the suffixes “sq” and “sqrt” stand for squared and square root, respectively. The % symbol indicates interactions between two variables

Appendix Table 4: Logit Model used to estimate propensity scores for NE

Logistic Regression			Number of observations		1155
			LR chi2(53)		219.26
			Prob>chi2		0
Log Likelihood		-616.944	Pseudo R2		0.1509

	Coef.	Std. Err.	z	P>z	[95% Conf. Interval]	
idhm	-53.5658	737.4867	-0.070	0.942	-1499.013	1391.882
totharv	-0.1271	0.0562	-2.260	0.024	-0.237	-0.017
popdens	-0.0080	0.0039	-2.030	0.042	-0.016	0.000
gdppc80	0.0012	0.0013	0.930	0.354	-0.001	0.004
gdppc96	-0.0039	0.0019	-2.080	0.038	-0.008	0.000
rented	-0.0488	0.1067	-0.460	0.648	-0.258	0.160
occupied	0.1998	0.1412	1.420	0.157	-0.077	0.476
rurpop	-0.2261	0.1457	-1.550	0.121	-0.512	0.060
areasq	0.0000	0.0000	-1.570	0.116	0.000	0.000
idhmsq	-24.3570	238.2455	-0.100	0.919	-491.310	442.596
sug_totharvsq	0.0002	0.0001	4.220	0.000	0.000	0.000
totharvsq	0.0008	0.0004	2.040	0.041	0.000	0.002
pasturesq	0.0004	0.0003	1.150	0.249	0.000	0.001
popdenssq	0.0000	0.0000	0.890	0.374	0.000	0.000
gdppc80sq	0.0000	0.0000	-0.020	0.980	0.000	0.000
gdppc96sq	0.0000	0.0000	-0.550	0.585	0.000	0.000
rentedsq	0.0003	0.0016	0.200	0.840	-0.003	0.004
occupiedsq	-0.0057	0.0030	-1.890	0.059	-0.012	0.000
ownedsq	-0.0003	0.0006	-0.460	0.646	-0.001	0.001
rurpopsq	0.0009	0.0006	1.580	0.115	0.000	0.002
highsq	0.0083	0.0090	0.920	0.356	-0.009	0.026
lowsq	-0.0036	0.0023	-1.560	0.118	-0.008	0.001
areasqrt	0.0482	0.0085	5.650	0.000	0.031	0.065
idhmsqrt	103.5393	706.6920	0.150	0.884	-1281.552	1488.630
totharvsqrt	0.6330	0.3142	2.010	0.044	0.017	1.249
pasturesqrt	0.2488	0.3638	0.680	0.494	-0.464	0.962
popdenssqrt	0.3440	0.0935	3.680	0.000	0.161	0.527
gdppc80sqrt	-0.0229	0.0447	-0.510	0.609	-0.110	0.065
gdppc96sqrt	-0.0361	0.0631	-0.570	0.567	-0.160	0.088
rentedsqrt	-0.2156	0.2844	-0.760	0.448	-0.773	0.342
occupiedsqrt	-0.9895	0.4397	-2.250	0.024	-1.851	-0.128
ownedsqrt	-0.1425	1.5319	-0.090	0.926	-3.145	2.860
rurpopsqrt	1.7834	1.2397	1.440	0.150	-0.646	4.213
highsqrt	2.0245	1.0191	1.990	0.047	0.027	4.022
lowsqrt	-0.5498	0.6014	-0.910	0.361	-1.728	0.629
high%totharv	-0.0551	0.0751	-0.730	0.463	-0.202	0.092
high%pasture	0.0041	0.0020	2.050	0.040	0.000	0.008
low%totharv	0.0064	0.0032	2.020	0.043	0.000	0.013
low%pasture	0.0006	0.0010	0.580	0.561	-0.001	0.003
rented%med	-0.0017	0.0012	-1.450	0.147	-0.004	0.001
rented%low	0.0294	0.0259	1.140	0.256	-0.021	0.080

Appendix Table 4 (Continued)

rented%pasture	0.0019	0.0015	1.310	0.190	-0.001	0.005
occupied%med	0.0106	0.0034	3.140	0.002	0.004	0.017
occupied%low	0.0019	0.0024	0.790	0.427	-0.003	0.007
occupied%pasture	-0.0014	0.0012	-1.180	0.240	-0.004	0.001
owned%high	-0.0067	0.0043	-1.570	0.116	-0.015	0.002
owned%med	-0.0002	0.0001	-1.400	0.162	0.000	0.000
owned%low	0.0015	0.0019	0.800	0.421	-0.002	0.005
owned%pasture	-0.0007	0.0006	-1.160	0.246	-0.002	0.000
gdppc80%idhm	-0.0019	0.0019	-1.020	0.310	-0.006	0.002
gdppc80%rur	0.0000	0.0000	0.180	0.860	0.000	0.000
gdppc96%idhm	0.0073	0.0028	2.600	0.009	0.002	0.013
gdppc96%rur	0.0000	0.0000	1.380	0.169	0.000	0.000
constant	-41.9226	190.9530	-0.220	0.826	-416.184	332.339

Note: the suffixes "sq" and "sqrt" stand for squared and square root, respectively. The % symbol indicates interactions between two variables

Appendix Table 5: Logit Model used to estimate propensity scores for CS

Logistic Regression		Number of observations	1494
		LR chi2(43)	334.72
		Prob>chi2	0
Log Likelihood	-829.633	Pseudo R2	0.1679

	Coef.	Std. Err.	z	P>z	[95% Conf. Interval]	
area	0.0000	0.0002	0.130	0.898	-0.0004	0.0005
sugharv	-0.0005	0.0001	-4.960	0.000	-0.0007	-0.0003
sug_totharv	-0.0066	0.0088	-0.750	0.455	-0.0239	0.0107
totharv	0.0113	0.0037	3.100	0.002	0.0042	0.0185
popdens	-0.0040	0.0014	-2.730	0.006	-0.0068	-0.0011
gdppc96	0.0000	0.0000	-0.170	0.868	-0.0001	0.0001
rented	-0.0582	0.1206	-0.480	0.629	-0.2945	0.1781
occupied	0.3146	0.2020	1.560	0.119	-0.0814	0.7106
owned	-1.1915	2.5209	-0.470	0.636	-6.1323	3.7494
rurpop	0.1400	0.0711	1.970	0.049	0.0006	0.2795
areasq	0.0000	0.0000	-0.300	0.763	0.0000	0.0000
sugharvsq	0.0000	0.0000	2.960	0.003	0.0000	0.0000
pasturesq	-0.0003	0.0002	-1.480	0.140	-0.0007	0.0001
popdenssq	0.0000	0.0000	1.650	0.099	0.0000	0.0000
gdppc80sq	0.0000	0.0000	-0.340	0.730	0.0000	0.0000
gdppc96sq	0.0000	0.0000	0.910	0.365	0.0000	0.0000
rentedsq	0.0014	0.0025	0.550	0.583	-0.0035	0.0062
occupiedsq	-0.0080	0.0077	-1.030	0.301	-0.0232	0.0072
ownedsq	0.0016	0.0054	0.300	0.767	-0.0090	0.0122
rurpopsq	-0.0006	0.0004	-1.770	0.076	-0.0013	0.0001
highsq	-0.0001	0.0002	-0.570	0.566	-0.0006	0.0003

Appendix Table 5 (Continued)

lowsq	0.0003	0.0009	0.300	0.763	-0.0015	0.0021
areasqrt	-0.0047	0.0171	-0.270	0.784	-0.0382	0.0288
sugharvsqrt	0.0684	0.0104	6.600	0.000	0.0481	0.0888
pasturesqrt	-0.2560	0.3172	-0.810	0.420	-0.8776	0.3656
rentedsqrt	0.1878	0.3689	0.510	0.611	-0.5352	0.9107
occupiedsqrt	-0.6892	0.4093	-1.680	0.092	-1.4913	0.1129
ownedsqrt	16.3916	29.9810	0.550	0.585	-42.3700	75.1533
rurpopsqrt	-1.1759	0.5180	-2.270	0.023	-2.1911	-0.1606
highsqrt	-0.0042	0.1424	-0.030	0.976	-0.2832	0.2748
sug_totharv%totharv	-0.0005	0.0003	-1.720	0.085	-0.0010	0.0001
high%pasture	0.0003	0.0002	1.590	0.111	-0.0001	0.0006
med%pasture	0.0001	0.0001	2.090	0.037	0.0000	0.0002
med%popdens	0.0001	0.0001	0.720	0.472	-0.0001	0.0002
low%sugharv	0.0000	0.0000	0.120	0.907	0.0000	0.0000
low%pasture	0.0017	0.0009	1.960	0.049	0.0000	0.0034
low%popdens	0.0001	0.0003	0.420	0.673	-0.0004	0.0007
occupied%high	-0.0028	0.0021	-1.320	0.188	-0.0069	0.0014
occupied%low	0.0019	0.0075	0.250	0.802	-0.0128	0.0166
occupied%pasture	-0.0005	0.0012	-0.420	0.678	-0.0029	0.0019
owned%high	0.0002	0.0004	0.590	0.555	-0.0005	0.0009
owned%low	-0.0012	0.0007	-1.850	0.065	-0.0026	0.0001
owned%pasture	0.0006	0.0004	1.410	0.160	-0.0002	0.0015
constant	-59.5709	101.6640	-0.590	0.558	-258.8287	139.6869

Note: the suffixes "sq" and "sqrt" stand for squared and square root, respectively. The % symbol indicates interactions between two variables

Appendix Table 6: Logit Model used to estimate propensity scores for CSex

Logistic Regression		Number of observations	1180
		LR chi2(71)	253.44
		Prob>chi2	0
Log Likelihood	-620.615	Pseudo R2	0.1696

	Coef.	Std. Err.	z	P>z	[95% Conf. Interval]	
area	0.0001	0.0003	0.450	0.650	-0.0004	0.0007
sugharv	0.0001	0.0006	0.200	0.839	-0.0011	0.0014
sug_totharv	-0.2209	0.0598	-3.690	0.000	-0.3381	-0.1036
totharv	0.0214	0.0158	1.350	0.176	-0.0096	0.0524
popdens	0.0017	0.0071	0.240	0.807	-0.0121	0.0156
metrop	-0.7798	0.3634	-2.150	0.032	-1.4922	-0.0675
gdppc80	0.0001	0.0002	0.590	0.555	-0.0003	0.0006
rented	-0.0641	0.1702	-0.380	0.707	-0.3977	0.2696
occupied	0.2897	0.2508	1.160	0.248	-0.2019	0.7813
owned	0.3740	4.2331	0.090	0.930	-7.9226	8.6707
rurpop	0.1415	0.0884	1.600	0.109	-0.0317	0.3148

Appendix Table 6 (Continued)

areasq	0.0000	0.0000	-0.650	0.518	0.0000	0.0000
sugharvsq	0.0000	0.0000	-1.070	0.286	0.0000	0.0000
sug_totharvsq	0.0015	0.0005	3.280	0.001	0.0006	0.0024
pasturesq	-0.0002	0.0004	-0.620	0.538	-0.0010	0.0005
popdenssq	0.0000	0.0000	-0.870	0.386	0.0000	0.0000
gdppc80sq	0.0000	0.0000	-0.600	0.550	0.0000	0.0000
gdppc96sq	0.0000	0.0000	1.550	0.120	0.0000	0.0000
rentedsq	0.0012	0.0036	0.330	0.745	-0.0058	0.0081
occupiedsq	-0.0047	0.0097	-0.490	0.624	-0.0237	0.0142
ownedsq	-0.0023	0.0087	-0.260	0.795	-0.0193	0.0148
rurpopsq	-0.0006	0.0004	-1.490	0.136	-0.0015	0.0002
highsq	0.0008	0.0005	1.440	0.149	-0.0003	0.0018
medsq	0.0002	0.0002	1.060	0.288	-0.0002	0.0006
lowsq	-0.0005	0.0015	-0.340	0.733	-0.0035	0.0025
areasqrt	-0.0093	0.0231	-0.400	0.687	-0.0546	0.0360
sugharvsqrt	0.0050	0.0369	0.140	0.892	-0.0674	0.0774
sug_totharvsqrt	1.1543	0.3332	3.460	0.001	0.5012	1.8073
totharvsqrt	0.0406	0.1834	0.220	0.825	-0.3189	0.4000
pasturesqrt	-0.1462	0.7432	-0.200	0.844	-1.6028	1.3104
popdensqrt	-0.1608	0.1169	-1.370	0.169	-0.3900	0.0684
gdppc80sqrt	-0.0188	0.0284	-0.660	0.507	-0.0745	0.0368
gdppc96sqrt	-0.0071	0.0133	-0.540	0.591	-0.0332	0.0189
rentedsqrt	0.3904	0.4516	0.860	0.387	-0.4947	1.2755
occupiedsqrt	-0.6809	0.5035	-1.350	0.176	-1.6677	0.3059
ownedsqrt	-0.0739	51.1568	0.000	0.999	-100.3394	100.1917
rurpopsqrt	-1.2776	0.6505	-1.960	0.050	-2.5526	-0.0026
highsqrt	0.2134	0.2545	0.840	0.402	-0.2855	0.7123
lowsqrt	-0.0915	0.3338	-0.270	0.784	-0.7458	0.5628
sug_totharv%totharv	-0.0008	0.0009	-0.880	0.378	-0.0025	0.0010
popdens%metrop	-0.0006	0.0034	-0.170	0.866	-0.0072	0.0061
high%totharv	-0.0003	0.0004	-0.710	0.475	-0.0010	0.0004
high%sug_totharv	0.0004	0.0009	0.460	0.646	-0.0013	0.0021
high%sugharv	0.0000	0.0000	-0.250	0.799	0.0000	0.0000
high%pasture	0.0001	0.0006	0.100	0.917	-0.0011	0.0012
high%popdens	-0.0001	0.0002	-0.530	0.593	-0.0006	0.0003
med%totharv	-0.0007	0.0003	-2.460	0.014	-0.0013	-0.0001
med%sug_totharv	-0.0019	0.0008	-2.360	0.018	-0.0035	-0.0003
med%sugharv	0.0000	0.0000	1.890	0.059	0.0000	0.0000
med%pasture	0.0000	0.0002	0.010	0.993	-0.0004	0.0004
med%popdens	0.0003	0.0001	2.410	0.016	0.0001	0.0006
low%totharv	0.0005	0.0013	0.360	0.722	-0.0022	0.0031
low%sug_totharv	0.0042	0.0026	1.590	0.111	-0.0010	0.0094
low%sugharv	0.0000	0.0000	-0.450	0.650	-0.0001	0.0001
mlow%pasture	0.0027	0.0012	2.280	0.023	0.0004	0.0050
low%popdens	0.0002	0.0003	0.900	0.370	-0.0003	0.0008
rented%high	0.0003	0.0015	0.200	0.840	-0.0026	0.0032
rented%med	0.0031	0.0012	2.600	0.009	0.0008	0.0055
rented%low	0.0027	0.0061	0.440	0.659	-0.0093	0.0147

Appendix Table 6 (Continued)

rented%pasture	-0.0021	0.0015	-1.430	0.152	-0.0049	0.0008
occupied%high	0.0002	0.0041	0.050	0.961	-0.0078	0.0082
occupied%med	-0.0017	0.0018	-0.970	0.330	-0.0052	0.0017
occupied%low	0.0018	0.0094	0.190	0.849	-0.0166	0.0202
occupied%pasture	-0.0004	0.0019	-0.230	0.815	-0.0042	0.0033
owned%high	-0.0006	0.0008	-0.840	0.401	-0.0021	0.0008
owned%low	-0.0015	0.0014	-1.100	0.272	-0.0043	0.0012
owned%pasture	0.0005	0.0010	0.570	0.566	-0.0013	0.0024
gdppc80%popdens	0.0000	0.0000	-0.480	0.633	0.0000	0.0000
gdppc80%rurpop	0.0000	0.0000	0.880	0.381	0.0000	0.0000
gdppc96%popdens	0.0000	0.0000	1.180	0.238	0.0000	0.0000
gdppc96%rurpop	0.0000	0.0000	-0.520	0.604	0.0000	0.0000
constant	-11.9659	175.3542	-0.070	0.946	-355.6538	331.7221

Note: the suffixes "sq" and "sqrt" stand for squared and square root, respectively. The % symbol indicates interactions between two variables

Appendix Table 7: Logit Model used to estimate propensity scores for SP

Logistic Regression		Number of observations	302
		LR chi2(29)	147.16
		Prob>chi2	0
Log Likelihood	-130.899	Pseudo R2	0.3598

	Coef.	Std. Err.	z	P>z	[95% Conf. Interval]	
area	-0.0053	0.0025	-2.130	0.033	-0.0102	-0.0004
sugharv	-0.0006	0.0006	-0.970	0.331	-0.0018	0.0006
totharv	-0.1829	0.1492	-1.230	0.220	-0.4754	0.1095
rented	-0.1549	0.4360	-0.360	0.722	-1.0095	0.6997
occupied	-0.5332	0.6569	-0.810	0.417	-1.8207	0.7543
owned	-12.3916	12.0250	-1.030	0.303	-35.9600	11.1769
rurpop	0.0195	0.0643	0.300	0.762	-0.1065	0.1454
sugharvsq	0.0000	0.0000	0.380	0.705	0.0000	0.0000
sug_totharvsq	-0.0002	0.0002	-0.980	0.329	-0.0007	0.0002
totharvsq	0.0004	0.0010	0.350	0.725	-0.0016	0.0023
popdenssq	0.0000	0.0000	0.950	0.344	0.0000	0.0000
gdppc80sq	0.0000	0.0000	-0.520	0.600	0.0000	0.0000
gdppc96sq	0.0000	0.0000	1.070	0.283	0.0000	0.0000
rentedsq	-0.0004	0.0073	-0.060	0.955	-0.0146	0.0138
occupiedsq	0.0038	0.0375	0.100	0.920	-0.0698	0.0773
ownedsq	0.0271	0.0258	1.050	0.292	-0.0234	0.0777
highsq	0.0001	0.0001	1.320	0.187	-0.0001	0.0003
medsq	-0.0002	0.0001	-1.470	0.142	-0.0004	0.0001
lowsq	0.0130	0.0195	0.670	0.504	-0.0252	0.0512
areasqrt	0.2819	0.1167	2.410	0.016	0.0531	0.5107
sugharvsqrt	0.0630	0.0595	1.060	0.290	-0.0536	0.1796

Appendix Table 7 (Continued)

sug_totharvsqrt	0.0849	0.2787	0.300	0.761	-0.4613	0.6311
totharvsqrt	1.5426	0.8605	1.790	0.073	-0.1441	3.2292
pasturesqrt	0.2531	0.1309	1.930	0.053	-0.0036	0.5097
popdenssqrt	-0.1639	0.0774	-2.120	0.034	-0.3157	-0.0121
rentedsqrt	-0.1216	1.5656	-0.080	0.938	-3.1901	2.9469
occupiedsqrt	0.7342	1.1940	0.610	0.539	-1.6060	3.0745
ownedsqrt	139.3097	140.4456	0.990	0.321	-135.9587	414.5781
rurpopsqrt	-0.5780	0.7295	-0.790	0.428	-2.0078	0.8517
constant	-429.4455	459.9351	-0.930	0.350	-1330.9020	472.0106

Note: the suffixes "sq" and "sqrt" stand for squared and square root, respectively. The % symbol indicates interactions between two variables

Appendix Table 8: Assessing the balance in covariates before and after reweighting based on the propensity score – region BR

variable	treated	control	t-stat	treated	control (weighted)	t-stat (weighted)
area	1753.86	1414.55	1.41	1753.86	1469.10	1.15
high	13.04	2.67	12.56	13.04	13.17	-0.14
med	13.34	6.34	9.07	13.34	12.38	1.14
low	1.70	0.84	3.50	1.70	1.77	-0.25
sugharv	990.71	410.55	5.95	990.71	1086.35	-0.91
sug_totharv	11.78	5.55	7.52	11.78	12.90	-1.25
totharv	17.48	14.41	3.77	17.48	16.12	1.63
pasture	42.74	35.66	7.32	42.74	42.03	0.73
rented	4.07	2.84	5.47	4.07	4.14	-0.28
occupied	2.71	4.09	-7.22	2.71	2.60	0.67
partner	1.88	1.68	1.38	1.88	1.79	0.61
owned	91.59	91.60	-0.04	91.59	91.69	-0.33
popdens	43.42	49.73	-1.99	43.42	44.93	-0.51
rurpop	41.97	49.85	-8.93	41.97	42.16	-0.22
idhm	0.64	0.60	10.15	0.64	0.64	-0.57
gdppc80	3533.43	2942.87	4.91	3533.43	3501.71	0.29
gdppc96	3455.58	2626.04	5.42	3455.58	3276.88	1.17

Appendix Table 9: Assessing the balance in covariates before and after reweighting based on the propensity score – region NE

variable	treated	control	t-stat	treated	control (weighted)	t-stat (weighted)
area	1311.02	917.03	3.27	1311.02	1324.45	-0.10
high	0.57	0.15	2.43	0.57	0.62	-0.26
med	4.07	1.71	3.87	4.07	4.91	-1.25

Appendix Table 9 (Continued)

low	1.19	0.36	2.47	1.19	0.66	1.56
sugharv	1911.33	479.60	5.56	1911.33	1917.77	-0.02
sug_totharv	21.63	6.36	8.39	21.63	22.66	-0.48
totharv	15.49	13.53	1.79	15.49	15.73	-0.21
pasture	24.97	31.05	-4.78	24.97	23.57	1.12
rented	2.80	1.55	3.22	2.80	2.82	-0.04
occupied	4.31	5.12	-2.81	4.31	4.19	0.43
partner	2.08	1.82	0.88	2.08	1.90	0.57
owned	91.15	91.75	-1.05	91.15	91.44	-0.48
popdens	67.89	53.45	1.73	67.89	65.39	0.32
rurpop	56.03	58.39	-1.96	56.03	55.69	0.27
idhm	0.52	0.51	1.30	0.52	0.51	0.56
gdppc80	1574.47	1195.15	3.61	1574.47	1475.62	0.92
gdppc96	1300.90	1112.09	2.72	1300.90	1253.41	0.67

Appendix Table 10: Assessing the balance in covariates before and after reweighting based on the propensity score – region CS

variable	treated	control	t-stat	treated	control (weighted)	t-stat (weighted)
area	988.56	1063.01	-0.75	988.56	1012.46	-0.26
high	18.71	5.02	11.39	18.71	18.69	0.01
med	17.60	10.99	5.88	17.60	17.32	0.23
low	1.95	1.15	2.53	1.95	2.06	-0.29
sugharv	828.98	337.42	4.77	828.98	892.25	-0.56
sug_totharv	9.27	5.12	4.96	9.27	9.79	-0.54
totharv	19.19	16.13	2.77	19.19	20.90	-1.44
pasture	50.90	41.35	8.25	50.90	50.03	0.76
rented	4.73	3.94	2.97	4.73	4.68	0.22
occupied	1.90	2.39	-3.61	1.90	1.90	-0.02
partner	1.73	1.60	0.93	1.73	1.63	0.71
owned	91.85	92.24	-1.17	91.85	91.98	-0.39
popdens	40.44	51.44	-2.46	40.44	42.62	-0.57
rurpop	36.65	42.41	-5.31	36.65	37.74	-1.06
idhm	0.68	0.68	2.48	0.68	0.68	0.69
gdppc80	4337.86	4355.14	-0.10	4337.86	4135.96	1.41
gdppc96	4316.49	3953.25	1.66	4316.49	3994.22	1.46

Appendix Table 11: Assessing the balance in covariates before and after reweighting based on the propensity score – region CSex

variable	treated	control	t-stat	treated	control (weighted)	t-stat (weighted)
area	1187.73	1187.78	0.00	1187.73	1152.83	0.29
high	8.11	3.13	4.67	8.11	8.07	0.04
med	17.35	10.49	5.18	17.35	15.93	1.02
low	2.78	1.27	3.42	2.78	2.13	1.36
sugharv	448.84	253.90	2.26	448.84	327.38	1.40
sug_totharv	5.43	4.26	1.64	5.43	4.71	1.00
totharv	20.22	17.10	2.16	20.22	19.66	0.37
pasture	49.67	42.84	5.09	49.67	48.80	0.64
rented	4.11	3.69	1.53	4.11	3.92	0.69
occupied	2.23	2.44	-1.24	2.23	2.19	0.23
partner	1.67	1.62	0.34	1.67	1.51	1.14
owned	92.21	92.44	-0.62	92.21	92.58	-1.01
popdens	37.90	44.64	-1.48	37.90	36.96	0.22
rurpop	39.89	43.94	-3.14	39.89	40.65	-0.61
idhm	0.67	0.67	-0.41	0.67	0.67	0.91
gdppc80	4095.35	4220.45	-0.63	4095.35	3916.76	1.04
gdppc96	4080.94	3830.25	0.83	4080.94	3599.55	1.61

Appendix Table 12: Assessing the balance in covariates before and after reweighting based on the propensity score – region SP

variable	treated	control	t-stat	treated	control (weighted)	t-stat (weighted)
area	490.90	416.79	1.80	490.90	497.20	-0.15
high	43.05	20.26	7.66	43.05	44.55	-0.47
med	19.20	18.16	0.33	19.20	17.52	0.65
low	0.27	0.14	0.93	0.27	0.08	1.47
sugharv	1703.52	1446.57	0.54	1703.52	1789.18	-0.24
sug_totharv	17.83	12.65	1.71	17.83	19.42	-0.56
totharv	16.82	13.78	1.66	16.82	16.89	-0.05
pasture	54.27	41.66	4.53	54.27	51.70	1.03
rented	6.23	5.81	0.60	6.23	6.08	0.25
occupied	1.07	1.67	-2.56	1.07	0.90	0.94
partner	1.91	1.63	0.75	1.91	1.83	0.22
owned	91.00	91.08	-0.09	91.00	91.57	-0.77
popdens	48.08	76.12	-2.01	48.08	57.19	-0.88
rurpop	28.64	32.15	-1.63	28.64	27.92	0.38
gdppc80	4785.98	4857.94	-0.26	4785.98	4883.10	-0.47
gdppc96	4922.97	4581.42	1.03	4922.97	4839.36	0.30

Appendix Table 13: Logit Model used to estimate propensity scores for region CSex –ATU analysis- scenario 1 (at least 1% growth)

Logistic Regression			Number of observations			1470
			LR chi2(79)			290.81
			Prob>chi2			0.0000
Log Likelihood		-869.253	Pseudo R2			0.1433

	Coef.	Std. Err.	z	P>z	[95% Conf.	Interval]
area	0.0004	0.0002	1.590	0.112	-0.0001	0.0009
sugharv	-0.0005	0.0003	-1.810	0.070	-0.0010	0.0000
sug_totharv	-0.1476	0.0410	-3.600	0.000	-0.2280	-0.0672
totharv	0.0098	0.0438	0.220	0.824	-0.0761	0.0957
pasture	-0.1263	0.1849	-0.680	0.495	-0.4886	0.2361
popdens	-0.0022	0.0055	-0.400	0.688	-0.0130	0.0086
metrop	-1.1112	0.3013	-3.690	0.000	-1.7017	-0.5208
gdppc80	0.0000	0.0002	0.210	0.834	-0.0003	0.0004
gdppc96	-0.0002	0.0002	-0.700	0.482	-0.0006	0.0003
rented	-0.0895	0.1550	-0.580	0.564	-0.3933	0.2143
occupied	0.1311	0.2024	0.650	0.517	-0.2656	0.5278
owned	1.9207	3.0218	0.640	0.525	-4.0019	7.8433
rurpop	0.0325	0.0707	0.460	0.645	-0.1061	0.1712
high	0.2681	0.1806	1.480	0.138	-0.0858	0.6220
med	0.4837	0.2636	1.830	0.067	-0.0331	1.0004
low	0.8200	0.7910	1.040	0.300	-0.7304	2.3704
areasq	0.0000	0.0000	-1.620	0.105	0.0000	0.0000
sugharvsq	0.0000	0.0000	-0.090	0.928	0.0000	0.0000
sug_totharvsq	0.0010	0.0003	2.950	0.003	0.0003	0.0016
totharvsq	0.0002	0.0003	0.960	0.336	-0.0003	0.0007
pasturesq	-0.0003	0.0003	-1.000	0.318	-0.0010	0.0003
popdenssq	0.0000	0.0000	-0.990	0.321	0.0000	0.0000
gdppc80sq	0.0000	0.0000	-0.240	0.809	0.0000	0.0000
gdppc96sq	0.0000	0.0000	1.480	0.139	0.0000	0.0000
rentedsq	0.0014	0.0030	0.470	0.636	-0.0045	0.0074
occupiedsq	-0.0011	0.0073	-0.140	0.885	-0.0154	0.0133
ownedsq	-0.0060	0.0063	-0.940	0.345	-0.0184	0.0064
rurpopsq	-0.0002	0.0003	-0.500	0.619	-0.0008	0.0005
highsq	0.0006	0.0005	1.330	0.184	-0.0003	0.0016
medsq	0.0002	0.0003	0.550	0.583	-0.0004	0.0008
lowsq	0.0000	0.0015	-0.020	0.987	-0.0029	0.0028
areasqrt	-0.0331	0.0191	-1.730	0.083	-0.0705	0.0043
sugharvsqrt	0.0495	0.0215	2.300	0.021	0.0074	0.0916
sug_totharvsqrt	0.7073	0.2162	3.270	0.001	0.2835	1.1311
totharvsqrt	-0.1457	0.2690	-0.540	0.588	-0.6729	0.3815
pasturesqrt	-0.5040	0.5877	-0.860	0.391	-1.6560	0.6479
popdenssqrt	-0.0396	0.0928	-0.430	0.670	-0.2214	0.1423
gdppc80sqrt	-0.0159	0.0253	-0.630	0.530	-0.0655	0.0337
gdppc96sqrt	0.0234	0.0284	0.820	0.410	-0.0323	0.0790

Appendix Table 13 (Continued)

rentedsqrt	0.3345	0.3690	0.910	0.365	-0.3888	1.0578
occupiedsqrt	-0.4557	0.4073	-1.120	0.263	-1.2540	0.3427
ownedsqrt	-16.4354	36.2879	-0.450	0.651	-87.5585	54.6876
rurpopsqrt	-0.3684	0.5272	-0.700	0.485	-1.4016	0.6649
highsqrt	0.1136	0.2334	0.490	0.626	-0.3438	0.5710
medsqrt	0.0011	0.1438	0.010	0.994	-0.2808	0.2830
lowsqrt	0.0139	0.2899	0.050	0.962	-0.5543	0.5821
sugharv%totharv	0.0000	0.0006	-0.080	0.935	-0.0011	0.0011
popdens%metrop	0.0034	0.0027	1.250	0.211	-0.0019	0.0088
high%totharv	-0.0002	0.0003	-0.540	0.591	-0.0009	0.0005
high%sug_totharv	0.0011	0.0009	1.280	0.200	-0.0006	0.0029
high%sugharv	0.0000	0.0000	-1.140	0.253	0.0000	0.0000
high%pasture	0.0007	0.0005	1.210	0.228	-0.0004	0.0017
high%popdens	0.0001	0.0001	0.700	0.485	-0.0002	0.0003
med%totharv	-0.0006	0.0003	-1.990	0.046	-0.0011	0.0000
med%sug_totharv	-0.0018	0.0006	-2.880	0.004	-0.0030	-0.0006
med%sugharv	0.0000	0.0000	2.190	0.029	0.0000	0.0000
med%pasture	0.0001	0.0002	0.480	0.632	-0.0004	0.0006
med%popdens	0.0002	0.0001	1.630	0.102	0.0000	0.0004
low%totharv	-0.0006	0.0014	-0.410	0.682	-0.0033	0.0022
low%sug_totharv	0.0013	0.0021	0.620	0.536	-0.0028	0.0054
low%sugharv	0.0000	0.0000	0.590	0.557	0.0000	0.0001
low%pasture	0.0023	0.0010	2.270	0.023	0.0003	0.0043
low%popdens	0.0001	0.0002	0.300	0.762	-0.0004	0.0005
rented%high	-0.0009	0.0022	-0.440	0.662	-0.0052	0.0033
rented%med	-0.0006	0.0028	-0.210	0.836	-0.0060	0.0049
rented%low	-0.0077	0.0102	-0.750	0.453	-0.0277	0.0124
rented%pasture	0.0004	0.0020	0.200	0.840	-0.0035	0.0043
occupied%high	-0.0043	0.0045	-0.960	0.335	-0.0132	0.0045
occupied%med	-0.0083	0.0035	-2.370	0.018	-0.0152	-0.0015
occupied%low	-0.0075	0.0124	-0.600	0.546	-0.0317	0.0168
occupied%pasture	0.0027	0.0022	1.250	0.212	-0.0015	0.0069
owned%high	-0.0037	0.0018	-2.030	0.042	-0.0073	-0.0001
owned%med	-0.0050	0.0027	-1.870	0.061	-0.0102	0.0002
owned%low	-0.0099	0.0080	-1.250	0.213	-0.0255	0.0057
owned%pasture	0.0021	0.0018	1.160	0.247	-0.0014	0.0055
gdppc80%popdens	0.0000	0.0000	-0.250	0.806	0.0000	0.0000
gdppc80%rurpop	0.0000	0.0000	1.670	0.094	0.0000	0.0000
gdppc96%popdens	0.0000	0.0000	0.800	0.426	0.0000	0.0000
gdppc96%rurpop	0.0000	0.0000	-1.250	0.211	0.0000	0.0000
constant	33.7550	124.1264	0.270	0.786	-209.5283	277.0383

Appendix Table 14: Logit Model used to estimate propensity scores for region CSex –ATU analysis- scenario 2 (at least 5% growth)

Logistic Regression				Number of observations		1291
				LR chi2(79)		278.02
				Prob>chi2		0.0000
Log Likelihood		-723.218		Pseudo R2		0.1612

	Coef.	Std. Err.	z	P>z	[95% Conf. Interval]	
area	0.0003	0.0003	1.220	0.224	-0.0002	0.0009
sugharv	-0.0004	0.0006	-0.810	0.417	-0.0015	0.0006
sug_totharv	-0.1601	0.0515	-3.110	0.002	-0.2609	-0.0592
totharv	-0.0034	0.0493	-0.070	0.945	-0.1001	0.0932
pasture	-0.0633	0.2070	-0.310	0.760	-0.4689	0.3424
popdens	-0.0037	0.0067	-0.550	0.586	-0.0169	0.0095
metrop	-0.9152	0.3423	-2.670	0.007	-1.5861	-0.2443
gdppc80	0.0001	0.0002	0.640	0.523	-0.0003	0.0006
gdppc96	-0.0002	0.0003	-0.690	0.491	-0.0007	0.0003
rented	-0.0679	0.1761	-0.390	0.700	-0.4131	0.2772
occupied	0.2051	0.2511	0.820	0.414	-0.2871	0.6972
owned	1.0501	4.2601	0.250	0.805	-7.2996	9.3997
rurpop	0.1274	0.0815	1.560	0.118	-0.0324	0.2872
high	0.2733	0.1946	1.400	0.160	-0.1081	0.6546
med	0.5312	0.2729	1.950	0.052	-0.0036	1.0660
low	0.9943	0.8847	1.120	0.261	-0.7396	2.7282
areasq	0.0000	0.0000	-1.240	0.217	0.0000	0.0000
sugharvsq	0.0000	0.0000	-0.280	0.780	0.0000	0.0000
sug_totharvsq	0.0012	0.0004	2.820	0.005	0.0004	0.0020
totharvsq	0.0003	0.0003	1.100	0.273	-0.0002	0.0009
pasturesq	-0.0006	0.0004	-1.600	0.110	-0.0014	0.0001
popdenssq	0.0000	0.0000	-0.310	0.757	0.0000	0.0000
gdppc80sq	0.0000	0.0000	-0.600	0.549	0.0000	0.0000
gdppc96sq	0.0000	0.0000	1.440	0.150	0.0000	0.0000
rentedsq	0.0014	0.0034	0.420	0.672	-0.0052	0.0081
occupiedsq	-0.0042	0.0093	-0.450	0.653	-0.0225	0.0141
ownedsq	-0.0045	0.0087	-0.520	0.604	-0.0215	0.0125
rurpopsq	-0.0006	0.0004	-1.450	0.146	-0.0013	0.0002
highsq	0.0005	0.0005	0.890	0.375	-0.0005	0.0014
medsq	0.0002	0.0003	0.550	0.581	-0.0005	0.0008
lowsq	-0.0004	0.0015	-0.250	0.800	-0.0034	0.0026
areasqrt	-0.0295	0.0212	-1.400	0.163	-0.0710	0.0119
sugharvsqrt	0.0517	0.0336	1.540	0.123	-0.0141	0.1176
sug_totharvsqrt	0.7156	0.2861	2.500	0.012	0.1550	1.2763
totharvsqrt	-0.0263	0.3047	-0.090	0.931	-0.6236	0.5710
pasturesqrt	-0.9589	0.6976	-1.370	0.169	-2.3263	0.4084
popdenssqrt	-0.0710	0.1085	-0.650	0.513	-0.2836	0.1417
gdppc80sqrt	-0.0284	0.0280	-1.020	0.310	-0.0833	0.0264
gdppc96sqrt	0.0212	0.0320	0.660	0.507	-0.0415	0.0839

Appendix Table 14 (Continued)

rentedsqrt	0.2421	0.4126	0.590	0.557	-0.5665	1.0507
occupiedsqrt	-0.6159	0.4672	-1.320	0.187	-1.5317	0.2998
ownedsqrt	-5.5061	51.7743	-0.110	0.915	-106.9819	95.9697
rurpopsqrt	-1.1558	0.6063	-1.910	0.057	-2.3442	0.0326
highsqrt	0.1112	0.2518	0.440	0.659	-0.3823	0.6046
medsqrt	-0.0541	0.1592	-0.340	0.734	-0.3660	0.2579
lowsqrt	-0.1322	0.3153	-0.420	0.675	-0.7501	0.4857
sugharv%totharv	-0.0002	0.0007	-0.220	0.828	-0.0015	0.0012
popdens%metrop	0.0014	0.0032	0.420	0.675	-0.0050	0.0077
high%totharv	-0.0003	0.0004	-0.670	0.504	-0.0010	0.0005
high%sug_totharv	0.0011	0.0010	1.110	0.268	-0.0009	0.0031
high%sugharv	0.0000	0.0000	-0.970	0.331	0.0000	0.0000
high%pasture	0.0004	0.0006	0.580	0.561	-0.0008	0.0015
high%popdens	-0.0002	0.0002	-0.840	0.402	-0.0007	0.0003
med%totharv	-0.0006	0.0003	-2.170	0.030	-0.0012	-0.0001
med%sug_totharv	-0.0017	0.0007	-2.430	0.015	-0.0030	-0.0003
med%sugharv	0.0000	0.0000	1.500	0.134	0.0000	0.0000
med%pasture	0.0001	0.0003	0.410	0.680	-0.0004	0.0006
med%popdens	0.0003	0.0001	2.290	0.022	0.0000	0.0005
low%totharv	-0.0010	0.0015	-0.660	0.507	-0.0041	0.0020
low%sug_totharv	0.0026	0.0025	1.020	0.308	-0.0024	0.0075
low%sugharv	0.0000	0.0000	-0.190	0.846	-0.0001	0.0001
low%pasture	0.0023	0.0011	2.020	0.044	0.0001	0.0045
low%popdens	0.0003	0.0003	0.970	0.331	-0.0003	0.0008
rented%high	-0.0018	0.0025	-0.700	0.483	-0.0067	0.0031
rented%med	-0.0010	0.0029	-0.350	0.728	-0.0068	0.0047
rented%low	-0.0077	0.0113	-0.680	0.495	-0.0299	0.0145
rented%pasture	0.0000	0.0022	0.000	0.999	-0.0044	0.0044
occupied%high	-0.0030	0.0048	-0.620	0.533	-0.0125	0.0065
occupied%med	-0.0088	0.0037	-2.390	0.017	-0.0160	-0.0016
occupied%low	-0.0105	0.0136	-0.780	0.438	-0.0372	0.0161
occupied%pasture	0.0025	0.0026	0.960	0.335	-0.0025	0.0075
owned%high	-0.0033	0.0019	-1.690	0.091	-0.0071	0.0005
owned%med	-0.0054	0.0027	-1.960	0.051	-0.0108	0.0000
owned%low	-0.0112	0.0089	-1.260	0.207	-0.0286	0.0062
owned%pasture	0.0021	0.0020	1.060	0.288	-0.0018	0.0060
gdppc80%popdens	0.0000	0.0000	-0.230	0.822	0.0000	0.0000
gdppc80%rurpop	0.0000	0.0000	1.290	0.198	0.0000	0.0000
gdppc96%popdens	0.0000	0.0000	0.800	0.423	0.0000	0.0000
gdppc96%rurpop	0.0000	0.0000	-1.090	0.274	0.0000	0.0000
constant	-0.6800	178.3455	0.000	0.997	-350.2307	348.8707

Appendix Table 15: Logit Model used to estimate propensity scores for region CSex –ATU analysis- scenario 3 (at least 10% growth)

Logistic Regression				Number of observations		1115
				LR chi2(79)		264.63
				Prob>chi2		0.0000
Log Likelihood		-544.943		Pseudo R2		0.1954

	Coef.	Std. Err.	z	P>z	[95% Conf. Interval]	
area	0.0001	0.0003	0.410	0.682	-0.0005	0.0008
sugharv	-0.0004	0.0008	-0.530	0.597	-0.0019	0.0011
sug_totharv	-0.2166	0.0668	-3.240	0.001	-0.3476	-0.0856
totharv	-0.0533	0.0582	-0.920	0.360	-0.1674	0.0608
pasture	-0.0668	0.2545	-0.260	0.793	-0.5656	0.4321
popdens	-0.0043	0.0092	-0.470	0.636	-0.0223	0.0137
metrop	-0.9232	0.4079	-2.260	0.024	-1.7227	-0.1237
gdppc80	0.0026	0.0010	2.680	0.007	0.0007	0.0046
gdppc96	-0.0002	0.0003	-0.530	0.598	-0.0008	0.0005
rented	-0.0938	0.2088	-0.450	0.653	-0.5030	0.3155
occupied	0.1769	0.3037	0.580	0.560	-0.4183	0.7721
owned	-2.4746	5.0167	-0.490	0.622	-12.3072	7.3579
rurpop	0.1611	0.0961	1.680	0.093	-0.0271	0.3494
high	0.2465	0.2110	1.170	0.243	-0.1671	0.6600
med	0.3542	0.3478	1.020	0.308	-0.3274	1.0358
low	1.1994	1.0224	1.170	0.241	-0.8045	3.2034
areasq	0.0000	0.0000	-0.660	0.509	0.0000	0.0000
sugharvsq	0.0000	0.0000	-0.850	0.396	0.0000	0.0000
sug_totharvsq	0.0016	0.0005	3.190	0.001	0.0006	0.0026
totharvsq	0.0005	0.0003	1.530	0.125	-0.0001	0.0011
pasturesq	-0.0007	0.0005	-1.360	0.175	-0.0016	0.0003
popdenssq	0.0000	0.0000	-1.910	0.056	0.0000	0.0000
gdppc80sq	0.0000	0.0000	-2.790	0.005	0.0000	0.0000
gdppc96sq	0.0000	0.0000	1.260	0.207	0.0000	0.0000
rentedsq	0.0022	0.0039	0.560	0.573	-0.0055	0.0099
occupiedsq	-0.0022	0.0115	-0.190	0.848	-0.0248	0.0204
ownedsq	0.0036	0.0103	0.360	0.723	-0.0165	0.0238
rurpopsq	-0.0008	0.0005	-1.670	0.094	-0.0017	0.0001
highsq	0.0006	0.0006	0.990	0.321	-0.0006	0.0017
medsq	0.0000	0.0004	0.020	0.985	-0.0007	0.0007
lowsq	-0.0008	0.0017	-0.460	0.649	-0.0040	0.0025
areasqrt	-0.0064	0.0252	-0.260	0.799	-0.0559	0.0430
sugharvsqrt	0.0187	0.0418	0.450	0.655	-0.0632	0.1006
sug_totharvsqrt	1.0998	0.3726	2.950	0.003	0.3695	1.8301
totharvsqrt	0.3919	0.3684	1.060	0.287	-0.3301	1.1138
pasturesqrt	-0.7121	0.8805	-0.810	0.419	-2.4378	1.0135
popdenssqrt	-0.1248	0.1349	-0.930	0.355	-0.3892	0.1395
gdppc80sqrt	-0.2331	0.0925	-2.520	0.012	-0.4144	-0.0518
gdppc96sqrt	0.0238	0.0412	0.580	0.563	-0.0569	0.1045

Appendix Table 15 (Continued)

rentedsqrt	0.4429	0.4991	0.890	0.375	-0.5353	1.4212
occupiedsqrt	-0.3802	0.5685	-0.670	0.504	-1.4945	0.7342
ownedsqrt	33.6982	60.7359	0.550	0.579	-85.3420	152.7384
rurpopsqrt	-1.3133	0.7046	-1.860	0.062	-2.6943	0.0676
highsqrt	0.1389	0.2965	0.470	0.639	-0.4421	0.7200
medsqrt	-0.1382	0.1867	-0.740	0.459	-0.5041	0.2277
lowsqrt	-0.1220	0.3658	-0.330	0.739	-0.8390	0.5950
sugharv%totharv	-0.0008	0.0011	-0.750	0.451	-0.0030	0.0014
popdens%metrop	0.0017	0.0041	0.410	0.683	-0.0064	0.0097
high%totharv	-0.0002	0.0004	-0.610	0.542	-0.0011	0.0006
high%sug_totharv	0.0002	0.0010	0.180	0.855	-0.0017	0.0021
high%sugharv	0.0000	0.0000	0.800	0.425	0.0000	0.0000
high%pasture	0.0003	0.0006	0.480	0.631	-0.0010	0.0016
high%popdens	-0.0002	0.0002	-0.820	0.412	-0.0007	0.0003
med%totharv	-0.0005	0.0003	-1.540	0.124	-0.0012	0.0001
med%sug_totharv	-0.0026	0.0010	-2.560	0.011	-0.0046	-0.0006
med%sugharv	0.0000	0.0000	2.510	0.012	0.0000	0.0000
med%pasture	0.0003	0.0003	0.880	0.378	-0.0003	0.0008
med%popdens	0.0003	0.0002	1.850	0.065	0.0000	0.0006
low%totharv	-0.0012	0.0017	-0.710	0.478	-0.0046	0.0022
low%sug_totharv	0.0044	0.0030	1.460	0.145	-0.0015	0.0103
low%sugharv	0.0000	0.0000	0.090	0.932	-0.0001	0.0001
low%pasture	0.0030	0.0013	2.210	0.027	0.0003	0.0056
low%popdens	0.0002	0.0003	0.630	0.527	-0.0004	0.0007
rented%high	-0.0027	0.0027	-1.010	0.311	-0.0080	0.0025
rented%med	-0.0007	0.0036	-0.200	0.843	-0.0077	0.0063
rented%low	-0.0107	0.0130	-0.820	0.410	-0.0362	0.0148
rented%pasture	-0.0003	0.0028	-0.110	0.911	-0.0057	0.0051
occupied%high	-0.0018	0.0051	-0.360	0.723	-0.0118	0.0082
occupied%med	-0.0070	0.0045	-1.560	0.118	-0.0158	0.0018
occupied%low	-0.0155	0.0152	-1.020	0.309	-0.0453	0.0143
occupied%pasture	0.0019	0.0030	0.640	0.524	-0.0040	0.0079
owned%high	-0.0030	0.0021	-1.430	0.154	-0.0071	0.0011
owned%med	-0.0034	0.0035	-0.970	0.331	-0.0103	0.0035
owned%low	-0.0135	0.0103	-1.310	0.190	-0.0336	0.0067
owned%pasture	0.0020	0.0025	0.820	0.415	-0.0028	0.0068
gdppc80%popdens	0.0000	0.0000	2.160	0.031	0.0000	0.0000
gdppc80%rurpop	0.0000	0.0000	1.030	0.301	0.0000	0.0000
gdppc96%popdens	0.0000	0.0000	0.030	0.972	0.0000	0.0000
gdppc96%rurpop	0.0000	0.0000	-1.430	0.151	0.0000	0.0000
constant	-119.8500	207.9239	-0.580	0.564	-527.3734	287.6733

Appendix Table 16: Assessing the balance in covariates before and after reweighting based on the propensity score – region CSex – estimating ATU – scenario 1 (at least 1% growth)

variable	control	treated	t-stat	control	treated (weighted)	t-stat (weighted)
area	1141.30	1087.49	0.50	1141.30	1070.72	0.66
high	2.92	5.02	-2.93	2.92	3.09	-0.27
med	9.89	13.41	-3.60	9.89	10.17	-0.32
low	1.47	1.62	-0.50	1.47	1.38	0.32
sugharv	382.51	454.95	-0.48	382.51	320.03	0.42
sug_totharv	4.41	6.22	-2.78	4.41	4.36	0.09
totharv	16.63	20.18	-2.92	16.63	15.80	0.80
pasture	40.96	45.58	-3.97	40.96	39.88	0.92
rented	3.84	3.85	-0.06	3.84	3.60	0.90
occupied	2.50	2.24	1.85	2.50	2.44	0.43
owned	92.13	92.54	-1.19	92.13	92.65	-1.45
popdens	65.77	44.18	2.61	65.77	52.82	1.48
rurpop	43.35	41.85	1.32	43.35	43.43	-0.06
gdppc80	4343.68	4185.40	0.81	4343.68	4288.97	0.30
gdppc96	3922.60	3756.04	1.31	3922.60	4049.43	-0.93

Appendix Table 17: Assessing the balance in covariates before and after reweighting based on the propensity score – region CSex – estimating ATU – scenario 2 (at least 5% growth)

variable	control	treated	t-stat	control	treated (weighted)	t-stat (weighted)
area	1137.64	1118.95	0.16	1137.64	1132.76	0.04
high	2.93	5.81	-3.36	2.93	3.24	-0.45
med	9.87	14.57	-4.19	9.87	10.19	-0.32
low	1.47	1.80	-0.93	1.47	1.30	0.59
sugharv	245.40	354.90	-1.91	245.40	350.24	-1.44
sug_totharv	4.29	5.32	-1.54	4.29	4.36	-0.10
totharv	16.62	20.68	-2.96	16.62	15.06	1.37
pasture	40.96	47.63	-5.26	40.96	39.39	1.21
rented	3.84	3.88	-0.16	3.84	3.50	1.19
occupied	2.50	2.24	1.71	2.50	2.38	0.79
owned	92.13	92.45	-0.86	92.13	92.78	-1.70
popdens	65.73	37.45	3.54	65.73	48.32	1.93
rurpop	43.39	41.66	1.41	43.39	44.36	-0.76
gdppc80	4344.43	4007.00	1.76	4344.43	4135.01	1.13
gdppc96	3923.39	3631.89	2.20	3923.39	3913.84	0.07

Appendix Table 18: Assessing the balance in covariates before and after reweighting based on the propensity score – region CSex – estimating ATU – scenario 3 (at least 10% growth)

variable	control	treated	t-stat	control	treated (weighted)	t-stat (weighted)
area	1143.47	1215.00	-0.53	1143.47	1109.88	0.26
high	2.94	7.51	-3.96	2.94	2.70	0.32
med	9.94	17.26	-5.12	9.94	10.23	-0.26
low	1.48	2.28	-1.75	1.48	1.56	-0.24
sugharv	246.93	275.51	-0.50	246.93	315.69	-1.00
sug_totharv	4.31	4.23	0.11	4.31	5.71	-1.49
totharv	16.65	20.05	-2.16	16.65	16.15	0.37
pasture	41.17	49.62	-5.81	41.17	41.57	-0.28
rented	3.80	3.98	-0.57	3.80	3.64	0.52
occupied	2.46	2.26	1.09	2.46	2.36	0.51
owned	92.31	92.39	-0.20	92.31	92.77	-1.11
popdens	54.36	37.95	2.57	54.36	46.53	1.21
rurpop	43.65	40.03	2.60	43.65	44.45	-0.54
gdppc80	4235.69	3901.41	2.22	4235.69	4187.56	0.25
gdppc96	3888.05	3607.36	1.93	3888.05	3906.59	-0.11

North-Northeast	**Center-South**
Tocantins	Mato Grosso
Maranhão	Goías
Piauí	Minas Gerais
Ceará	Mato Grosso do Sul
Rio Grande do Norte	Espírito Santo
Paraíba	Rio de Janeiro
Pernambuco	Paraná
Alagaos	Santa Catarina
Sergipe	Rio Grande do Sul
Bahia	São Paulo

Appendix Figure 1: Map of Brazilian states and list of states in North-Northeast and Center-South

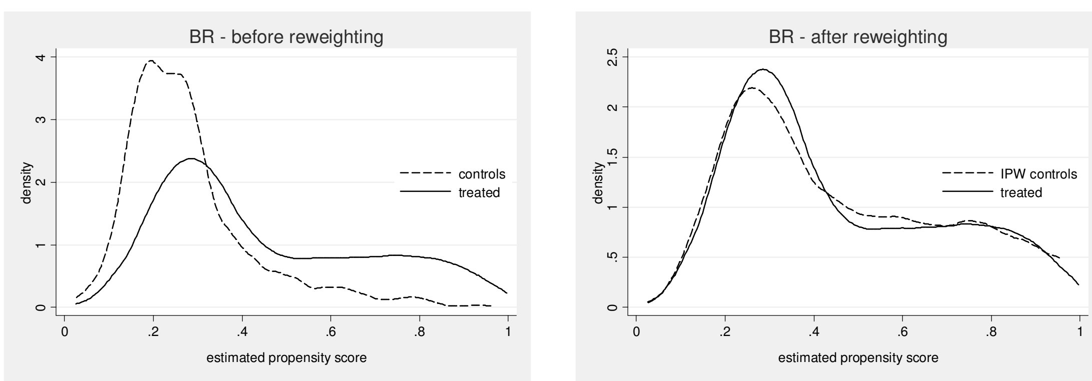

Appendix Figure 2: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (right) – region BR

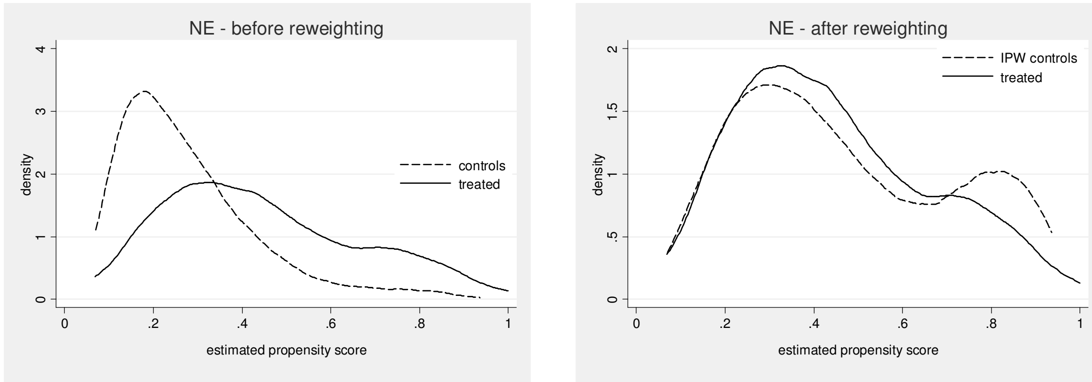

Appendix Figure 3: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (right) – region NE

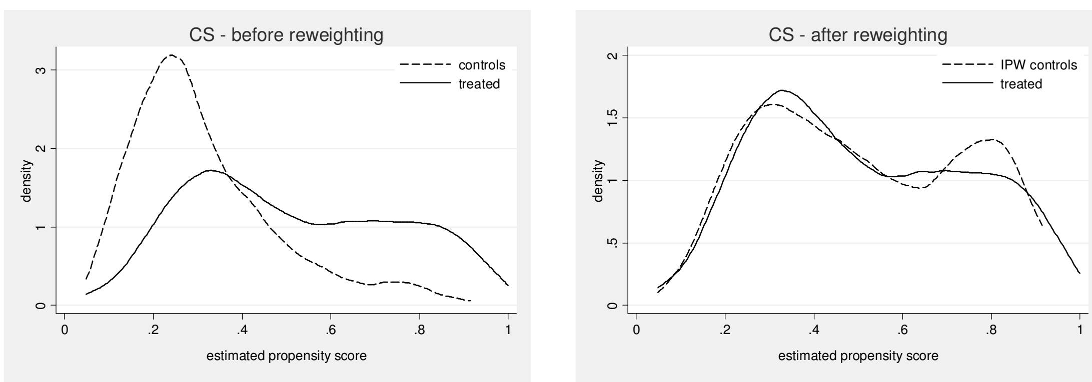

Appendix Figure 4: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (right) – region CS

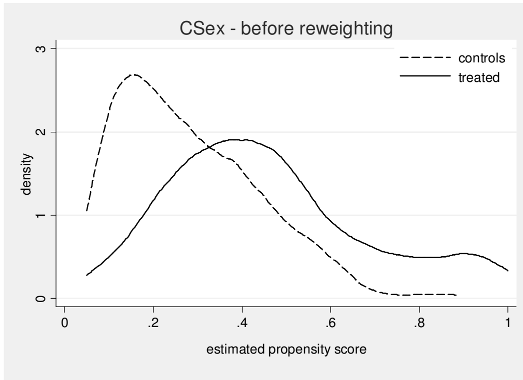

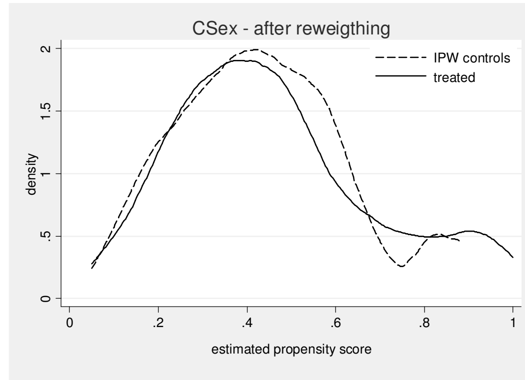

Appendix Figure 5: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (right) – region CSex

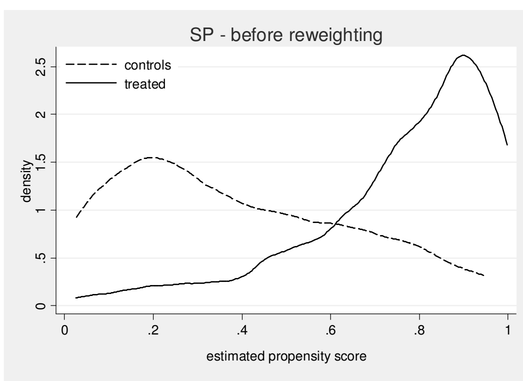

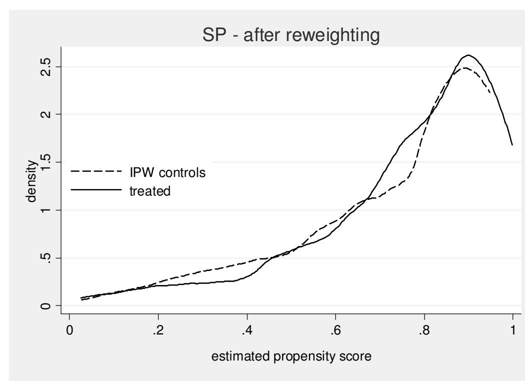

Appendix Figure 6: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (right) – region SP

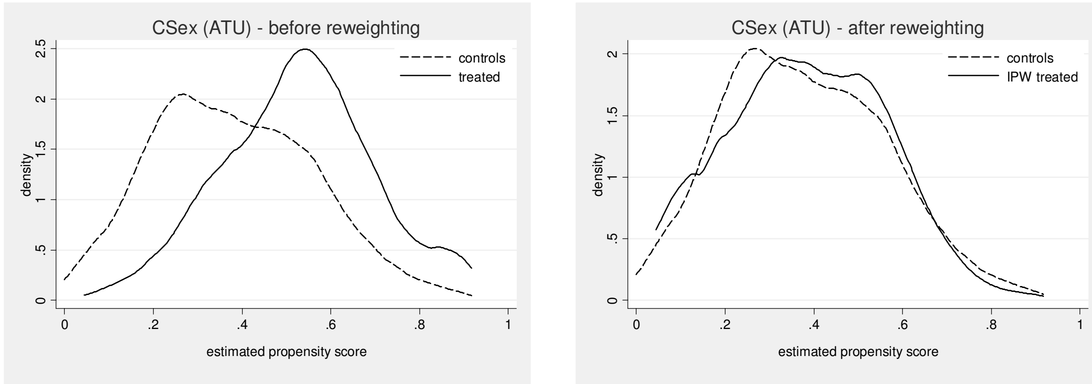

Appendix Figure 7: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (right) – region CSex – estimating ATU - scenario 1 (at least 1% growth)

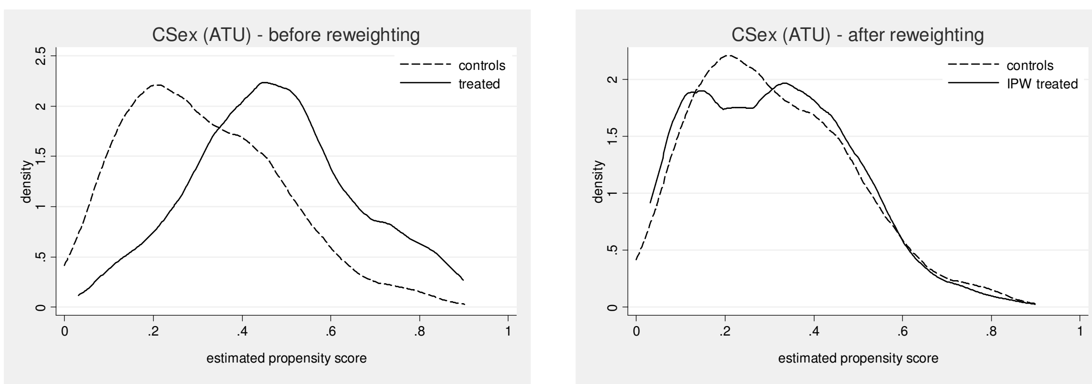

Appendix Figure 8: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (right) – region CSex – estimating ATU - scenario 2 (at least 5% growth)

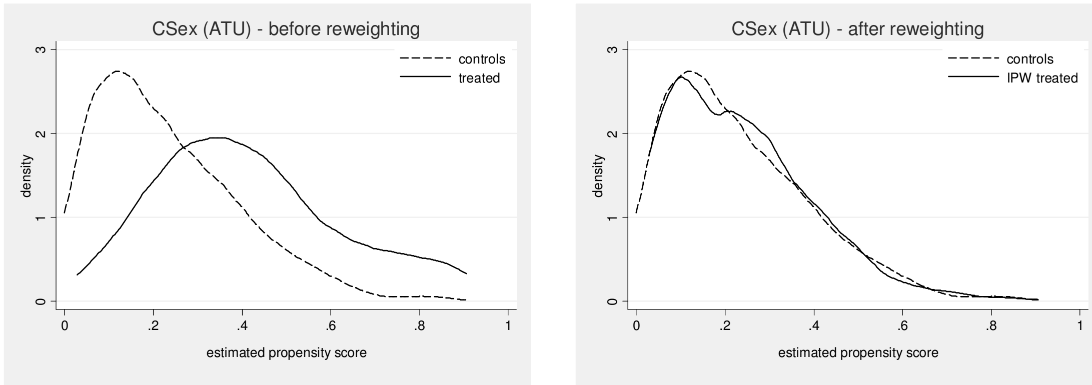

Appendix Figure 9: Kernel densities of estimated propensity scores before reweighting (left) and after reweighting (right) – region CSex – estimating ATU - scenario 3 (at least 10% growth)

REFERENCES

Altieri, A. 2008. Brazilian perspectives on biofuels. Presentation at the Business School of São Paulo. 8 Jan 2008, São Paulo. (unpublished).

Amaral, T.M. and M.F. Neves. 2003. As Cadeias Agroalimentares do Açúcar: Um Estudo Comparative entre França e Brasil. Anais do 2º Congresso do Instituto Franco-Brasileiro de Administração de Empresas, IFBEA, Franca, 26 e 27 de maio de 2003.

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

Brandão, A.S.P. 2007. The Sugar/Ethanol Complex in Brazil: Development and Future. Paper delivered at "Global Sugar Markets, Policies and Reform Options", State University of Rio de Janeiro, June 1, 2007, Reforming Global Sugar Trade: What Do Changes in Sugar Policy Mean for the Future?. International Food Policy Research Institute (IFPRI), Washington DC.

Burnquist, H., C.C. Cabral da Costa and J.J. Martins Guilhoto. 2004. Impacts of Changes in Regional Sugar and Ethanol Exports upon Brazilian overall Economy. Paper presented at the International Conference on "Input-Output and General Equilibrium: Data, Modeling and Policy Analysis", Free University of Brussels, Brussels, September 2–4, 2004.

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.

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

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.

Fredo, C.E., M.C.M.Vicente, C.S.L. Baptistella and J.E.R. Veiga. 2007. Índice de Mecanização na Colheita da Cana-de-Açúcar no Estado de São Paulo e nas Regiões Produtoras Paulistas, Junho de 2007. *Análises e Indicadores do Agronegócio* 3(28)

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. 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. 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., 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. Amsterdam: Elsevier, Vol. III, pp. 1865-2097.

Hoffman, R. and F.C.R. de Oliveira. 2008a. Remuneração, características das pessoas ocupadas na agroindústria canavieira do Brasil, 2002 a 2006. Paper presented at the Workshop "Mercado de Trabalho da Agroindústria Canavieira: Indicadores Sociais e Perspectivas Futuras", Escola Superior de Agricultura "Luiz de Queiroz" (ESALQ-USP), Piracicaba, Brazil, 25 April, 2008.

Hoffman, R. and F.C.R. de Oliveira. 2008b. Evolução da remuneração das pessoas empregadas na cana-de-açúcar e em outras lavouras, no Brasil e em São Paulo. Paper presented at the Workshop "Mercado de Trabalho da Agroindústria Canavieira: Indicadores Sociais e Perspectivas Futuras", Escola Superior de Agricultura "Luiz de Queiroz" (ESALQ-USP), Piracicaba, Brazil, 25 April, 2008.

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

IBGE (Instituto Brasileiro de Geografia e Estatística). 2008. Produto Interno Bruto dos Municípios - 2a edição. Séries Relatórios Metodológicos número 29. Rio de Janeiro, 2008.

IBGE (Instituto Brasileiro de Geografia e Estatística). 2010. Online database. Available at: http://www.ibge.gov.br. Accessed August 2009 (Chapter 1 and Chapter 3) and May 2010 (Chapter 2)

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

Imbens, G.W. and J. Wooldridge. 2007. What's New in Econometrics?. NBER 2007 Summer Institute Notes. Available at: http://www.nber.org/confer/2007/si2007/wneprg.html

IPEA. 2009. Online database. Available at: http://www.ipeadata.gov.br. Accessed August 2009 (Chapter 1 and Chapter 3) and May 2010 (Chapter 2)

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

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

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

Krivonos, E. and M. Olarreaga. 2006. Sugar Prices, Labour Income and Poverty in Brazil. World Bank Policy Research Working Paper No. 3874

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

Manzatto, C.V., E.D. Assad, J.F.Mansilla Bacca, M.J. Zaroni and S.E Marschhausen Pereira. 2009. Zoneamento agroecológico da cana-de-açucar. Rio de Janeiro: Embrapa Solos, 2009. Available online: http://www.cnps.embrapa.br/zoneamento_cana_de_acucar/ZonCana.pdf

MAPA (Ministério da Agricultura, Pecuária e Abastecimento), 2010. Relação das Unidades Produtoras Cadastradas no Departamento da Cana-de-açucar e Agroenergia. Secretaria de Produção e Agroenergia, Departamento da Cana-de-açucar e Agroenergia. Brasilia, Brasil, Posição 1 April, 2010.

Moraes, M.A. and M. Pessini. 2004. Analysis of the labor market of the Brazilian sugar alcohol sector. Mimeo. University of São Paulo, São Paulo, Brazil.

Moraes, M.A. 2007. Indicadores do mercado de trabalho do sistema agroindustrial da cana-de-açúcar do Brasil no período 1992-2005. *Estud. Econ*. 37(4):875-902.

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

Robins, J., A. Rotnitzky and L.P. Zhao. 1995. Analysis of semiparametric regression models for repeated outcomes under the presence of missing data. *Journal of the American Statistical Association* 90:106-121

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

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

Rosenbaum, P.R. and D.B. Rubin. 1983b. Assessing the Sensitivity to an Unobserved Binary Covariate in an Observational Study with Binary Outcome. *Journal of the Royal Statistical Society* 45(B):212-218.

Roy, A. 1951. Some thoughts on the distribution of earnings. *Oxford Economic Paper* 3(2):135-145

Rubin, D.B. 1974. Estimating causal effects to treatments in randomised and nonrandomised studies. *Journal of Educational Psychology* 66:688-701

Rudorff, B.F.T, D.A. de Aguiar, W.F. da Silva, L.M. Sugawara, M. Adami and M.A. Moreira. 2010. Studies on the Rapid Expansion of Sugarcane for Ethanol Production in São Paulo State (Brazil) Using Landsat Data. *Remote Sensing* 2:1057-1076

Sachs, R.C.C. 2007. Renumeração da tonelada de cana-de-açúcar no estado de São Paulo. *Informações Econômicas*, 37(2): 55-66.

SECEX/MDIC (Secretaria de Comércio Exterior/Ministério do Desenvolvimento, Indústria e Comércio Exterior). 2010. Online database. Available at: http://aliceweb.desenvolvimento.gov.br/default.asp. Accessed May 2010

Smeets, E., M. Junginger, A. Faaij, A. Walter and P. Dolzan. 2006. Sustainability of Brazilian bio-ethanol. Report commissioned by SenterNovem, the Netherlands Agency for Sustainable Development and Innovation. Unit Science, Technology and Society, Copernicus Institute, Utrecht University, Utrecht, the Netherlands.

Sparovek, G., A. Barretto, G. Berndes, S. Martins and R. Maule. 2009. Environmental, land-use and economic implications of Brazilian sugarcane expansion 1996-2006. *Mitigation and Adaptation Strategies for Global Change* 14(3):285-298.

Todd, P. 2008. Evaluating Social Programs with Endogenous Program Placement and Selection of the Treated. In T.P. Schultz and J. Strauss (eds) Handbook of Development Economics (Vol. 4, pp. 3847–94). New York: North-Holland.

Ueki, Y. 2007. Industrial Development and the Innovation System of the Ethanol Sector in Brazil. Discussion paper 109. Institute of Developing Economies

(IDE), Japan External Trade Organization (JETRO). Chiba.

Veiga Filho, A. and P. Ramos. 2006. Proálcool e evidências de concentração na produção e processamento de cana-de-açúcar. *Informações Econômicas* 36(7):48-61.

Walter, A. 2008. “Sustainability of Ethanol Production in Brazil.” Paper presented at the “Workshop internacional Brasil-Japão em biocombustível, meio ambiente e novos productos da biomassa”, UNICAMP, Campinas, Brazil, 29 October–1 November.