

Brazil in Transition: Agriculture, Labor, and Health

by

Charles Austin Davis

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2017

Doctoral Committee:

Professor Martha Bailey, Co-Chair
Professor David Lam, Co-Chair
Assistant Professor Achyuta R. Adhvaryu
Professor Paul W. Rhode

Charles Austin Davis
caustind@umich.edu
ORCID ID: 0000-0002-7450-0986

© Charles Austin Davis 2017

Acknowledgements

I am grateful for the patience, generosity, and wisdom of my chairs, Martha Bailey and David Lam. Thanks to my advisors Achyuta Adhvaryu and Paul Rhode for various and valuable lessons. I relied on many colleagues, but I am deeply moved by the support and friendship of Gaurav Khanna, Margaret Lay, and Katherine Lim. I was buoyed throughout by the love of my family. Caroline D'Angelo: you bring me joy. Você está sempre no meu coração.

TABLE OF CONTENTS

Acknowledgments	ii
List of Figures	v
List of Tables	vi
Abstract	viii
 Chapter	
1 Why Did Sugarcane Growers Suddenly Adopt Existing Technology?	1
1.1 Introduction	1
1.2 Evaluating the Regulation	4
1.2.1 Where is Sugarcane Grown?	4
1.2.2 Background on Regulation	5
1.2.3 Regression Discontinuity Evidence	7
1.2.4 Difference-in-differences Evidence	12
1.3 The Role of Wages	16
1.3.1 The Market for Unskilled Labor in Brazil, 1999–2013	16
1.3.2 Sugarcane Cultivation	20
1.3.3 Modeling the Choice of Harvest Technologies	21
1.3.4 Estimating The Effect of Wages on Mechanization	26
1.4 Conclusions	30
1.5 Appendix	32
1.5.1 Labor Demand	32
1.5.2 Regression Discontinuity Evidence: The Steepness Threshold	34
1.5.3 Supplemental Analysis for Regression Discontinuity Estimates Using Size Threshold	37
1.5.4 Difference-in-Difference Estimates Using the Size Threshold	48
2 Putting Out the Fires: Measuring the Local Effects of Air Pollution on Health	51
2.1 Introduction	51
2.2 Context and Mechanisms	53
2.2.1 Sugarcane Burning and Health in South Central Brazil	53
2.2.2 Harvesting Transition	56
2.3 Data and Descriptive Statistics	57
2.3.1 Airborne Particulates	57
2.3.2 Natality and Mortality	62
2.3.3 Hospital admissions	63
2.4 Estimation	66

2.4.1	Birth Outcomes	66
2.4.2	Infant Mortality	73
2.4.3	Hospital Admissions	75
2.5	Conclusion	77
2.6	Additional Tables and Figures	79
3	Worker Transitions in a Time of Technological Change	89
3.1	Introduction	89
3.2	Sugarcane Cultivation and the Adoption of Mechanical Harvesting	91
3.3	Data	93
3.3.1	Collection and Coverage	93
3.3.2	Variable Definitions	94
3.3.3	Sample	95
3.4	Changes in the Composition of Workers and Firms	95
3.4.1	How Did the Workforce Change During Mechanization?	95
3.4.2	Did Mechanization Encourage Consolidation? Which Firms Mechanized?	97
3.5	Mobility of Workers in the Sugarcane Industry	101
3.5.1	Employment, Fluidity, and Tenure	101
3.5.2	Where Do Sugarcane Workers Come From? Go?	104
3.5.3	How Did the Characteristics and Outcomes of Workers Change?	108
3.5.4	Who Successfully Transitions Out of Manual Sugarcane?	116
3.6	Conclusion	120
	Bibliography	121

LIST OF FIGURES

1.1	Major sugarcane producing states and the study region	5
1.2	Binned Scatter Plot and Local Linear Estimates	10
1.3	Required Mechanization as a Fraction of Land Area	15
1.4	Aggregate Employment and Real Wages from 1999–2013	17
1.5	Change in Wage and Employment 2002–2013 by Industry	18
1.6	Change in Wage and Employment 2002–2013 by Occupation	18
1.7	Employment and Wages in Sugarcane from 1999–2013	19
1.9	Empirical Support for S-Shaped Adoption	25
1.10	Quadratic and Linear IV Results	29
1.11	Labor Demand with Different Correlations ρ	34
1.12	Burning Propensity and Fitted Polynomial (2007)	36
1.13	Estimated Average Treatment Effect at the Threshold by Year	36
1.14	p-values for Continuity of the Density of the Running Variable	37
1.15	Histogram of Establishment Area (rounded to nearest hectare)	40
1.16	Point Estimates and CIs at Various Symmetric Bandwidths	41
1.17	Binned Scatter Plots and Local Linear Estimates	43
1.18	Binned Scatter Plots and Local Linear Estimates for Manual Harvesting Inputs	45
1.19	Binned Scatter Plots and Local Linear Estimates for Mechanical Harvesting Inputs	46
2.1	Major sugarcane producing states and the study region	54
2.2	The Transition to Mechanized Harvesting	56
2.3	AAI Grid and Municipio Values	60
2.4	Daily and Seasonal AAI Averaged over Study Region	61
2.5	Density for Birth Weight (2005–2012)	62
2.6	Trends and Distributions in Hospital Admissions (2000–2012)	64
2.7	CDF of Regression Residuals by Pollution Exposure	71
2.8	Quarterly Births (2005–2012)	87
2.9	All-Cause, one-year Infant Mortality Rate	87
2.10	Output and Area Cultivated for Large Sugarcane Producers	88
2.11	Output and Area Cultivated for Sugarcane-Producing Regions of Brazil	88
3.1	Sugarcane Production and Area	92
3.2	Labor intensity of sugarcane harvesting	93
3.3	Aggregate Employment in Sugarcane by Quarter	102
3.4	Distribution of Job Tenure Over Time	102
3.5	Hire and Separation Rates by Quarter	103

LIST OF TABLES

1.1	The Effect of Regulation on Labor Intensity (DiD)	14
1.2	IV Estimates of Wage Elasticity of Labor Demand	28
1.3	The Effect of Regulation on Harvesting Techniques (RD)	42
1.4	The Effect of Regulation on Input Use (RD)	47
1.5	Likelihood of Harvesting Only Manually	49
1.6	Likelihood of Harvesting Only Mechanically	49
1.7	Likelihood of Harvesting Manually and Mechanically	50
2.1	Summary Statistics for Study Region (2000)	58
2.2	Distribution of AAI by Season	59
2.3	Gestational Age	63
2.4	The Effect of AAI on Birth Outcomes	67
2.5	Distribution of Mean AAI Across All Measured Pregnancies	68
2.6	The Effect of Extreme AAI on Birth Outcomes	69
2.7	The Effect of AAI on Birth Outcomes without Muni FE	70
2.8	The Effect of AAI on Birth Outcomes without Major Metros	70
2.9	Distribution of Mean AAI by Metro Region	71
2.10	Distribution of Lagged, Concurrent, and Lead AAI	73
2.11	Effect of AAI on Mortality Rate and Cohort Size	74
2.12	The Effect of AAI on Hospital Admissions	76
2.13	Largest Sugarcane Growers	79
2.14	Area under cultivation (millions of Ha)	80
2.15	Output (millions of Tonnes)	80
2.16	Total revenue (billions of Reais - nominal)	81
2.17	Median Municipio-level Yield (tons / Ha)	81
2.18	Robustness Checks for Mortality	82
2.19	The Effect of AAI on Birth outcomes (all coefficients)	83
2.19	The Effect of AAI on Birth outcomes (all coefficients)	84
2.19	The Effect of AAI on Birth outcomes (all coefficients)	85
2.19	The Effect of AAI on Birth outcomes (all coefficients)	86
3.1	Manual worker characteristics by year	96
3.2	Mechanical worker characteristics by year	96
3.3	Other worker characteristics by year	97
3.4	Firm-level workforce characteristics by firm employment	99
3.4	Firm-level workforce characteristics by firm employment (<i>continued</i>)	100
3.5	Occupation and Industry Transition Rates for Manual Sugarcane Workers (period means)	105
3.6	Industries and Occupations in $t - 1$ for Manual Cane Workers in t ($t \in [1999, 2007]$)	106

3.7	Industries and Occupations in $t - 1$ for Manual Cane Workers in t ($t \in [2008, 2014]$)	107
3.8	Industries and Occupations in $t + 1$ for Manual Cane Workers in t ($t \in [1998, 2006]$)	107
3.9	Industries and Occupations in $t + 1$ for Manual Cane Workers in t ($t \in [2007, 2013]$)	108
3.10	Characteristics of Manual Sugarcane Workers in t by Employment in $t - 1$	110
3.11	Characteristics of Manual Sugarcane Workers in t by Employment in $t + 1$	111
3.12	Change in Characteristics of Manual Cane Workers $Y_t - Y_{t-1}$ by Employment in $t - 1$	114
3.13	Change in Characteristics of Manual Cane Workers $Y_{t+1} - Y_t$ by Employment in $t + 1$	115
3.14	Transition Probabilities by Individual and Municipality Characteristics (1998–2006)	118
3.15	Transition Probabilities by Individual and Municipality Characteristics (2007–2013)	119

ABSTRACT

Brazil in Transition: Agriculture, Labor, and Health

by

Charles Austin Davis

Chair: Martha Bailey and David Lam

My dissertation studies major transitions in labor, agriculture, and health in Brazil. In recent decades, Brazil moved from a relatively poor, highly unequal country to an increasingly prosperous and egalitarian middle-income country. These successes, combined with the breadth and detail of available data, offer the opportunity to discover new insights about process of development. In each of three chapters, I focus on different but related topics in the growing Brazilian economy: the adoption of mechanical harvesting in one of Brazil's major crops, measuring the effect of air pollution on health, and the transitions of workers exposed to technological change. I combine confidential, country-wide microdata with modern statistical methods to overcome identification challenges and to draw lessons from the experience of Brazil. My results emphasize that market forces can lead to technology adoption when government intervention does not, that remote sensing data is an inadequate substitute for ground-based pollution monitoring, and that local labor market conditions are primary determinants of workers' success in coping with technological displacement.

CHAPTER 1

Why Did Sugarcane Growers Suddenly Adopt Existing Technology?

I investigate the role of regulation and factor prices in the rapid, widespread adoption of mechanical harvesting technology by Brazilian sugarcane growers. I use worker- and establishment-level data to test the effect of regulation using complementary regression discontinuity and difference-in-differences approaches. I find that regulation is, at best, a partial explanation, accounting for no more than one quarter of the dramatic change in harvesting practices. I develop a tipping-point model to show how rising wages may have played an important role even though the change in wages was gradual and the change in harvesting was abrupt; instrumental variables estimates imply that increasing wages alone are sufficient to explain the adoption of green technology.

1.1 Introduction

In the developing world, a large fraction of the population engages in low-productivity agriculture (Gollin, 2010). Widely available technologies like fertilizer have the potential to improve productivity but are not widely used (Morris et al., 2007). Following this pattern, Brazilian sugarcane growers relied on manual harvesting as late as 2007 even though harvesting machines had been available for many years. But, by 2013, almost all sugarcane was harvested mechanically.¹ This paper studies how this dramatic transformation was achieved.

¹Adoption was relatively quick in this case, especially considering that the technology was not new. It took about 22 years for American farmers to completely adopt diesel tractors (White). Mansfield (1961) studies the diffusion of 12 innovations of “outstanding importance” among major firms in several American industries; the average time between initial use and complete adoption is over 18 years. As another point of comparison, den Bulte (2000) finds that widespread consumer adoption of durable goods typically takes 7 to 14 years, depending on characteristics of the good and on the economic and demographic environment at the time of introduction.

Recent literature has emphasized several different factors affecting technology adoption in agriculture, from social learning to behavioral biases (Duflo et al., 2011; Conley and Udry, 2010). From interviews with various stakeholders, I identify two other factors of primary importance in Brazilian sugarcane. The first, government regulation, is understudied in this literature and the second, factor prices, is foundational. Manual harvesting was gradually banned by state governments because of pollution associated with the practice. This regulation coincides with the mechanization of harvesting, lending credence to the government's claim that the regulation caused growers to adopt mechanical harvesting. However, I find that the regulation accounts for little of the change in harvesting practices. I find that a strong labor market, where rising real wages made manual harvesting more expensive, is sufficient to explain mechanization.

Both publicly and in interviews, state government advertises regulation as the causal factor driving the adoption of machine harvesting in sugarcane. Between 2002 and 2014, all of the sugarcane-growing states responded to constituent concerns about air pollution by passing gradual bans of the straw burning associated with manual harvesting. These bans, along with enforcement efforts, were covered in national newspapers and they coincided with the period of rapid mechanization.² In various communications, environmental regulators in the largest sugarcane growing state, So Paulo, took credit for mechanization.³ So Paulo's environmental ministry has a detailed website attesting to the success of the regulation and Brazil's space agency publishes satellite-based monitoring data online.⁴

I test the effects of regulation on harvesting practices using two highly detailed, confidential data sources that provide near universal coverage of the industry. The first data source, known by its Portuguese-language acronym RAIS, captures detailed employment information for all formal-sector workers in Brazil from 1998 to 2014.⁵ The second data source is the 2006 Census of Agriculture. I use these data sources to conduct complementary tests of the regulation.

Regression discontinuity estimates from the Census of Agriculture show a small effect of regulation on harvesting practices. I take advantage of an area threshold that exempted small growers to provide complementary evaluations of the regulation, comparing harvesting techniques and input use between unregulated establishments just below the area threshold to regulated establishments just above.⁶ If the true effect lies at the extreme of the confidence interval, I find that regulation explains no more than a quarter of the change in harvesting practices. Moreover, regulated farmers show no changes in input use

²For reporting on passage of the law, see, e.g., Spinelli (2001); Osse (2002); *Gazeta Mercantil* (2002). For coverage of enforcement efforts, see, e.g. Spinelli (2001); Samora (2006); Credendio (2008); *Folha de S. Paulo* (2008); Henrique (2008); Coissi (2008); Tomazela (2016); *O Globo G1* (2016).

³For example, I recently received an email from environmental regulators claiming "As a result of this policy ... 83% of the 2013/2014 crop was harvested without burning." Officials made similar claims during in-person interviews.

⁴See <http://www.ambiente.sp.gov.br/etanolverde/> and <http://www.dsr.inpe.br/laf/canasat/>.

⁵RAIS stands for *Relaio Anual de Informaes Sociais*, which roughly translates to Annual Report of Social Information. For historical reasons, most sugarcane workers are formal.

⁶As of 2006, roughly 80 percent of establishments are not regulated as they fall below the 150 hectare threshold. However, regulated establishments control about 80 percent of sugarcane area.

that would be consistent with mechanization.

The regression discontinuity estimates measure the behavior of establishments near the area threshold; I find a similarly small effect of regulation using a difference-in-differences approach which captures the behavior of larger growers. Specifically, I use data from RAIS to estimate how changes in the stringency of the regulation across states and over time affected changes in county-level labor intensity. Here again, if the true effect lies at the extreme of the confidence interval, I find that regulation can account for at most one quarter of the observed decline in labor intensities.

Besides regulation, what caused the rapid transition of harvesting techniques? My detailed interviews with industry participants suggest that rising wages were an important motivation for mechanized harvesting. Using administrative and survey data, I show that, from 1998 to 2014, increasing labor demand from large sectors like construction helped drive real wages up by almost 50 percent for harvest workers. However, this timing does not obviously support wages as a driver of mechanization. Wages rose continuously from 1998 while widespread mechanization began only in 2007.

I develop a tipping-point model of grower behavior that reconciles these dynamics. For each parcel of land, there is a threshold wage, above which a profit-maximizing grower will harvest the parcel mechanically. This threshold can be different for each parcel, depending on characteristics of the land. If wages are well below the threshold for a majority of parcels, wages may increase steadily without affecting harvesting techniques. Eventually, as wages rise, they will cross the switching thresholds of many parcels, causing widespread mechanization.

To test wages as an explanation for mechanization, I estimate the wage elasticity of labor demand in sugarcane using a set of instruments similar to the instrument developed by Dube and Vargas (2013). These instruments capture shifts in county-level agricultural labor supply. Specifically, for each of the four other crops grown in the sugarcane region, I interact the historical county-level acreage of that crop with a measure of the crop's international price. As a first stage, these instruments predict county-level sugarcane wages. I estimate the wage elasticity of labor demand by regressing the county-level quantity of labor on the predicted wages from the first stage. The estimated elasticity is large, suggesting that the observed increase in wages is sufficient to explain the mechanization of sugarcane harvesting.

As countries become richer, fewer people work in an increasingly productive agricultural sector. This relationship holds across countries and within a country over time; its ubiquity makes it one of the fundamental facts of development (Gollin, 2010). But the nature of this relationship remains uncertain. Do improvements to agricultural productivity stimulate growth in other sectors? Or vice versa? What is the role of policy? Do other factors drive change?

Brazilian sugarcane is fascinating because we observe this development process: a large agricultural sector transitioning from a labor-intensive, low-technology production to capital-intensive, high-technology production. I find that government regulation had an extremely limited role, at least in this context. I

provide evidence that growth in other sectors, operating through increased labor demand, motivated sugarcane growers to adopt new technology. These results suggest that, while agriculture is a large part of developing economies, creating labor-market opportunities in other sectors may be an effective way to improve productivity and induce technology adoption in agriculture.

This paper owes much to prior work but offers novel insights by disentangling government and market forces. The research question is similar to various studies of mechanization among farmers in the antebellum United States (David, 1966; Olmstead and Rhode, 1993, 1995). These papers also discuss the importance of factor prices but, beyond differences in time, location, crop and technology, mechanization was not regulated in the historical US. There is also a deep literature on technological adoption in developing country agriculture (Foster and Rosenzweig, 1995; Conley and Udry, 2010; Duflo et al., 2011). However, learning and network effects are primary in these papers and, thanks to the sophistication of Brazilian sugarcane growers and the provision of extension services, these issues are less important here. Finally, prior work has shown that it is difficult to end polluting practices in the developing world; both government regulation and NGO-backed technological fixes have failed (Davis, 2008; Greenstone and Hanna, 2014). At least in this context, reductions to pollution were achieved through development itself, as embodied by higher wages for some of Brazil's poorest workers.

1.2 Evaluating the Regulation

1.2.1 Where is Sugarcane Grown?

Although sugarcane is grown in two separate regions of the country, I focus on a group of six contiguous states in South-Central Brazil. These states, which I refer to as the “study region,” account for about 80 percent of Brazil's output and experienced the changes in harvesting practices that motivate this paper.⁷ By contrast, the smaller, two-state sugarcane growing region in the Northeast has long used a more rudimentary form of mechanized harvesting due to its history and ecology.⁸ Figure 2.1 is a map of Brazil that indicates the study region.

The study region is home to about half of Brazil's 200 million inhabitants and is a major producer of sugarcane, coffee, oranges, soybeans, and maize. From 2006-2010, the study region accounted for 32 percent of world sugarcane production, 21 percent of world coffee production, 23 percent of world orange production, and 4 percent of world maize production.⁹ According to household survey data, each crop

⁷The study region accounted for 80 percent of total sugarcane tonnage between 1990 and 2010, according to the Brazilian Census Bureau's *Produção Agrícola Municipal (PAM)* data.

⁸The six states in the study region are Goiás, Minas Gerais, Paraná, Mato Grosso do Sul, Rio de Janeiro, São Paulo. The two other sugarcane-producing states are Pernambuco and Alagoas, both located in the northeast region. Pernambuco and Alagoas do not have the optimal soil and climate for growing sugarcane and, according to anecdotal reports, the otherwise unprofitable industry endures thanks to subsidies from the state government.

⁹Brazilian output from PAM. World output from FAO.

employs between 0.5 and 2.5 million workers.¹⁰

Figure 1.1: Major sugarcane producing states and the study region



1.2.2 Background on Regulation

Sugarcane fields are burned in preparation for manual harvesting but not mechanical harvesting. The resulting pollution motivated state governments to restrict pre-harvest burning, effectively mandating mechanization. By eliminating pests and extra vegetative matter, burning allows harvest workers to move more quickly through the fields, approximately tripling their productivity. Following constituent complaints about salient pollution and health concerns, the state of So Paulo passed a gradual ban of pre-harvest burning in 2002. Under the regulations, property owners are permitted to burn only a fraction of each property.¹¹ In 2002, owners were permitted to burn 80 percent of each property. The regulation scheduled future reductions to this fraction: 70 percent in 2006, 50 percent in 2011, 20 percent in 2016,

¹⁰Calculated from PNAD.

¹¹The regulation applied to each property as listed in the local property registry.

and 0 percent in 2021. This regulation became more strict in 2007, advancing all target reductions, with the complete cessation of burning required by 2014.¹² Between 2008 and 2014, the five neighboring states passed similarly structured regulation.

The regulations included meaningful incentives to change grower behavior. Violations of the So Paulo regulation could be punished by large fines. The regulation demands that growers pay a fine for each hectare burned in excess of their allowed fraction which, in 2002, amounted to 13 percent of average per-hectare revenue. The fine was revised upward every year, roughly tracking inflation. Additionally, failure to comply with burning restrictions jeopardized mandatory state environmental licenses. Enforcement strategies vary across the other states. Some impose large fines and threaten jail time while others instead offer incentives for compliance.

Anecdotal evidence suggests that the regulations were salient and enforced. National media publicized the passage of the regulations themselves while also documenting enforcement efforts. The So Paulo government partnered with the Brazilian space agency INPE to monitor harvesting practices via satellite images; this technique has been used to identify and fine violators. Citizens and journalists also reported violations of the law. Small and large growers have received fines throughout the years, according to newspaper reports. In these cases, amounts ranged from \$3,000 for a small grower to over \$1 million for a large grower.

Regulators claim these regulations were successful in changing harvesting practices while sugarcane growers use compliance to advertise their environmental stewardship. In verbal and written communications, environmental officials in So Paulo attribute mechanization to the regulation. The state of So Paulo maintains a website about the regulation which makes a number of claims, including: “with [regulation], all the mechanizeable area will be harvested ... without burning,”¹³ that the regulation avoided millions of tons of pollutants,¹⁴ and that the regulation quadrupled the number of harvesting machines in the state.¹⁵ The Brazilian sugarcane industry’s English-language website advertises sustainable practices, writing that “[m]echanization already exceeds 90 percent of the harvest in So Paulo, Brazil’s top cane-producing state. It will be the only means of harvesting in So Paulo by 2017, thanks to [regulation].”¹⁶

Two features of the regulations allow me to evaluate these claims. First, the regulation in So Paulo was much less stringent for growers with less than 150 hectares. I use this variation to evaluate the So Paulo regulation via regression discontinuity. Second, the stringency of regulation varies across states and over

¹²Technically, the 2007 revision to the law was a voluntary agreement between sugarcane growers and the environmental regulator. The environmental regulator held leverage over sugarcane growers in the form of environmental licensing, other regulations, and lawmakers’ threats of stricter legislation. The two parties agreed to more aggressive restrictions on burning while avoiding an uncertain and costly legislative process.

¹³<http://www.ambiente.sp.gov.br/etanolverde/protocolo-agroambiental/ganhos-ambientais/>

¹⁴<http://www.ambiente.sp.gov.br/etanolverde/files/2016/06/Etanol-Verde-Relatorio-Safra-15-16.pdf>.

¹⁵Ibid.

¹⁶<http://sugarcane.org/sustainability/best-practices>

time, allowing me to evaluate the regulation in a difference-in-differences framework.

1.2.3 Regression Discontinuity Evidence

Taking advantage of a size threshold built into the regulation, I estimate the effect of the regulation on harvesting practices via regression discontinuity.

1.2.3.1 Measuring Size and Harvesting Practices

The 2006 Census of Agriculture is a rich source with which to evaluate the regulation because the data are disaggregated, include a great breadth of information, and offer near-universal coverage. The Census of Agriculture records a variety of information about every agricultural establishment in Brazil, including their location, size, the crops grown, and harvesting technique. Thus, a researcher can identify regulated growers and study a range of relevant outcomes. The Brazilian Census Bureau endeavors to survey every agricultural establishment in the country. Finally, since these data describe each agricultural establishment, the unit of observation corresponds to the decision-making unit.¹⁷

The primary outcome in this analysis is an indicator variable for harvesting practices that should respond discontinuously to the regulation. Specifically, the primary outcome assumes a value of one if the establishment used manual harvesting only and zero if the establishment used any mechanical harvesting. Recall that the regulation required every establishment to harvest 20 percent of the land mechanically, so compliant establishments should have a zero. This measure may overstate compliance since it cannot distinguish establishments that mechanize some fraction below the required 20 percent.

While the Census does not record a continuous measure of mechanization, several continuously measured inputs serve as secondary outcomes in Appendix 1.5.3.3. I consider five machine-related outcomes: expenditure on contracting services, fuel expenditure, the number of harvesting machines, machine rental expenditure, the value of all vehicles. I also consider three labor-related outcomes: days paid to temporary workers, the number of temporary workers, and the total number of workers.

The assignment variable is reported establishment area. The Census does not collect administrative data on establishment area, so the Census variable may not be correctly identify treatment status. As such, I test for strategic manipulation of the assignment variable but find no evidence of such manipulation. See Appendix 1.5.3.1 for details.

At the time of the 2006 Census, So Paulo establishments had been regulated for the three previous years, but no other state had yet introduced regulation. Consequently, this analysis focuses exclusively on

¹⁷Given the level of detail and disaggregation, these data are confidential. To access them, I traveled to a Census Bureau facility in Rio de Janeiro. I was permitted to analyze the data only in a secure room and I was only allowed to remove programs and output files, all of which were inspected by Census Bureau employees to ensure the anonymity of respondents.

establishments in the state of So Paulo. Unfortunately, it is not feasible to incorporate additional years since the 2016 Census has not been completed and there are substantive differences between the 1996 and 2006 Censuses.

1.2.3.2 RD Estimation

The regression discontinuity design compares So Paulo establishments above and below the 150 hectare regulatory threshold. This exercise yields the average treatment effect at the threshold. To the extent that the regulation has different effects for different size establishments, this parameter is most informative about establishments that are close to the 150 hectare threshold.¹⁸

Using the method described in Calonico et al. (2014), I estimate the treatment effect of the regulation as

$$\begin{aligned}\tau &= \lim_{x \rightarrow 150^+} \left(\mathbb{E}[Y_i | X_i = x] \right) - \lim_{x \rightarrow 150^-} \left(\mathbb{E}[Y_i | X_i = x] \right) \\ &= \mu_{Y^+} - \mu_{Y^-}\end{aligned}$$

where Y is an outcome variable, x is establishment area, and the threshold is 150 hectares. Separate local polynomials of degree p are estimated above and below the threshold:

$$\begin{aligned}\hat{\tau}_p(h_n) &= \hat{\mu}_{+,p} - \hat{\mu}_{-,p} \\ \hat{\mu}_{+,p} &= e_0 \hat{\beta}_{+,p}(h_n) \\ \hat{\mu}_{-,p} &= e_0 \hat{\beta}_{-,p}(h_n)\end{aligned}$$

with

$$\begin{aligned}\hat{\beta}_{+,p}(h) &= \arg \min_{\beta} \sum_{i=1}^n \mathbb{1}(X_i \geq 0) \{Y_i - r_p(X_i)' \beta\}^2 K_{h_n} X_i \\ \hat{\beta}_{-,p}(h) &= \arg \min_{\beta} \sum_{i=1}^n \mathbb{1}(X_i < 0) \{Y_i - r_p(X_i)' \beta\}^2 K_{h_n} X_i\end{aligned}$$


where $r_p(x) = (1, x, \dots, x^p)'$, $e_0 = (1, 0, \dots, 0)$, e K_{h_n} is a kernel function with a series of bandwidths h_n .

In choosing the parameters of the analysis, I follow the recommendations of Calonico et al. (2014) but varying these choices does not substantively alter the results. Specifically: i) bandwidths for the point

¹⁸In Appendix 1.5.4, I include difference-in-differences estimates which compare the difference between So Paulo establishments above and below the the regulatory threshold, to the same difference among establishments outside of So Paulo. The resulting parameter may be interpreted as an average treatment effect on the treated. The results from this approach are substantively similar to the RD results.

estimate are selected using the method developed by Calonico et al. (2014), which minimizes mean squared error, ii) bandwidths for the bias correction are selected using the method developed by Calonico et al. (2016), which minimizes coverage error, iii) I use the triangular kernel function, iv) I estimate local linear regressions ($p = 1$), and v) variance is estimated using a nearest-neighbor approach clustered by municipality. I use separate bandwidths on either side of the cutoff because the density of X is decreasing in this region; forcing symmetric bandwidths results in many observations below the threshold and few above. Finally, I report a 95% confidence interval instead of an estimated standard error because correct inference requires that the confidence interval be recentered to account for misspecification bias.

Figure 1.2: Binned Scatter Plot and Local Linear Estimates



Source: 2006 Census of Agriculture.

Data from 2006 Agricultural Census. Local linear estimates are plotted for $x \in [150 - h_-, \min\{300, 150 + h_+\}]$. Bin sizes set to mimic the variance of the underlying data (see Calonico et al. (2015)).

The regression discontinuity estimates in Figure 1.2 show small, insignificant declines in exclusive manual harvesting. Growers just above the area threshold, i.e. regulated growers, are about 4 percentage points less likely to use manual harvesting only. The magnitude of the estimates suggests that regulation accounts for less than a quarter of the near-complete mechanization observed by 2014. According to the law, no regulated establishment could rely exclusively on manual harvesting. Since more than 60 percent of regulated establishments report manual harvesting only, a regulation that worked as intended would have a much larger treatment effect. More importantly, we observe that virtually all sugarcane harvesting was mechanized by 2014. At the extreme of the confidence interval, the estimates admit the possibility that the regulation caused about 14 percent of growers to shift away from using manual harvesting only. Fourteen percent is small compared to both the design of the regulation and observed changes in behavior.

Even these small estimates may overstate the effect of the regulation since the binary outcome does not capture the extent of mechanization. If the regulation causes some establishments to mechanize, then these establishments may meet or exceed the fraction required by regulation. In this case, the regulation would still account for less than a quarter of the near-complete mechanization observed in later years. But these establishments may instead mechanize a small fraction of their land. I attempt to measure the extent of mechanization by estimating the regression discontinuity using various agricultural inputs as continuous proxy measures of mechanization. These estimates, detailed in Appendix 1.5.3.1, are noisy but show no evidence that regulated establishments used inputs differently than unregulated establishments.

Appendix 1.5.3.1 contains supplementary analysis, including estimates using a range of bandwidths and density tests.

Overall, the evidence shows that regulated establishments were slightly less likely to use manual harvesting alone. However, the regulation is far from sufficient to explain the rapid, widespread change in harvesting techniques.

In interpreting these results, it is important to acknowledge their limitations. These estimates could miss an effect of the regulation under certain circumstances. For example, suppose a small number of large growers switched harvesting techniques because of a targeted government enforcement effort. If the targeted growers are far above the threshold area, they are unlikely to affect the regression discontinuity estimates. However, if these growers control a large fraction of total sugarcane area, they could drive large-scale changes in harvesting practices. This is a real possibility since half of all sugarcane area is controlled by the largest 2 percent of establishments.

In the next section, I address this concern by combining a continuous, size-weighted measure of mechanization with an identification strategy that does not rely on a comparison of establishments near the 150 hectare threshold.

1.2.4 Difference-in-differences Evidence

1.2.4.1 Constructing Labor Intensity from Administrative Labor Data and Survey Data on Sugarcane Cultivation

I combine two data sources to construct a continuous, size-weighted measure of mechanization. I measure mechanization as the hours of labor supplied by manual laborers in the sugarcane industry (L) divided by the total area devoted to sugarcane cultivation (T).¹⁹ Labor intensity is calculated for each municipio in each year.²⁰ Aggregating to the municipio level effectively weights mechanization by establishment size; if regulators change the behavior of the largest growers, this change might not be apparent when comparing growers near the 150 hectare threshold, but it should be visible in aggregate labor intensity. The numerator is drawn from confidential, administrative micro data maintained by the Brazilian Ministry of Labor.²¹ Known by its Portuguese-language acronym RAIS, the micro data are compiled annually and comprise the universe of formal employment.²² ²³ Each record in the dataset corresponds to an employment spell, providing information about the employee and the employer, including the municipality of work, wages, the duration of employment, hours worked, plus detailed industry and job classifications. For each municipio-year, I measure the quantity of labor L as the sum of hours worked by manual laborers in the sugarcane industry. The denominator T is drawn from the PAM survey conducted by the Brazilian Census Bureau IBGE; IBGE employees contact local producers and other industry participants to determine the land area devoted to sugarcane cultivation in each municipio-year.

1.2.4.2 Identifying Variation from Cross-Sectional and Time-Series Differences in Regulation

Cross-sectional and time-series variation in the stringency of the regulation enable its evaluation via a differences-in-differences approach. The first regulation was introduced in So Paulo in 2002 and became more stringent over time. Other states introduced similarly-structured regulation between 2008 and 2014. The fraction of land that growers were permitted to burn or, equivalently, the fraction of land they were

¹⁹I exclude most workers involved in mechanical harvesting by limiting the analysis to workers whose occupation is “manual laborer.”

²⁰In terms of area, population, and governance, Brazilian municipios are roughly equivalent to US counties.

²¹I combine two variables to calculate L : the length of an employment spell in months and the contracted hours per week. In practice, the vast majority of sugarcane workers report 44 hours per week so the variation is primarily driven by the length of employment.

²²Owing to the sensitive, identifiable information stored in RAIS, these data are confidential. I obtained permission from the Brazilian Ministry of Labor to store and analyze the data at a secure facility maintained by the University of Michigan.

²³In interviews, farmers and farm workers indicate that labor in the sugarcane sector is predominantly formal sector and unionized. As a consequence, RAIS captures roughly 60 to 75 percent of the sugarcane and coffee employment recorded in household survey data. Direct comparisons between household survey data, the PNAD, and RAIS are complicated for a several of reasons: i) the unit of analysis in each dataset is different, ii) quantity of labor is measured differently, and iii) many sugarcane workers are seasonal migrants. The figure I report here, 60 to 75 percent, is the national count of employment spells from RAIS divided by national count of individuals from PNAD. I use the national counts because RAIS records the place of work and PNAD records the place of residence. For seasonal migrants, these will not be the same so a single individual might appear in different places in each dataset.

required to mechanize, varies across states and over time. Figure 1.3 shows the mechanization requirement in each state from 1999 to 2028. A value of zero, shaded red, means there is no mechanization requirement. A value of one, shaded green, means one hundred percent of property area must be mechanized. I use this mechanization requirement (Pct) as the treatment variable in a continuous differences-in-differences design. This design assumes that changes in states with no change in stringency provide a counterfactual for states with a change in stringency.

1.2.4.3 Difference-in-differences Estimation

If regulation caused mechanization, we would expect the labor intensity of sugarcane harvesting to fall in states where the regulation became more stringent as compared to neighboring states without changes in regulation. I estimate the effect of the regulation on labor intensity using a differences-in-differences approach adapted for a continuous treatment variable, namely the stringency of regulation:

$$(L/T)_{j,s,t} = \delta_s + \gamma_t + \omega \text{Pct}_{s,t} + \varepsilon_{j,s,t}, \quad (1.1)$$

where j indexes municipality, s indexes state, and t indexes year. The outcome L/T is labor intensity, defined as hours of labor contributed by manual workers in the sugarcane industry divided by area harvested. The fixed effects δ and γ capture state-specific and year-specific unobservables. Some specifications include municipality fixed effects instead of state fixed effects. Finally, Pct measures the required fraction of mechanization; these values are shown in Figure 1.3. The coefficient of interest is ω , which captures how the changing stringency of the regulation affects labor intensity. Some specifications include a lag and lead of Pct to capture anticipatory or delayed responses to the regulation.

The results, presented in Table 1.1, suggest that regulation can explain no more than a quarter of the observed decline in labor intensity. The reported estimates show how much labor intensity would change, in terms of hours per hectare, moving from unregulated harvesting to a complete ban on burning. The point estimates from columns (1) and (2) suggest that a complete ban on burning would actually increase by about 20 hours per hectare. Assume the true effect of the regulation lies at the lower bound of the 95 percent confidence interval in column 1, i.e. $\omega = 21.43 - 1.96 \times 21.22 = -21.2$. The average value of Pct, regulation stringency, was 0.573 in 2013. This implies that regulation reduced labor intensity by $0.573 \times 21.2 = 12.2$ hours per hectare in 2013. This amounts to one quarter of the observed reduction; in aggregate, labor intensity declined by 47 hours per hectare between 2007 and 2013.

The estimates are stable with respect to controls, offering some hope that the estimated effect of the regulation is not biased by omitted variables. The estimated effect of the regulation is essentially unaffected by municipality fixed effects.²⁴ The estimates are also substantively similar when controlling for state-

²⁴Adding controls for municipio wages for manual laborers in sugarcane also has no affect on $\hat{\omega}$. Results available on request.

Table 1.1: The Effect of Regulation on Labor Intensity (DiD)

	(1)	(2)	(3)	(4)
	L / T	L / T	L / T	L / T
Pct_{t-1}			4.700 (17.97)	2.801 (18.24)
Pct_t	21.43 (21.22)	21.03 (21.68)	40.00 (20.89)	39.99 (21.06)
Pct_{t+1}			-2.830 (20.39)	1.315 (20.95)
N	8,263	8,263	6,097	6,097
\bar{y}	82.0	82.0	84.7	84.7
Muni FE		Y		Y

Pct is the legal mechanization requirement.

Quantity of labor (L) from RAIS.

Area harvested (T) from PAM.

Outcome Winsorized at the 1st and 99th percentiles.


SEs clustered by municipality.

specific time trends (results omitted; available on request). The point estimates are larger in columns 4-6, which include a lag and a lead of regulatory stringency. The larger estimates may result from the smaller sample, since observations from 1999 and 2013 are omitted. In any case, they are not statistically different from the corresponding estimates in columns 1-3.

These results suggest that regulation cannot explain more than a quarter of the observed decline in labor intensity. The point estimates from columns (1) and (2) suggest that moving from 0 to a 100 percent mechanization requirement would actually increase labor intensity by about 20 hours per hectare.

Compared to the regression discontinuity, the difference-in-differences analysis estimates a conceptually different parameter from a separate data, but both approaches lead to the same conclusion: regulation accounted for a small fraction of the change in harvesting techniques. The difference-in-differences estimates can be thought of as an average treatment effect on the treated, measuring how labor intensity responded to regulation, on average, in regulated counties. That, together with the use of municipio labor intensity as an outcome, makes this approach sensitive to the behavior of large growers. Still, the effect of the regulation appears limited.

Figure 1.3: Required Mechanization as a Fraction of Land Area



1.3 The Role of Wages

While regulation may have played some role, the mechanization of sugarcane remains largely unexplained. I turn now to the role of wages. I begin by describing the labor markets that produced a 50 percent increase in real wages for manual laborers in sugarcane between 1999 to 2013. I then present a model that shows how the tipping-point behavior observed in the data can emerge even in a frictionless, full-information environment. Finally, I estimate sugarcane growers' responsiveness to wage changes using an instrumental variables strategy.

1.3.1 The Market for Unskilled Labor in Brazil, 1999–2013

The late 1980s and early 1990s were a period of political change and economic uncertainty in Brazil. The country emerged from military dictatorship in 1989 and the first democratically elected president was impeached for corruption in 1992. Meanwhile, inflation ranged from 100 percent to over 30,000 percent between 1980 and 1995. As political and economic conditions stabilized in the late 1990s, a period of rapid, sustained growth took hold in the early 2000s.

Sugarcane was part of and subject to a broad-based increase in labor demand; from 2002 to 2013, real wages rose substantially in all occupations and industries while hours worked increased almost everywhere but agriculture. Figure 1.4 shows that, economy-wide, median hourly wages increased by more than 50 percent in real terms. Hours worked increased by more than 20 percent. Figure 1.5 shows the growth of real wages and hours worked by industry. Figure 1.6 gives the same information by occupation. All industries and occupations experienced meaningful wage growth from 2002–2013. The quantity of labor increased in all industries except agriculture and domestic services. The quantity of labor increased in all occupations except agricultural workers. Hours worked in agriculture fell by about 20 percent during this period.²⁵

Growers faced steadily increasing wages for sugarcane workers before and during the period of rapid mechanization, while labor supply appears to decline in later years. Even adjusting for inflation, wages for sugarcane workers nearly doubled between 1999 and 2013.²⁶ Increases in hours worked through 2007 imply increases in labor demand that coincide with a large increase in area harvested. Subsequent decreases in hours, combined with higher wages, suggest contractions in labor supply from 2008 to 2013 (see Figure 1.7).

²⁵The information in this paragraph, along with Figures 1.4, 1.5, and 1.6, is drawn from a nationally-representative household survey called the Pesquisa Nacional por Amostra de Domiclios (PNAD). Wages and hours worked are from a reference week, typically in early September, that includes the harvest season for many crops, including sugarcane.


²⁶Author's calculations using administrative data (RAIS). Results are substantively similar using household survey data (PNAD).

Figure 1.4: Aggregate Employment and Real Wages from 1999–2013




(a) Hours on the horizontal axis to emphasize movements of supply and demand

(b) Years on the horizontal axis to emphasize evolution of each series




From PNAD micro data; PNAD not conducted in 2010 because that was a census year. Includes all paid workers. Hourly wages measured in 2003 R\$.

Figure 1.5: Change in Wage and Employment 2002–2013 by Industry




From PNAD micro data. Includes all paid workers. Hourly wages measured in 2003 R\$. Industries are Agriculture (Ag), Industry (Ind), Manufacturing (Man), Construction (Con), Trade & repair (T/R), Hospitality (Hosp), Transport, communication, & storage (Trans), Public administration (Pub), Education, health, & social services (Ed/Med), Domestic services (Dom), Other services (Serv), Other (Oth). Changes to PNAD industry codes prohibit easy comparisons to earlier years.

Figure 1.6: Change in Wage and Employment 2002–2013 by Occupation




From PNAD micro data. Includes all paid workers. Hourly wages measured in 2003 R\$. Occupations are Managers (Man), Professionals in arts and sciences (Prof), Mid-level technicians (Tech), Administrative workers (Adm), Service workers (Serv), Sales and business services (Sale), Agricultural workers (Ag), Production, construction, industrial, repair workers (Prod), Military (Mil). Changes to PNAD occupation codes prohibit easy comparisons to earlier years.

Figure 1.7: Employment and Wages in Sugarcane from 1999–2013



(a) Hours on the horizontal axis to emphasize movements of supply and demand

(b) Years on the horizontal axis to emphasize evolution of each series



From RAIS administrative data; results from PNAD survey data are substantively similar. Includes workers in the sugarcane-growing states of South-Central Brazil, limited by industry (sugarcane cultivation) and occupation (agricultural worker). Hourly wages measured in 2003 R\$.

1.3.2 Sugarcane Cultivation

To inform a model of technology adoption, I summarize some relevant features of sugarcane cultivation.

Sugarcane growers may harvest manually or mechanically, each with its own inputs. Using the manual technology, unskilled workers cut sugarcane stalks with basic tools like machetes. With the mechanical technology, trained machine operators use sophisticated harvesters to cut and clean the stalks. Because the types of labor and capital are substantively different between the two technologies, I treat them as separate inputs.

The capital and labor shares are different between the two harvesting technologies. The observed capital to labor ratio is higher for the mechanical harvesting technology. In interviews, growers and machine manufacturers report that harvesting the same area requires 70 to 90 percent less labor using the mechanical technology as compared to the manual technology.

The sharing of harvesting machines is extensive so I will assume growers do not face fixed costs of acquiring machines. In practice, there are several ways to share machines, including contract harvest services, machine rental, and land rental. Data indicate that sharing is widespread: according to the 2006 Census of Agriculture, fewer than 10 percent of establishments that use mechanical harvesting actually own a harvesting machine.

The productivity of the manual technology is greatly enhanced by pre-harvest burning. Restrictions on burning effectively decrease total factor productivity of the manual technology. In a burned field, workers can clear 8 to 10 tons of sugarcane per hour. In an unburned field, slowed by heavy dense vegetation, workers clear around 3 tons of sugarcane per hour. While it is possible, pre-harvest burning is almost never combined with mechanical harvesting. Machine harvesters do move slightly faster in a burned field but one important feature of this technology is that, relative to manual harvesting, burning offers a very small increase in productivity.

The productivity of each harvesting technique depends on characteristics of the land. For example, turning large harvesting machines is slow so the size, shape, and layout of a parcel affect the rate of harvesting. Machines must proceed slowly and cautiously along hillsides. Mechanical harvesters cut as little as 450 tons per day on poorly prepared fields and as much as 1,000 tons per day on ideally prepared fields. While the productivity of manual harvesting also depends on parcel characteristics, it is much less sensitive to those characteristics.

1.3.3 Modeling the Choice of Harvest Technologies

One key fact in the data is that wages rose continuously for many years while mechanical harvesting showed a kind of S-shaped adoption: initially flat and then sharply increasing. S-shaped adoption is a common empirical finding in technology adoption so models of technology adoption inevitably account for this behavior somehow. In one class of models, this behavior originates from information diffusion or learning by doing (see, e.g., Foster and Rosenzweig (1995)). In another class of models, a fixed cost results in a size threshold for adoption and the S-shape comes from the interaction between changing factor prices and the size distribution (see, e.g., David (1969)). Manuelli and Seshadri (2014) develop a model of tractor adoption that, as in this paper, has no fixed costs or information diffusion. They argue that changing factor prices and improvements to the technology itself were the major drivers of adoption.

In the model presented below, I show that S-shaped adoption can be explained by factor prices alone, emerging in a full information environment with no fixed costs and no improvements to technology. S-shaped adoption in this model follows from three assumptions. First, that aggregate production is the sum of two production functions, one for each harvesting technology. Second, that each of the two production functions has constant or increasing returns to scale. These two assumptions imply that, for a given parcel of land, a profit-maximizing grower faces a threshold wage. Above the threshold, the grower will harvest the whole parcel mechanically and below the threshold the grower will harvest manually. The third key assumption is that these thresholds depend on characteristics of the land, like steepness, which have an S-shaped distribution. Thus, steadily rising wages can give rise to abrupt changes in the rate of mechanization.

The productivity of each harvesting techniques depends on characteristics of the land, so the model considers how a grower allocates a homogenous parcel of land between manual and mechanical harvesting techniques. Although analyzing homogenous pieces of land might seem unrealistic, it enhances the model's flexibility. The total area owned by a single grower may be broken into several parcels that are internally homogenous and the grower may decide to harvest each parcel differently.

As described above, growers do not face fixed costs to acquire a harvesting machine. Other fixed costs of mechanization, if any, are incorporated into the productivity of mechanical harvesting. For example, changing the layout of rows can increase the productivity of machine harvesting but, in many cases, it is still possible to harvest mechanically with a suboptimal layout. So, instead of including a fixed cost to change the layout, a parcel with a suboptimal layout will simply have a low productivity of mechanical harvesting.

The model reflects the fact that manual and mechanical harvesting apply fundamentally different types of labor and capital to the same land. The manual harvesting tool is a machete while the mechanical harvesting tool is a sort of combine. Those machines are operated by trained drivers while the manual

harvesting is accomplished by unskilled workers. Thus, each technology (manual is denoted p for “person” and mechanical is denoted m for “machine”) has separate inputs and factor prices: L_m, K_m, L_p, K_p and w_m, r_m, w_p, r_p . Land, T , is a normalized fixed factor $T_m + T_p = 1$. Production functions are constant elasticity of substitution (CES) with constant returns to scale (CRS):²⁷

$$Y_m = A_m \left(\alpha K_m^{-\eta} + \beta L_m^{-\eta} + (1 - \alpha - \beta) T_m^{-\eta} \right)^{\frac{-1}{\eta}} \quad \text{and} \quad (1.2)$$

$$Y_p = A_p \left(\gamma K_p^{-\zeta} + \delta L_p^{-\zeta} + (1 - \gamma - \delta) T_p^{-\zeta} \right)^{\frac{-1}{\zeta}}. \quad (1.3)$$

Growers solve the following profit maximization problem:

$$\max_{L, K, T} Y_m + Y_p - w_m L_m - r_m K_m - w_p L_p - r_p K_p \quad (1.4)$$

$$\text{s.t. } T_p + T_m = 1. \quad (1.5)$$

I interpret the mechanical productivity term A_m as the suitability of a parcel for mechanical harvesting. This parameter will be high for flat parcels with perfectly arranged rows. It will be low for steep parcels or parcels with less-than-ideal layouts. The major source of variation in the manual productivity term A_p will be burning. Manual productivity is high for a burned parcel and low for an unburned parcel. I assume that wages and rental rates are not affected by the decisions made for any individual parcel, i.e. factor prices are exogenous.

In this model, growers will generally allocate an entire parcel to only one harvesting technique. This result follows from the constant returns to scale in each production function combined with linear costs; depending on productivities and factor prices, one harvesting technique will be cheaper than the other.²⁸ This observation leads to a proposition.

Proposition 1 *There exists a threshold manual wage for each parcel, above which growers mechanize and below which they harvest manually. For some combination of parameters and factor prices, growers are indifferent between techniques, which can be expressed as a threshold wage:*

$$w_p = \phi(A_p, A_m, r_p, r_m, w_m, \alpha, \beta, \eta, \gamma, \delta, \zeta) \quad (1.6)$$

$$\text{where } \frac{\partial \phi}{\partial A_p} > 0, \frac{\partial \phi}{\partial A_m} < 0. \quad (1.7)$$

Naturally, higher wages will encourage mechanization but, for any given parcel, the threshold wage will

²⁷The CES form offers some generality. Because it is difficult to substitute between capital and labor within each technology, a Leontieff production function may be the most plausible but that is a limiting case of the CES production function.

²⁸The same result emerges with increasing returns to scale. With decreasing returns to scale, it is still possible that growers will choose only one technique because land is a fixed factor.

depend on several parameters. Figure 1.8a graphically describes a grower's optimal choices. For wages below the threshold, growers mechanize none of the parcel ($T_m = 0$). For wages about the threshold, growers mechanize all of the parcel ($T_m = 1$). Each parcel may have a different threshold wage ϕ which is determined by several parameters. Taking parcel steepness, a determinant of A_m , as an example, flatter parcels will mechanize at lower wages than steeper parcels.


Since each parcel may have a different threshold wage ϕ , the aggregate response to wage changes depends on the distribution of ϕ . Since parcels differ in terms of characteristics like steepness, productivities A_p, A_m differ across parcels. Different productivities imply different thresholds ϕ . Thus, differences across parcels give rise a distribution of threshold wages ϕ .

If the cumulative distribution of ϕ is S-shaped, then steady increases in wages can lead to abrupt changes in the rate of mechanization. In places where the distribution is flat, wage changes will have limited effects on harvesting techniques. This is a situation where, for most farms, manual harvesting is either so cheap or so expensive that small wage changes have no impact on the choice of technique. In places where the distribution is steep, wage changes will have large effects because here many parcels are nearly indifferent between the two techniques. Figure 1.8b illustrates this point. $G(\phi)$ is a hypothetical CDF of the threshold ϕ across all parcels. A vertical line gives the current manual wage w_p . All parcels with a threshold ϕ below w_p will mechanize. All parcels above will harvest manually. An increase in wage between periods $t = 1$ and $t = 2$ induces very few parcels to mechanize. An identical increase in wage between periods $t = 2$ and $t = 3$ induces many parcels to mechanize as the wage $w_{p,3}$ is now above the threshold for a large fraction of parcels.

This prediction is underpinned by the S shape of the CDF of threshold wages; I find some support for this shape in the data. First, it's worth noting that every CDF has steep and flat portions, unless the underlying variable is uniformly distributed. Beyond that theoretical point, I find S-shaped CDFs in two parcel characteristics that affect the threshold wage through productivity A_m : parcel steepness and parcel size. Recall that steeper plots require machines to move more slowly, lowering A_m and smaller plots also lower A_m by requiring machines to turn more frequently. Figure 1.9a shows the empirical CDF of parcel grades, Figure 1.9b shows the empirical CDF of parcel area, and Figure 1.9c shows the joint density of parcel grade and area.²⁹ Most sugarcane parcels are observed to have similar steepness and area. If these characteristics are important determinants of the threshold wage, then we would expect most sugarcane parcels to have similar threshold wages. Future drafts will present direct evidence regarding the relationship between grade, area, and threshold wages.

Thus, profit maximizing growers, even in the absence of information frictions or fixed costs, may respond to continuously rising wages by sharply increasing the rate of mechanization.

²⁹Note that the unit of observation is a parcel and not an establishment as in the Census of Agriculture. A parcel is defined as a unbroken area of sugarcane cultivation that is harvested using the same technique. One establishment may be composed of many parcels.



(a) Land allocation under the model

(b) Harvesting techniques and the distribution of ϕ







Figure 1.9: Empirical Support for S-Shaped Adoption



(a) Parcel-level CDF of Steepness (2010)



(b) Parcel-level CDF of Area (2010)



(c) Joint Density of Steepness and Area (2010)

1.3.4 Estimating The Effect of Wages on Mechanization

This section will estimate the change in labor intensity that would be expected solely due to the observed change in wages. Measuring the effect of wages requires exogenous variation to solve a basic problem of reverse causality, namely that mechanization itself could affect wages by shifting labor demand. This section details a method for estimating the elasticity of labor demand in sugarcane using instrumental variables based on fluctuations in labor demand from other agricultural sectors.

The study region is a major producer of sugarcane, coffee, oranges, soybeans, and maize. From 2006–2010, the study region accounted for 32 percent of world sugarcane production, 21 percent of world coffee production, 23 percent of world orange production, and 4 percent of world maize production.³⁰ According to household survey data, each crop employs between 0.5 and 2.5 million workers.³¹

I construct instruments for sugarcane wages that are conceptually similar to the instrument used in Dube and Vargas (2013); exogenous fluctuations in the markets for other crops generate variation in the wage for sugarcane workers. The intuition is straightforward. There are four other crops that i) are grown in the same region as sugarcane, ii) have a large land area devoted to their cultivation, iii) employ at least as many workers as sugarcane, and iv) employ similar types of workers as sugarcane. They are maize, soybeans, coffee, and oranges. Shocks to the production of these crops will likely shift the supply of labor facing sugarcane producers. For instance, unfavorable rains in Indonesia lead to lower coffee output, raising the international price. Coffee farmers hire more labor, increasing their output with extra pruning and tending. Because the newly-hired coffee workers might have harvested sugarcane, low coffee output in Indonesia generates a contraction in the labor supply faced by sugarcane farmers.

Taking coffee as an example, the instruments are constructed as the interaction between two components: coffee output from the top three other producers interacted with the historical area cultivated of coffee. The first component of each instrument is the output from the top 3 producers of coffee, excluding Brazil.³² This provides time-varying shifts in the labor supply faced by sugarcane farmers. The magnitude of that shift will depend on the importance of coffee to the local economy, offering a source of cross-sectional variation. Therefore, the second component of each instrument is the historical area cultivated of coffee, calculated as the 1994–1998 average.

I include a set of lags for each crop to account for the possibility that agricultural labor demand does not respond to contemporaneous price shocks. The speed of the response might vary by crops so I use

³⁰Brazilian output from PAM. World output from FAO.

³¹Calculated from PNAD.

³²The international price might influence agricultural labor supply more directly but this region of Brazil grows 20 to 30 percent of the world's coffee, soybeans, oranges, and sugarcane. Therefore, Brazilian output can have meaningful impacts on the international price of those crops. Imagine there's a frost in Brazil that damages both coffee and sugarcane plants. As a consequence, international coffee prices rise and sugarcane farmers hire less labor. In this case, the frost is directly affecting both the international price and the outcome. International prices cannot satisfy the exclusion restriction required of a valid instrument. I use output from other large producers to address this concern. However, this strategy may not totally eliminate the confounding effects of weather if those large producers change their output meaningfully in response to events in Brazil.

several different lag structures. For the temporary crops maize and soy, which can be harvested the same season they are planted, I use contemporaneous and previous year data. For the permanent crops coffee and oranges, which take 3 to 5 years to bear fruit, I use contemporaneous data and four lags. The vector of excluded instruments Z is given by

$$\begin{aligned}
Z_{j,t} \equiv \{ & \text{Int}_j^{\text{maize}} \times Y_t^{\text{maize}}, \text{Int}_j^{\text{maize}} \times Y_{t-1}^{\text{maize}}, \\
& \text{Int}_j^{\text{soy}} \times Y_t^{\text{soy}}, \text{Int}_j^{\text{soy}} \times Y_{t-1}^{\text{soy}}, \\
& \text{Int}_j^{\text{cof}} \times Y_t^{\text{cof}}, \text{Int}_j^{\text{cof}} \times Y_{t-1}^{\text{cof}}, \dots, \text{Int}_j^{\text{cof}} \times Y_{t-4}^{\text{cof}}, \\
& \text{Int}_j^{\text{orng}} \times Y_t^{\text{orng}}, \text{Int}_j^{\text{orng}} \times Y_{t-1}^{\text{orng}}, \dots, \text{Int}_j^{\text{orng}} \times Y_{t-4}^{\text{orng}} \},
\end{aligned} \tag{1.8}$$

where Y represents logged output from the top 3 producers of crop c excluding Brazil. Intensity of cultivation is defined as the historical average area of crop c harvested in municipality j .

1.3.4.1 Measuring Wages and Quantities Labor, Constructing the Instrument

As in Section Section 1.2.4.1, I turn to confidential administrative data (RAIS) for annual, municipio-level wages and quantities of labor. Recall that these data include the universe of formal-sector workers and most sugarcane workers are in the formal sector. Using detailed occupation and industry codes, I measure the wages and employment of only the relevant workers: manual laborers in the sugarcane industry. I draw other-country crop output from the UN Food and Agriculture Organization's FAOSTAT database. The PAM survey described in Section 1.2.4.1 contains annual, municipio-level area harvested for all crops; I draw contemporaneous sugarcane area and historical area for other crops are from PAM.

1.3.4.2 Estimating the Elasticity of Labor Demand via IV

The first-stage regression takes the form:

$$\log w_{j,t}^S = \gamma_0 + \gamma_1 X_{j,t} + \Gamma Z_{j,t} + u, \tag{1.9}$$

where the superscript S indicates sugarcane, j indexes municipality, and t indexes year. I measure wages w^S as the median real wage for manual laborers employed in the sugarcane industry. This study is primarily concerned with the rapid decline in labor intensity associated with mechanization. Sugarcane area harvested increase meaningfully between 1999 and 2013. To control for the associated increase in labor demand, X includes the natural logarithm of sugarcane area harvested.

The second stage regression is given by:

$$\log L_{j,t}^S = \beta_0 + \beta_1 X_{j,t} + \beta_2 \widehat{\log w_{j,t}^S} + \varepsilon \tag{1.10}$$

The outcome variable L^S is the hours worked by manual laborers employed by sugarcane growers. The object of interest is β_2 which, when multiplied by the observe wage change, will provide an estimate of the change in labor predicted by the change in wages.

Table 1.2 below presents the estimation results from the second stage. The estimates in the first row show the increase in log hours associated with a one log point increase in area harvested. The second row displays estimates of the wage elasticity of labor demand. Columns (1)-(5) estimate equation ((1.10)) using different sets of instruments. Column (1) uses the maize instruments only. Columns (2)-(4) use soybeans only, coffee only, and oranges only. Column (5) uses the instruments from all four crops. Column (6) uses the instruments from all four crops but adds a squared wage term to capture the curvature in labor demand predicted by the model.³³

Table 1.2: IV Estimates of Wage Elasticity of Labor Demand

	(1) Maize	(2) Soybeans	(3) Coffee	(4) Oranges	(5) All	(6) All
$\log(A^S)$	0.852*** (0.0653)	0.598*** (0.0900)	0.813*** (0.0438)	0.868*** (0.0560)	0.840*** (0.0405)	0.794*** (0.0446)
$\log(w^S)$	-2.938** (0.992)	1.423 (1.477)	-2.204*** (0.551)	-3.201*** (0.726)	-2.729*** (0.420)	38.86* (17.01)
$[\log(w^S)]^2$						-3.383* (1.387)
N	8,233	8,233	8,179	8,179	8,179	8,179
F stat	19.9	13.9	15.1	5.1	9.5	


SEs clustered by county.

The estimates of the wage elasticity of labor demand are statistically significant with a consistent and credible magnitude. Columns (1)-(4) use only the instruments associated with one crop. Each of the four estimates relies on a different source of variation; the cross-sectional intensities and price series for each crop are different. In three out of the four crops, the estimated elasticities are negative, statistically significant, and similar in magnitude. The estimate from soybeans is imprecise; although the point estimate is positive 1.4, the lower bound of the 95 percent confidence interval is negative 1.5, roughly. For oranges, the excluded instruments are not as strong as one would like, with a first-stage F-statistic of 5.1. Column (5) uses the instruments for all crops, again finding a strongly significant elasticity near -3.

There is some evidence of the curvature in labor demand that is predicted by the model. Column (6) accounts for this curvature by adding the square of log wages as second endogenous variable. As shown

³³The sample includes the years 1999-2013. About 400 municipalities appear in the early years, increasing to almost 800 by the end of the sample. This increase likely corresponds to the approximate doubling of sugarcane cultivation during that time but, as the labor data come from an administrative dataset, it could also be the result of increased reporting.

Figure 1.10: Quadratic and Linear IV Results



in Figure 1.10, the models in columns (5) and (6) make roughly similar predictions for the quantity of labor over the observed range of wages.³⁴ However, the model from column (6) predicts a smaller elasticity at lower wages and a larger elasticity at higher wages. This pattern is consistent with both the model and the aggregate pattern of mechanization.

1.3.4.3 Expected Change in Labor Intensity from Observed Change in Wages

We can predict the change in labor intensity attributable to the wage changes alone by multiplying the observed wage change by the estimated elasticity. This calculation assumes higher wages result from a decrease in labor supply; area harvested is meant to control for any shifts in labor demand. Differencing the second stage, equation ((1.10)) yields:

$$\widehat{\Delta \log L^s} = \hat{\beta}_1 \times \Delta \log A^s + \hat{\beta}_2 \times \Delta \log w^s \quad (1.11)$$

Now, I substitute the estimates for $\hat{\beta}_1, \hat{\beta}_2$ from Column (5) above, the change in the log of aggregate area cultivated for $\Delta \log A^s$, and the change in the log of aggregate median real wages for $\Delta \log w^s$:

$$= 0.84 \times 0.7 + -2.7 \times 0.5 = -0.76 \text{ with 95\% CI } [-1.13; -0.39] \quad (1.12)$$

Given an observed change in the log of aggregate quantity of labor of -0.49, the wage changes are sufficient to predict the decline in labor intensity.

³⁴The log of real median wages for sugarcane workers was near 6 in 1999 and close to 6.5 by 2013.

1.4 Conclusions

The course of development involves a shift from low productivity agriculture that employs many to high productivity agriculture that employs few. Brazilian sugarcane offers a window into this process with the recent adoption of mechanized harvesting.

As lawmakers contemplated regulation, economists predicted that mechanization would dramatically depress employment (Osse, 2002). Instead, the adoption of labor-saving technology coincided with a period of tightness in the labor market, characterized by large wage increases and increases in aggregate employment. By making manual harvesting more expensive, these wage changes potentially contributed to mechanization but mechanization began after years of rising wages.

Wages are not the only candidate explanation; beginning in 2002, state governments passed regulation that prioritized reductions to pollution and its associated health benefits. At the predicted cost of hundreds of thousands of jobs, sugarcane growers were obligated to reduce and eventually halt the pre-harvest burning that facilitates manual harvesting. Increasingly stringent regulation coincides with the adoption of machine harvesting.

However, the empirical analysis developed in this paper argues that development itself, in the form of rising wages for some of the poorest workers in Brazil, pulled labor out of agriculture. A range of econometric evaluations find limited evidence that the regulation contributed to mechanization. By contrast, my estimate of the elasticity of labor demand suggests that the observed change in wages is sufficient to explain mechanization.

These empirical results are supported by a theoretical framework that reconciles the trends in wages and mechanization. The model has two key insights. First, that there exists a threshold wage for each parcel of land; for wages below that threshold it is cheaper to harvest manually but for wages above that threshold it is cheaper to harvest mechanically. Therefore, if wages do not rise above the threshold, they can increase without causing a change in harvesting techniques. The second insight is that each parcel of land will have a different threshold based on the characteristics of that parcel. Because each parcel has a different threshold wage, mechanized and manual harvesting will coexist at certain wages.

This paper argues that mechanization was caused by increasing wages without thoroughly investigating why wages were increasing. That wages were increasing in spite of widespread mechanization is surprising and it suggests large increases in labor demand, especially for low-skilled workers. Future research will explore changes in Brazilian labor demand, identifying the sectors responsible and exploring the consequences for the wage distribution. A companion paper studies the health benefits that might be expected from the reduction in burning.

Another way to view these results is as a success of sustainable development. It is often assumed that more economic output means more pollution but, in this context, major development markers, like wages

for the poor and agricultural productivity, improved alongside environmental outcomes. Sugarcane farmers once burned an area the size of New Jersey every year. Mechanization has significantly curtailed the air pollution associated with sugarcane harvesting even as agricultural workers earn substantially more. Development may well have led to a better environment in this case.

1.5 Appendix

1.5.1 Labor Demand

It is possible to derive an analytical expression for labor demand by assuming a form of the production function and assuming a joint distribution of threshold wage and parcel area.

I begin by assuming that the production functions are Leontieff. From section 1.3.3, farmers choose to harvest manually when the manual wage w_p is below the threshold wage ψ . For parcels which are harvested manually, i.e. for which $\psi \geq w_p$, a Leontieff production function has two implications i) farmers will devote the entire parcel area to manual harvesting ($T_p = T$), ii) L_p will always be employed in a fixed proportion to land area $L_p = \lambda_p T_p = \lambda_p T$. Thus, for any given parcel, labor demand is given by:

$$L_p = \begin{cases} 0 & \text{if } \psi \leq w_p \\ \lambda_p T & \text{if } \psi > w_p \end{cases} \quad (1.13)$$

Moving from the parcel-level labor demand to aggregate labor demand requires an assumption about the joint distribution of ψ and land area T . For analytical convenience, I assume that ψ , T have a joint normal distribution:

$$\begin{pmatrix} \psi \\ T \end{pmatrix} \sim N \left[\begin{pmatrix} \mu_\psi \\ \mu_T \end{pmatrix}, \begin{pmatrix} \sigma_\psi^2 & \rho\sigma_\psi\sigma_T \\ \rho\sigma_\psi\sigma_T & \sigma_T^2 \end{pmatrix} \right] \quad (1.14)$$

Aggregate labor demand will be given by (with i indexing N total parcels):

$$\mathbb{E} \left[\sum_i L_p \right] = \sum_i \mathbb{E} [L_p] = N \mathbf{P}(\psi > w_p) \lambda_p \mathbb{E} [T \mid \psi > w_p] \quad (1.15)$$

$$= N \left(1 - \Phi \left(\frac{w_p - \mu_\psi}{\sigma_\psi} \right) \right) \lambda_p \mathbb{E} [T \mid \psi > w_p] \quad (1.16)$$

where $\Phi(\cdot)$ is the standard normal CDF. To derive an expression for the last expectation, begin with the conditional expectation of a jointly distributed normal variable

$$\mathbb{E} [T \mid \psi = w_p] = \mu_T + \rho \frac{\sigma_T}{\sigma_\psi} \left(\mathbb{E} [\psi \mid \psi = w_p] - \mu_\psi \right) \quad (1.17)$$

Then, by the Law of Iterated Expectations,

$$\mathbb{E}[T \mid \psi > w_p] = \mu_T + \rho \frac{\sigma_T}{\sigma_\psi} \left(\mathbb{E}[\psi \mid \psi > w_p] - \mu_\psi \right) \quad (1.18)$$

$$= \mu_T + \rho \frac{\sigma_T}{\sigma_\psi} \left(\mu_\psi + \sigma_\psi \frac{\phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)}{1 - \Phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)} - \mu_\psi \right) \quad (1.19)$$

$$= \mu_T + \rho \sigma_T \frac{\phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)}{1 - \Phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)} \quad (1.20)$$

where $\phi(\cdot)$ is the standard normal PDF. Returning to equation (1.16),


$$\mathbb{E}\left[\sum_i L_p\right] = \sum_i \mathbb{E}[L_p] = N \left(1 - \Phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)\right) \lambda_p \left(\mu_T + \rho \sigma_T \frac{\phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)}{1 - \Phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)}\right) \quad (1.21)$$

$$= \underbrace{N \lambda_p \mu_T}_{\text{scale}} \left[\underbrace{\left(1 - \Phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)\right)}_{\text{Crossing thresholds}} + \underbrace{\rho \frac{\sigma_T}{\mu_T} \frac{\phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)}{1 - \Phi\left(\frac{w_p - \mu_\psi}{\sigma_\psi}\right)}}_{\text{threshold area corr.}} \right] \quad (1.22)$$

We can interpret this expression as the combination of three effects. The first term, labeled “scale,” scales the quantity of labor based on the number of parcels N and the labor required to harvest a parcel of average area, $\lambda_p \mu_T$. The second term, labeled “crossing thresholds,” measures the fraction of parcels that are harvesting manually, i.e. those with thresholds ψ above the observed wage w_p . The final term, labeled “threshold area corr.,” adjusts quantity of labor based on the correlation between parcel threshold ψ and parcel area T . Because parcels only harvest manually if the threshold ψ is above the observed wage w_p , labor is only demanded by parcels in a truncated portion of the threshold distribution. If parcel thresholds ψ are correlated with parcel area T , knowing that a parcel threshold ψ is above the observed wage w_p also reveals something about the area of parcels that are demanding labor. For example, if large parcels have lower thresholds, i.e. T and ψ are negatively correlated, we must adjust the quantity of labor to reflect the fact that manually harvested parcels will be the smallest parcels.

This expression provides general intuition about the competing forces that determine aggregate demand for sugarcane harvest labor. However, deriving this expression required two non-trivial assumptions. Leontieff production functions are reasonable in this context as inputs are not readily substitutable. The joint normal distribution, however, is less defensible. The wage threshold ψ is unobservable and the distribution of parcel areas is not normal. The resulting labor demand has an undesirable feature: it can have a positive slope if parcel area T has a high coefficient of variation and / or the correlation ρ is large in absolute value. Figure 1.11 plots Equation (1.22) for several values of ρ with a coefficient of variation equal to 2.5, which lies at the low end of the observed value across the years 2006–2010.

Figure 1.11: Labor Demand with Different Correlations ρ



1.5.2 Regression Discontinuity Evidence: The Steepness Threshold

Below, I describe a supplemental regression discontinuity analysis that takes advantage of another threshold in So Paulo's regulation. In addition to the 150 hectare area threshold, growers were exempted from strict regulation if their plots had a steepness of at least 12 percent.

Exploiting the steepness threshold built into the regulation, I evaluate the regulation with a regression discontinuity design with burning as the outcome. From satellite data analyzed by the Brazilian space agency INPE, I directly observe the location and harvesting method for all sugarcane cultivation in the state of São Paulo from 2006 to 2012. I calculate the mean slope of each plot using a high-resolution digital elevation model produced by NASA.³⁵

As in Section 1.2.3.2, I use the method described in Calonico et al. (2014). The outcome Y takes a value of one if the field is unburned and zero otherwise and the running variable parcel steepness.³⁶

³⁵In principle, the area of 150 hectares offers another threshold which might be tested using a RD analysis. Unfortunately, the parcels identified in the satellite data do not necessarily correspond to legally defined property boundaries. Comparing the distribution of areas in the satellite data to the distribution of farm areas from the 2006 Agricultural Census, parcels in the satellite data are generally smaller than the farm area. Since the regulation was meant to apply to each farm, using satellite-measured area would systematically understate the running variable, undermining any resulting RD estimates. In future drafts, I will exploit this 150 hectare area threshold using confidential micro data from the 2006 Agricultural Census. This issue will also affect the measurement of steepness. However, with steepness, the splitting of farms may just add noise to the measured running variable, rather than systematically mismeasuring it.

³⁶I calculate parcel steepness using a digital elevation map from NASA's Shuttle Radar Topography Mission. These data record elevation on a 30 meter by 30 meter grid. I calculate percent grade at each grid point and take the mean of all grid points within each parcel. Results do not change substantively if I use the 75th percentile instead of the mean.

We would expect to observe a discontinuity in harvesting practices at the grade and size thresholds if i) farmers at each threshold would choose manual harvests in the absence of the regulation and ii) the regulation was effective. We would expect these discontinuities to appear in 2002 and to widen in 2006 and 2007, corresponding to changes in the law.

The regulation does not appear to affect harvesting practices near the threshold; there is no statistically significant difference in burning propensity across the area or steepness thresholds. Figure 1.12 show the binned averages, a fitted global polynomial control function, and the threshold for 2007.³⁷ This figure is largely representative of the other years; somewhere between 40 and 60 percent of fields are burned and there is no jump in burning at the thresholds. Figure 1.13 plot the estimated average treatment effect at the threshold for each year, along with 95 percent confidence intervals. The estimates represents a percentage point change in the likelihood of mechanization associated with being subject to the regulation, i.e. below the steepness threshold. If the regulation caused mechanization below the threshold, we would expect statistically significant positive results. Instead, we find negative point estimates, none of which is significantly different from zero.


The estimates are somewhat imprecise but they rule out effects of a size we would expect if the regulation were binding and well-enforced. The RD estimates the change in the likelihood that a parcel will be mechanized moving from the unregulated to the regulated side of the threshold. Because the area of parcels will be equal, in expectation, on either side of the threshold, the estimate also corresponds to the fraction of area mechanized because of the regulation.³⁸ Between 2007 and 2009, the agreement required the mechanization of 50 percent of farm area for farms below 12 percent grade and no mechanization of farms above 12 percent grade. Thus, if the regulation were binding and well-enforced near the threshold, we would expect a RD estimate of 0.5 in those years. The upper limit of the confidence interval is only about one quarter that size. Similar logic and conclusions apply to the estimates from 2006 and 2010.

There are a several potential explanations for these results. One is that the regulation had no effect, at least not near the thresholds. Another is that I mismeasure the running variable. For this analysis, each “parcel” is a polygon outlined from a satellite image that may or may not correspond to the unit at which the policy was applied. As outlined above, this limitation of the data introduces noise into the running variable. This noise will decrease the precision of the estimates and attenuate the point estimates so, if there is a small effect of the regulation, the analysis might fail to detect it. A final explanation is that, instead of changing harvesting practices, farmers on the regulated side threshold switched crops or lay fallow. I will consider changes in land use as in outcome in future drafts.

³⁷The global polynomial is for illustrative purposes only. The RD estimates are the difference between intercepts of a local polynomial.


³⁸Running the RD analysis with area as the outcome variable reveals no statistically significant differences in area across the threshold.

Figure 1.12: Burning Propensity and Fitted Polynomial (2007)



Excludes plots below the 150 ha threshold for area; Triangular kernel; BW and bin sizes estimated as in Calonico, Cattaneo, and Titiunik.

Figure 1.13: Estimated Average Treatment Effect at the Threshold by Year



Excludes plots below the 150 ha threshold for area; Triangular kernel; BW and bin sizes estimated as in Calonico, Cattaneo, and Titiunik.


1.5.3 Supplemental Analysis for Regression Discontinuity Estimates Using Size Threshold

1.5.3.1 Bandwidths and Running Variable Density

This section will address two issues relevant to the validity of regression discontinuity estimates: i) continuity of the running variable at the threshold and ii) bandwidth selection.

The density of establishment area is discontinuous at the regulatory threshold; rather than manipulation to avoid regulation, this discontinuity appears to be the result of heaping at multiples of 10. I test for continuity of the density using the procedure described by Cattaneo et al. (2016) but results are substantively similar using McCrary (2008). Figure 1.14 shows the p-values for tests of continuity conducted at every multiple of 10 between 100 and 200 hectares, both inside and outside of So Paulo. At the 95 percent level, the test rejects continuity for 8 of 11 area thresholds inside So Paulo. Outside So Paulo, among neighboring states without a burning regulation at the time, the test rejects continuity for 7 of 11 area thresholds. The histogram shown in Figure 1.15 makes clear that observations are concentrated near multiples of 5. Moreover, the histogram does not suggest a large mass of establishments strategically underreporting their area to avoid mechanization.

Figure 1.14: p-values for Continuity of the Density of the Running Variable



Data from 2006 Census of Agriculture.

Heaping is unlikely to introduce a material bias in the RD estimates. Imagine that the regulator and the grower know the establishment's true area but, responding to the Census, the grower rounds area to the nearest multiple of 5. The RD estimate assumes that establishments who report 150 regulated

when, in this scenario, some are not. In principle, the RD could underestimate the effect of regulation by counting some unregulated establishments as regulated. However, this bias is probably small because, in spite of the statistically significant discontinuity in the density, less than 40 observations report exactly 150 hectares, only a fraction of these would be misclassified, the estimation samples include over 4,000 observations, and the outcome is binary.

Heaping may explain the discontinuities in the density of the running variable but it does not rule out manipulation; if manipulation is present, it suggests that the already small RD estimates may be an upper bound on the effect of regulation. Perhaps some of the observations reporting 145 hectares truly have 150. Intuition suggests that such manipulation would lead to an overestimate of the effect of regulation. The growers most likely to strategically underreport their area are the growers with the highest costs of complying with the regulation and, consequently, the least likely to comply. Since manipulation would remove from the “treatment” group the establishments that are least likely to comply, manipulation would introduce an upward bias into the RD estimate.


Bandwidth selection has important impacts on point estimation and inference. Generally speaking, a larger bandwidth increases the precision of a point estimate at the cost of increasing bias. Recent work on RD estimation by Cattaneo and coauthors argues that estimating this misspecification bias is critical for correct inference. Confidence intervals must be recentered, to adjust for the bias, and scaled, to account for the variability introduced by estimating the bias. Estimating the bias also necessitates selecting a bandwidth. Since each bandwidth may be different on either side of the threshold, point estimation and inference can use up to 4 bandwidths.

Because bandwidth selection is important, I report estimates based on systematic, objective procedures that produce bandwidths with desirable properties. To estimate the RD parameter τ , I use the bandwidth selection procedure from Calonico et al. (2014) which balances bias and variance to minimize MSE. To generate confidence intervals, I use the bandwidth selection procedure from Calonico et al. (2016), which minimizes coverage error. For point estimating and inference, I choose separate bandwidths on either side of the threshold. I do so because the density of the running variable is decreasing rapidly around the cutoff, resulting in a large number of observations below the threshold and few above. Symmetric bandwidths would give a precise but biased estimate below the threshold and an imprecise but relatively unbiased estimate above the threshold. Using separate bandwidths ensures that estimates above and below the threshold optimally balance bias and variance.

The regulation is unable to explain mechanization under conservative assumptions and a range of bandwidths. Figure 1.16 shows the point estimates with confidence intervals for a range of symmetric bandwidths, including manually selected and MSE-optimal bandwidths. I show symmetric bandwidths here for ease of exposition and because they yield larger point estimates and wider confidence intervals. Since the evidence overall argues against a large effect for the regulation, I take these estimates as conservative. No point estimate is larger than 10 percentage points in absolute value and the point estimates tend to zero


as the bandwidth increases. None of the 95 percent confidence intervals exclude zero. The confidence intervals do exclude a decrease of more than 30 percentage points in manual harvesting, an increase of more than 30 percentage points in manual and mechanical, and an increase of more than 20 percentage points in mechanical harvesting. Recall that, by 2014, virtually all harvesting was done mechanically. To account for this change, the regulation would have to reduce manual harvesting by roughly 65 percentage points, with corresponding changes in the other outcomes

Figure 1.15: Histogram of Establishment Area (rounded to nearest hectare)




(a) So Paulo

(b) Neighbors




Data from 2006 Census of Agriculture. Disclosure rules prevent me from displaying bins with less than 10 observations. This level is marked on the graph. The true value for missing bins may be anywhere below the line.


Figure 1.16: Point Estimates and CIs at Various Symmetric Bandwidths



(a) Manual only



(b) Manual and mechanical



(c) Mechanical only

Data from 2006 Agricultural Census. Point estimate from local polynomial of degree 1. Bias correction from local polynomial of degree 2. Variance estimated using nearest neighbor method clustered by municipality. There are multiple point estimates at each manually selected bandwidth h ; these correspond to different bandwidths b used to generate robust bias-corrected confidence intervals. The bias-correction bandwidth b must be at least as large as the point-estimate bandwidth h . For each symmetric h , I estimate confidence intervals using $h \leq b \in \{25, 50, 75, 100, 125, 150\}$ and display the widest.

Table 1.3: The Effect of Regulation on Harvesting Techniques (RD)


	$\mathbb{1}(\text{Manual})$	$\mathbb{1}(\text{Both})$	$\mathbb{1}(\text{Mechanical})$
$\hat{\tau}$	-0.041 [-0.139; 0.047]	0.013 [-0.068; 0.099]	0.025 [-0.035; 0.093]
h_-	75	87	78
h_+	755	697	1,270
N	14,795	14,795	14,795
N_-	1,863	2,375	2,070
N_+	2,402	2,348	2,594

Point estimate from local polynomial of degree 1. Bias correction from local polynomial of degree 2. Variance estimated using nearest neighbor method clustered by municipality. Reported bandwidths h selected to generate the MSE optimal point estimates $\hat{\tau}$. The CER optimal CIs use somewhat smaller bandwidths. Assignment variable is Establishment Area.


1.5.3.2 The Full Set of Binary Outcomes

The 2006 Census of Agriculture asks respondents to report whether they use manual harvesting only, mechanical harvesting only, or both. This is the only direct question on harvesting practices. I convert this question into three indicator variables: $\mathbb{1}(\text{Manual})$, $\mathbb{1}(\text{Mechanical})$, and $\mathbb{1}(\text{Both})$. For brevity, Section 1.2.3.2 presents the results of estimation for only the first of those three indicator variables. In Table 1.3 and Figure 1.17, I provide the results for all three. Estimation procedures are the same as described in Section 1.2.3.2.


Figure 1.17: Binned Scatter Plots and Local Linear Estimates



(a) Manual only



(b) Manual and mechanical




(c) Mechanical only

Data from 2006 Agricultural Census. Local linear estimates are plotted for $x \in [150 - h_-, \min\{300, 150 + h_+\}]$. Bin sizes set to mimic the variance of the underlying data (see Calonico et al. (2015)).


1.5.3.3 Continuously-varying Inputs as Outcome Variables

If growers were shifting large amounts of area from manual to mechanical harvesting, we would expect to see similarly large changes in a variety of inputs that are measured by the Census of Agriculture. Repeating the estimation procedure from Section 1.2.3.2, I consider three labor-related inputs (days paid to temporary workers, the number of temporary workers, and the total number of workers) and five machine-related inputs (expenditure on contracting services, fuel expenditure, the number of harvesting machines, machine rental expenditure, the value of all vehicles). Figures 1.18 and 1.19 show binned scatterplots of the data along with regression lines. In Table 1.4, I present estimates from the same regression discontinuity procedure as above to see if regulated growers used different inputs than their unregulated counterparts. The estimates are noisy so, in general, the confidence intervals do not exclude large effects of regulation. However, only one of the point estimates are statistically significant and some have counterintuitive signs. The number of employees is higher among regulated establishments, and significant at the 5 percent level, while fuel expenditure and the number of harvesting machines is lower. Finally, the graphs provide no visual evidence of input changes at or near the threshold.


Figure 1.18: Binned Scatter Plots and Local Linear Estimates for Manual Harvesting Inputs



(a) Count of Temporary Workers




(b) Days Worked by Temporary Workers




(c) Count of All Workers

Data from 2006 Agricultural Census. Local linear estimates are plotted for $x \in [150 - h_-, \min\{300, 150 + h_+\}]$. Bin sizes set to mimic the variance of the underlying data (see Calonico et al. (2015)).


Figure 1.19: Binned Scatter Plots and Local Linear Estimates for Mechanical Harvesting Inputs




(a) Contracting Expenditure




(b) Fuel Expenditure



(c) Count of Harvesting Machines



(d) Value of Vehicles



(e) Machine Rental Expenditure

Data from 2006 Agricultural Census. Local linear estimates are plotted for $x \in [150 - h_-, \min\{300, 150 + h_+\}]$. Bin sizes set to mimic the variance of the underlying data (see Calonico et al. (2015)).

Table 1.4: The Effect of Regulation on Input Use (RD)

Input	$\hat{\tau}$	95% CI	\bar{Y}	h_-	h_+	N_-	N_+
Temp. days paid	146	[-167; 305]	197	51	1,214	1,047	2,590
Temp. employ.	7.25	[-2.03; 15.7]	3.6	43	443	869	2,097
Tot. employ.	7.77	[1.23; 21]	5.22	34	340	707	1,946
Cont. Exp.	2.63	[-7.09; 13.9]	14.9	45	466	906	2,139
Fuel Exp.	-10.3	[-12.7; 8.9]	.	67	57	978	429
Harvesters	-.00976	[-.149; .0911]	.118	51	392	1,040	2,031
Mach. Rent Exp.	506	[-2,827; 2,396]	491	66	269	1,555	1,767
Val. Vehicles	10.7	[-63.9; 172]	132	14	556	235	2,239

1.5.4 Difference-in-Difference Estimates Using the Size Threshold

The area threshold in the So Paulo regulation also enables me to estimate the effect of regulation via difference-in-differences. Recall that So Paulo was the only regulated state at the time of the 2006 Census of Agriculture and that establishments below 150 hectares were exempted from the regulation. I compare the difference between So Paulo establishments above and below the 150 hectare threshold to the same difference among non-So Paulo establishments. Specifically, I estimate:

$$Y_i = \beta_0 + \beta_1 \mathbb{1}(T_i > 150 \text{ ha}) + \beta_2 \mathbb{1}(\text{State}_i = \text{SP}) + \beta_3 [\mathbb{1}(T_i > 150 \text{ ha}) \times \mathbb{1}(\text{State}_i = \text{SP})] + \varepsilon_i \quad (1.23)$$

where β_3 is the effect of the regulation. The unit of observation is an agricultural establishment, indexed by i . The outcome Y_i may be one of three variables: indicators for i) establishments that use manual harvesting only, ii) establishments that use mechanical harvesting only, or iii) establishments that use both. In the reported specifications, I add state fixed effects and a polynomial in area.

Note that this identification strategy yields a different parameter than the RD strategy in Section 1.2.3.2 even though they both use the same 150 hectare threshold. The difference-in-differences estimate can be interpreted as an average treatment effect on the treated so, in principle, it may be more sensitive to changes away from 150 hectare threshold (see, e.g., Athey and Imbens (2006)). That said, most establishments are near or below the threshold so this procedure may also fail to detect an effect on the small number of large establishments.

The estimates, shown in Tables 1.5, 1.6, and 1.7, argue that regulation encouraged growers to move towards mechanical harvesting. The estimates of β_3 indicate the average effect of regulation on harvesting practices for regulation establishments. Regulated establishments are about 10 percentage points less likely to use manually only. Regulated establishments are 3 percentage points more likely to use mechanical only and 7 percentage points more likely to use both techniques. The estimated effect of the regulation is significant in all specifications.

While these effects seem large compared to the means, they are within the confidence intervals of the RD estimates in Section 1.2.3.2 and they are small compared to the intended effect of the regulation and the eventual outcome of complete mechanization. Recall that, above the threshold, more than 60 percent of establishments report manual harvesting only. According to the regulation, none of them should rely exclusively on manual harvesting. Moreover, establishments over 150 hectares are almost completely mechanized by 2014. If the regulation causes 10 percent of establishments to use some mechanized harvesting, then the regulation is certainly insufficient to explain the change in harvesting techniques.

Table 1.5: Likelihood of Harvesting Only Manually

	(1)	(2)	(3)	(4)
Area > 150ha (β_1)	0.084*** 0.020	0.098*** 0.030	-0.044 0.037	-0.044*** 0.010
São Paulo (β_2)	-0.188 0.021	-0.078 0.061	-0.090 0.063	-0.090*** 0.005
Regulated (β_3)	-0.104*** 0.024	-0.134*** 0.041	-0.159*** 0.041	-0.159*** 0.013
\bar{Y}	0.781	0.781	0.781	0.781
σ_Y	0.413	0.413	0.413	0.413
Clust. SE	Y	Y	Y	
Area poly	Y	Y		
State FE	Y			
N	30,423	30,423	30,423	30,423

Agricultural establishments are the unit of observation. Removed observations in the top and bottom 1% by area. The outcome variable is an indicator for Harvesting Only Manually. Assignment variable is Establishment Area. In the indicated columns, SE estimates are clustered by county.

Table 1.6: Likelihood of Harvesting Only Mechanically

	(1)	(2)	(3)	(4)
Area > 150ha (β_1)	-0.019** 0.009	-0.021** 0.009	0.019 0.014	0.019*** 0.006
São Paulo (β_2)	0.063 0.014	0.043** 0.017	0.047*** 0.018	0.047*** 0.003
Regulated (β_3)	0.030* 0.016	0.034* 0.018	0.040** 0.018	0.040*** 0.008
\bar{Y}	0.067	0.067	0.067	0.067
σ_Y	0.249	0.249	0.249	0.249
Clust. SE	Y	Y	Y	
Area poly	Y	Y		
State FE	Y			
N	30,423	30,423	30,423	30,423

Agricultural establishments are the unit of observation. Removed observations in the top and bottom 1% by area. The outcome variable is an indicator for Harvesting Only Mechanically. Assignment variable is Establishment Area. In the indicated columns, SE estimates are clustered by county.

Table 1.7: Likelihood of Harvesting Manually and Mechanically

	(1)	(2)	(3)	(4)
Area > 150ha (β_1)	-0.064*** 0.016	-0.077*** 0.025	0.025 0.029	0.025*** 0.008
São Paulo (β_2)	0.125 0.015	0.035 0.050	0.042 0.051	0.042*** 0.004
Regulated (β_3)	0.074*** 0.019	0.100*** 0.032	0.119*** 0.032	0.119*** 0.011
\bar{Y}	0.152	0.152	0.152	0.152
σ_Y	0.359	0.359	0.359	0.359
Clust. SE	Y	Y	Y	
Area poly	Y	Y		
State FE	Y			
N	30,423	30,423	30,423	30,423

Agricultural establishments are the unit of observation. Removed observations in the top and bottom 1% by area. The outcome variable is an indicator for Harvesting Manually and Mechanically. Assignment variable is Establishment Area. In the indicated columns, SE estimates are clustered by county.

CHAPTER 2

Putting Out the Fires: Measuring the Local Effects of Air Pollution on Health

The effects of air pollution on health vary by local conditions, but the sparsity of air quality monitoring has prevented researchers from producing locally-relevant estimates for much of the world. We evaluate an extremely accessible method to do so, estimating a simple regression using freely available satellite and vital statistics data. While this method could be implemented almost anywhere, we choose South Central Brazil for two reasons. First, it is an area where recent policy debates would have benefited from estimates of the health costs of pollution. Second, it is also the setting for the more standard approach in Rangel and Vogl (2016), which serves as a benchmark. Disappointingly, the method we test substantively underestimates the health consequences of air pollution relative to the benchmark paper and the existing literature generally. We discuss potential refinements to this method, but our results highlight the value of ground-based monitors and valid instruments for pollution exposure.

2.1 Introduction

Data from air quality networks have yielded substantial insights about the effects of air pollution on health. One of these insights is that the effects are context-specific, depending on local conditions like weather and the concentration and chemical composition of pollutants (WHO, 2013). Unfortunately, air quality networks are sparse, concentrated in urban areas and wealthy countries, which leaves many regions without locally-relevant estimates of the effects of air pollution (Hsu et al., 2016; Mathiesen, 2016). For example, many developing countries suffer very high concentrations of pollutants, often from different sources of pollution like agricultural burning, but lack a comprehensive monitoring system. In this paper, we evaluate a method for estimating the health effects of pollution that can be easily implemented

Co-authored with Sebastian Miller of the Inter-American Development Bank.

in many contexts by combining remote sensing data with basic vital statistics.

We evaluate this approach by estimating the health effects of air pollution in a middle-income context where the same has been done using data from ground monitors. Specifically, we study the health effects of pollution from agricultural burning in South-Central Brazil. We compare these results to those from Rangel and Vogl (2016), who relied on data from a six-site network of ground monitors. In general, the health effects estimated via remote sensing data are very close to zero, notably smaller than those from ground monitors. There are several possible explanations for these differences, including: i) the remote sensing data covers a larger geographic area and a longer period of time, possibly masking some heterogeneity, ii) attenuating measurement error in the remote sensing data, and iii) confounding variables that are difficult to address with remote sensing data. While these results do not immediately recommend remote sensing data, refinements to this methodology may yet yield a more complete understanding of how air pollution affects health.

The effects of pollution on health are heterogeneous, varying based on local conditions (WHO, 2013). The sources of local pollution matter because the health effects of pollution depend on the mixture of pollutants at a particular location. For example, laboratory studies show that the combination of dust and ozone has a much greater effect on respiratory function than either dust or ozone alone (Mihave et al., 2005; Tamás et al., 2006). Weather conditions like humidity and temperature also modify the effects of air pollution (Vanos et al., 2015).

Air pollution is unmeasured in much of the world; given the heterogeneity in the health effects, a lack of data on air pollution means the health consequences are also unmeasured in much of the world. Some 92 countries with a total population of 1 billion lack the capacity to monitor particulate matter, which is known to have negative effects on health. Another 33 countries have only one or two monitors in large urban areas. These countries tend to be poor (Hsu et al., 2016; Mathiesen, 2016). Monitoring in rural areas is extremely rare, even though rural air quality suffers as a result of wild fires and agricultural burning.

Understanding the health benefits of reduced pollution is critical to policymakers, for instance, as they weigh measure to curb agricultural burning. In the late 1990s and early 2000s, a public debate played out in the sugarcane growing region of Brazil. Some voices argued for the elimination of sugarcane burning, citing its environmental and health impacts, while others opposed intervention, predicting 300,000 sugarcane workers would lose their jobs. Agricultural burning has also become a political issue in India, where it helped make New Delhi the most polluted city in the world. But, as in Brazil, Indian farmers worry that burning restrictions will drive them out of business (Anand, 2016).

Remote sensing data has shown promise in filling this gap, but this paper is the first to compare a broadly applicable methodology based on remote sensing to well-identified causal estimate from ground-level monitors. The global scale of satellite measurements has allowed researchers to estimate the health effects of air pollution in places where that would otherwise be impossible (see, e.g., Jayachandran (2009);

Johnston et al. (2012)). In this paper, we exploit the frequency and resolution of satellite data to estimate the health effects via a fixed effects strategy. This methodology could be applied broadly, since the health outcomes are drawn from widely available vital statistics data, and NASA freely distributes satellite data with near global coverage.

We apply the methodology to the sugarcane growing region of Brazil for two reasons. First, the recent work by Rangel and Vogl (2016) serves as a benchmark. Using data from a six-station network installed in 2009, the authors generate plausibly exogenous variation in pollution exposure by combining fire detections with changes in wind direction. We compare the results to those of Rangel and Vogl (2016) to evaluate our broadly applicable remote sensing methodology. Second, this region exemplifies the need; it has a large rural population that experiences substantial air pollution from agricultural burning, but, until recently, had no monitoring network. This lack of data left government officials with limited information on which to base policy.

While there is room for refinement, results from the remote sensing methodology do not agree with Rangel and Vogl (2016) or previous literature. Our estimates show a near zero effect of prenatal exposure on early life outcomes, and our estimates are precise. By contrast, Rangel and Vogl (2016), along with much of the prior literature, show significant reductions in birth weight, gestational age, and cohort size. The satellite data cover a larger area and a longer time span, possible obscuring negative effects observed in the narrower setting studied by Rangel and Vogl (2016). The satellite data could suffer from attenuating measurement error. Finally, even with daily data on exposure, we cannot definitively exclude the possibility that unobserved variables confound our estimates. This method may yet prove useful to policy makers who currently lack solid information about air pollution and health, but not without improvements.

2.2 Context and Mechanisms

2.2.1 Sugarcane Burning and Health in South Central Brazil

Brazil grows more sugarcane than any other country on earth. Since 2000, it has accounted for one quarter to one third of global output, with almost 10 million hectares cultivated in 2012.¹ Sugarcane is grown extensively in 8 states: Goiás, Minas Gerais, Paraná, Mato Grosso do Sul, Rio de Janeiro, São Paulo, Pernambuco, and Alagoas. This study will focus on the first six, which form a contiguous group in South-Central Brazil, as shown in Figure 2.1). Pernambuco and Alagoas are on the Northeastern coast and will be excluded from the analysis. Since 2000, the study region has accounted for at least three

¹Data from FAOSTAT. See Figure 2.10 for details.

quarters of Brazil's output, producing more sugarcane than India from 2003 to 2012.²

Figure 2.1: Major sugarcane producing states and the study region



There are two primary harvesting techniques; the historically predominant technique involves first burning the cane fields to dispose of dead plant matter and eliminate pests. Sugarcane is harvested during the dry season that runs from April through November although most of the pre-harvest burning occurs in August, September, and October. The harvest can be done manually or with harvesting machines. Pre-harvest burning substantially improves the efficiency of manual harvesting but not mechanical harvesting. For some decades, the vast majority of sugarcane has been harvested manually, preceded by burning.

Smoke from biomass burning is bad for human health. The pre-harvest burning of sugarcane adds a number of chemicals to the local atmosphere, including volatile organic compounds, carbon monoxide, carbon dioxide, oxides of nitrogen and sulfur, and particulate matter (PM) (Tsao et al., 2011; Pope III and Dockery, 2006; Andreae et al., 1991).³ Both the mixture and the individual components are known to have negative health effects (WHO, 2013). Particulate matter is especially harmful, even at relatively low concentrations, and sugarcane burning is a major source of PM.⁴ Daily fluctuations in PM have been associated with mortality in a number of cities. Long-term exposure correlates with increased infant

²Data from FAOSTAT and the Produção Agrícola Municipal (PAM) database maintained by Brazil's census bureau, IBGE. See Figures 2.10, 2.11 for details.

³"Particulate matter ... is a complex mixture of extremely small particles and liquid droplets. Particle pollution is made up of a number of components, including acids (such as nitrates and sulfates), organic chemicals, metals, and soil or dust particles." PM is generally separated into two categories: inhalable coarse particles, with diameters between 2.5 and 10 micrometers, and fine particles, with diameters of less than 2.5 micrometers (US Environmental Protection Agency).

⁴Sugarcane burning accounts for as much as 60 percent of the PM in central São Paulo and sugarcane burning coincides with dramatic spikes in local PM concentrations (Lara et al., 2005).

mortality, respiratory illness among children, and cardiopulmonary deaths among adults. Hospitalizations for respiratory illness, lung function, and school absences also tend to increase with the level of PM. The effects of acute exposure are generally small; for exposure periods of 5 days or less, relative mortality risk increases by about 1 percent for every $10 \mu\text{g}/\text{m}^3$ increase in fine particulate matter.⁵ A similar increase in concentrations over the course of a month is associated with a 3 percent increase in mortality risk. Over the course of years, the increase in risk is on the order of 10 percent. For both acute and chronic exposure, the relationship between PM concentrations and the percent increase in deaths is approximately linear over a wide range of PM concentrations (Pope III and Dockery, 2006).⁶

Epidemiological research has verified these findings for some cities in South-Central Brazil: sugarcane burning is associated with increased concentrations of particulate matter which, in turn, predict higher hospital admissions for hypertension and respiratory ailments (Arbex et al., 2000, 2007, 2010). Consistent with the findings from other areas, the relationship in South-Central Brazil is strongest for children and the elderly (Cançado et al., 2006).

In the economics literature, various types of air pollution have been identified as detrimental to infant health. Kenneth Chay and Michael Greenstone made early contributions in this area, showing that particulate matter increases infant mortality. They estimate elasticities, finding that a 1 percent decline in particulate matter is associated with a 0.35 to 0.5 percent decline in infant mortality (Chay and Greenstone, 2003b,a). Janet Currie and coauthors have used a variety of identification strategies to demonstrate that carbon monoxide tends to shorten gestation, decrease birth weight, and increase infant mortality. These papers find, among other things, that the roll-out of EZPass in the northeastern United States lowered the incidence of low birth weight babies by 10 percent, thanks to reductions in CO-generating vehicle traffic and that observed reductions in California's CO levels saved the lives of 1,000 infants during the 1990s (Currie and Walker, 2011; Currie and Neidell, 2005; Currie et al., 2009). Previous work by Seema Jayachandran is similar to our enterprise in both topic and methods; she also uses satellite data to assess the health consequences of emissions from biomass burning. Her results suggest that widespread forest fires in Indonesia increased infant and child mortality, reducing the size of affected cohorts by 1.2 percent, or 15,600 children (Jayachandran, 2009).

⁵Rangel and Vogl (2016) summarize data from ground monitors in rural São Paulo. The station with the lowest levels of PM shows an average reading of about $21 \mu\text{g}/\text{m}^3$ with a standard deviation of about 11. The station with the highest levels of PM shows an average of $38 \mu\text{g}/\text{m}^3$ with a standard deviation of 19.


⁶While the findings of the epidemiology research are robust, the underlying biology is not completely understood. See Pope III and Dockery (2006). One study, conducted in Brazil, exposed one group of mice to pollution from vehicle traffic and another group to pollution from sugarcane burning. Relative to controls, both groups of mice suffered deterioration in the mechanical function of their lungs (Mazzoli-Rocha et al., 2008).

2.2.2 Harvesting Transition

This is a compelling context in part because it highlights the need for local estimates of the health effects of air pollution. As detailed in Chapter 1, sugarcane burning became a policy target because of its effects on air quality, but policymakers acted with limited evidence about the health effects. Responding to constituent concerns about air pollution, the state of São Paulo debated regulation in the late 1990s, ultimately passing a gradual ban of sugarcane burning 2002. The five surrounding states passed similar regulations between 2008 and 2014. The first study we located on air pollution and health in the region was published in 2000 (Arbex et al., 2000). Follow up work appeared in 2006, 2007, and 2010 (Cançado et al., 2006; Arbex et al., 2007, 2010). All of these were correlational. The first study with strong causal identification appeared years later, with the first draft of Rangel and Vogl (2016) appearing in 2014. Having a timely estimate of the health effects would have been useful to policymakers balancing health against the genuine concerns about job losses associated with mechanization.

The policy target was achieved, although the policy played a limited role: sugarcane harvesting was rapidly mechanized beginning around 2007, significantly curtailing the practice of pre-harvest burning.⁷ This trend is visible both in terms of employment and the sales of harvesting machines. Combining administrative employment records with crop data collected by the Ministry of Agriculture, Figure 2.2a shows the hours worked per hectare of sugarcane harvested. Through 2007, harvesting required 60 to 70 hours per hectare. By 2013, harvesting required less than 20 hours per hectare. Interviews with a manufacturer of sugarcane harvesters indicate that the decline in labor intensity coincided with a large increase in sales of machinery. Beginning in 2006, São Paulo used satellite imagery to directly measure the burning of sugarcane. Figure 2.2b shows the area of sugarcane harvested by harvest type.

Figure 2.2: The Transition to Mechanized Harvesting



The area of burned cane declines substantially while the area of unburned cane quadruples; about 66

⁷Explaining the transition is the focus of Chapter 1. Changes in factor prices appear to be the main causal driver, not regulation.

percent of sugarcane was burned in 2006 but only 16 percent was burned in 2013. If farmers had burned 66 percent of their 2013 harvest, they would have set fire to an additional 24,000 square kilometers in the state of São Paulo alone. For scale, the state of Maryland covers about 25,000 square kilometers and the Indonesian forest fire studied by Jayachandran burned an area of approximately 48,000 square kilometers. What health benefits can we attribute to this stunning decline in polluting activity? Can policymakers in other contexts turn to a straightforward combination of satellite data and vital statistics for much needed insights?

2.3 Data and Descriptive Statistics

Demographics

Table 2.1 presents summary statistics of the study region from the 2000 demographic census. The first column includes all individuals in the study region. The second column is limited to the São Paulo and Rio de Janeiro metropolitan areas. For the remaining columns, each statistic is calculated at the municipal level and the 1,460 municipalities are ranked. The second column presents the 10th percentile of each statistic across municipalities, the third column presents the median, and the fourth column presents the 90th percentile.

Over seventy seven million people resided there in 2000. It's 90.5 percent urban with a median household income of 720 reais per month and practically everyone has electricity. The region is home to two of Brazil's largest cities: São Paulo and Rio de Janeiro. Together, these metro areas account for more than a third of the study region's population. Not surprisingly, they are more urbanized than the rest of the study region. Adults are more educated, on average, and children are more likely to be in school. Households are richer in these urban areas and the income distribution appears more skewed. Racial composition varies meaningfully over municipalities, although the vast majority are "branca." Generally, there is meaningful diversity both within and between the communities of this region.

2.3.1 Airborne Particulates

Smoke from biomass burning is detectable by satellite imagery. The presence and concentration of tropospheric smoke, ash, or desert dust can be quantified based on their ultraviolet signature.⁸ The measure itself is called the Absorbing Aerosol Index (AAI), a unit-less quantity calculated from the ratio of radiances at two different wavelengths in the ultraviolet spectrum. One of the canonical uses for the AAI is to gauge smoke from biomass burning and we employ it as a measure human exposure to

⁸The troposphere is the portion of the atmosphere closest to the earth's surface. It has an average depth of 11-12 miles.

Table 2.1: Summary Statistics for Study Region (2000)

	Population	SP & RJ	P10 Muni.	Med. Muni.	P90 Muni
Pct. Urban	0.905	0.971	0.449	0.766	0.955
Pct. ≥ 18	0.672	0.687	0.606	0.659	0.702
Pct. Elec.	0.989	0.999	0.904	0.991	1.000
Adult Educ. (yrs)	7.0	7.7	3.6	4.4	5.4
Pct. in School (<18)	0.689	0.701	0.251	0.286	0.320
Race: Branca	0.681	0.613	0.460	0.704	0.857
Race: Preta	0.052	0.071	0.014	0.034	0.076
Race: Amarela	0.009	0.011	0.000	0.002	0.015
Race: Parda	0.256	0.303	0.109	0.241	0.473
Race: Indigena	0.003	0.002	0.000	0.001	0.005
HH Inc. (P10)	151	200	90	151	230
HH Inc. (med)	720	850	302	451	724
HH Inc. (P90)	3,000	3,785	1,000	1,501	2,450
Tot. Pop.	77,639,115	28,774,153	3,383	13,150	82,014

Sources: 2000 Demographic Census.

air pollution from the pre-harvest burning of sugarcane. AAI is able to distinguish clouds from smoke (Herman et al., 1997; Gleason et al., 1998; Duncan et al., 2003; Hsu et al., 1996).

The index is defined with a minimum value of zero and higher values generally correspond to higher concentrations of aerosols, and it is well correlated with ground-based measurements. At least two other properties of an aerosol can increase its AAI: its exact chemical composition and its altitude.⁹ For a layer of smoke at a fixed concentration, AAI *increases* with altitude. For instance, a high AAI value may correspond to a high concentration of smoke near the ground or a lower concentration of smoke at altitude. Since most people spend most of their time near the ground, AAI may overstate human exposure to smoke. However, AAI values agree with other measurements of aerosol concentrations, including satellite-based Aerosol Optical Thickness and ground-based sun photometry (Herman et al., 1997; Gleason et al., 1998; Hsu et al., 1996; Torres et al., 1998, 2005).

We use satellite measures of AAI to determine daily exposure to pollution from biomass burning. Using the Ozone Monitoring Instrument on its Aura satellite, NASA offers daily AAI values on high-resolution grid. The series begins on October 1, 2004 and continues as of publication. The grid is 0.25 degrees in latitude by 0.25 degrees in longitude, which roughly equates to 70 square kilometers per grid cell. To measure air pollution at the municipality level, the AAI grid is overlaid atop a map of the municipalities in the study region. As shown in Figure 2.3, each municipality is assigned the area-weighted average AAI for all grid cells that overlap the municipality.¹⁰ Many municipalities are smaller than a single grid

⁹Details of the sophisticated physical science underpinning the AAI are beyond the scope of this article. See, e.g., Torres et al. (1998); Herman et al. (1997); Gleason et al. (1998).

¹⁰There is a noticeable linear discontinuity that runs north-to-south in the figure. On the western side of the line, AAI is visibly higher than it is on the eastern side. This is because it took two orbits for the satellite to image the entire study region,

cell but, because of the shape and alignment of the municipalities relative to the grid cells, fewer than 10 percent of municipalities draw their AAI values from a single cell. Roughly half draw their AAI values from 4 or more cells.

Observed values of the index range from zero to approximately 4 for the municipality-days in the sample. As expected, AAI tends to be higher during the harvest season:


Table 2.2: Distribution of AAI by Season

	25 th pct	Mean	75 th pct
Not harvest season	0.00	0.33	0.57
Harvest season	0.00	0.42	0.73

Even at a high level of geographic aggregation, there’s substantial variation in AAI over time and some evidence of seasonality. Figure 2.4a shows daily AAI values averaged over the entire study region. Figure 2.4b shows seasonal patterns in AAI values. It is consistent with the table above in that AAI appears to be higher and more extreme during the harvest season, especially in August and September when burning is most common. It’s important to note, however, that the table is calculated based on municipality-day observations while Figure 2.4b is based on daily observations averaged over the entire study region.

so the eastern and western sides were measured at different times of day. Relative to the earth’s surface, the satellite’s orbits shift somewhat every day so the discontinuity shifts every day. On certain days, the entire study region may be measured in a single orbit. Every 16 days, the satellite retraces the same path over the earth’s surface. We do not believe this 16 day cycle aligns with any potential confounders.

Figure 2.3: AAI Grid and Municipio Values



(a) Municipio map with AAI grid overlay

(b) Area-weighted municipio average AAI






Figure 2.4: Daily and Seasonal AAI Averaged over Study Region



(a) Daily Average AAI over Study Region

(b) Monthly Average AAI over Study Region




2.3.2 Natality and Mortality

Birth weight, gestational age, and cohort size form one set of measures we use to investigate the relationship between air pollution and health. The Brazilian Ministry of Health publishes vital statistics including the universe of birth records and the universe of death records. We use data from 2004 to 2012. There are 300,000 to 350,000 births recorded each quarter, and the data provide a range of information about the mother, the pregnancy, and the infant. Information about the mother includes education, race, marital status, number of children, and municipality of residence. Gestational age is recorded, along with birth weight, the number of clinic visits, pregnancy type (e.g. singleton, twins), delivery type, date, and time. Death records include sex, date of birth, date of death, cause of death, and municipality of residence. We assume throughout the analysis that mothers reside in the same municipality throughout their pregnancy and children do not move before they reach one year of age.

For estimation, we measure the birth weight outcome as both a continuous and a categorical variable; the gestational age outcome is an indicator for pre-term births. In addition to the birth weight in grams, we use an indicator variable for low birth weight (below 2.5kg). The density of birth weight in the sample is shown in Figure 2.5, along with the thresholds for low birth weight, very low birth weight (below 1.5kg), and extremely low birth weight (below 1kg). We consider any child born earlier than 37 weeks to be pre-term.

Figure 2.5: Density for Birth Weight (2005–2012)



There is a strong biological link between birth weight and gestational age; if air pollution leads to lower birth weights, that could be because air pollution lowers birth weight per se, air pollution induces earlier

Table 2.3: Gestational Age

Years	28-31 weeks	32-36 weeks	37-41 weeks	42+ weeks
2004–2006	.008	.061	.919	.007
2007–2009	.008	.065	.917	.005
2009-2012	.01	.084	.884	.017

births, or both.¹¹ To separate these effects, we estimate regressions of birth weight on air pollution with and without controls for gestational age. Comparing the coefficients on air pollution between these specifications shows the extent to which any effects on birth weight are mediated through gestational age.

Gestational age plays another key role: identifying the date of conception and, hence, the period of pollution exposure. Gestational age is measured as the number of weeks since the mother’s last normal period, but the data only store bins: (i) less than 22 weeks, (ii) 22 to 27 weeks, (iii) 28 to 31 weeks, (iv) 32 to 36 weeks, (v) 37 to 41 weeks, (vi) 42 weeks or more. To estimate the date of conception, we assume exact gestational ages for each bin. The precise density of gestational age is still under study by medical researchers, but it is believed to be sharply increasing up to 40 weeks and sharply decreasing after 40 weeks. For any window less than 40 weeks, the preponderance of children will be born at the end of the window. Therefore, we use (i) 21 weeks, (ii) 27 weeks, (iii) 31 weeks, (iv) 35 weeks, (v) 40 weeks, (vi) 41 weeks. Table 2.3 shows the frequency of the 4 most common categories over time. Forty weeks is considered normal and medical intervention generally precludes longer pregnancies. In the sample, about 90 percent of births occur in the normal range. Another 6 to 8 percent occur between the 32nd and 36th weeks.

2.3.3 Hospital admissions


As additional outcomes, we explore the relationship between pre-harvest burning and hospital admissions. Specifically, the outcomes of interest are weekly counts of (i) total hospital admissions and (ii) hospital admissions due to respiratory illness by municipality of residence. The Brazilian Ministry of Health maintains a database of hospital admissions. Each record contains information about the patient, including age, municipality of residence, and primary diagnosis.

Hospital admissions are highly variable, showing a slight downward trend during the study period and strong seasonal variation. Figure 2.6a plots the total number of admissions per month and respiratory admissions per month from 2000 to 2012. Total hospital admissions are near 1 million while respiratory


¹¹In the data, here is a strong positive correlation between gestational age and birth weight. Babies born between 32 and 36 weeks are a full kilogram heavier, on average, than babies born between 28 and 31 weeks. Full term babies, born between 37 and 41 weeks, gain another 0.7 kilograms.

admissions are generally below 250,000. Both series are highly variable, both show a very slight downward trend, and both show seasonal patterns. The seasonal patterns are especially strong for respiratory admissions, which peak May, as shown in Figure 2.6b.

Figure 2.6: Trends and Distributions in Hospital Admissions (2000–2012)




(a) Total Monthly Admissions




(b) Seasonal Patterns in Admissions

(c) Muni.-level Admissions (all)

(d) Muni.-level Admissions (resp.)



Order statistics calculated from all municipalities in each month.



Order statistics calculated from all municipalities in each month.

Primarily due to differences in population, admissions counts vary meaningfully across municipalities. Figure 2.6c shows the distribution of monthly hospital admissions across municipalities for all causes, while Figure 2.6d shows only respiratory causes. For each month, the figures plot the 25th, 50th, and 75th percentiles of the distribution of the municipality-level admissions count. The municipality-level median is just above 50 admissions per month for all causes. The 25th percentile municipality has about 25 admissions per month while the 75th percentile municipality falls near 130 in most months. For respiratory causes, the corresponding statistics are 18 for the median municipality, 4 for the 25th percentile mu-

nicipality, and 25 for the 75th percentile municipality. Seasonal patterns are more obvious in respiratory admissions and there's a pronounced downward trend, especially at the median and 75th percentile.

The figures only include admissions for patients at facilities participating in Brazil's Universal Health System (SUS). Data for admissions outside of the SUS are available from 2008 to 2010. Over that period, non-SUS admissions accounted for approximately 27 percent of total admissions and 24 percent of respiratory admissions. The prevalence of non-SUS admissions varies substantially across municipalities: less than 4 percent of admissions are non-SUS in the 25th percentile municipality, about 12 percent in the median municipality, and 24 percent in the 75th percentile municipality. The fraction of non-SUS admissions is higher in richer, more urban municipalities.¹²

Since non-SUS admissions are only available for a limited time period, we rely on SUS admissions, excluding patients and institutions who do not participate in the public healthcare system. To the extent that admissions caused by air pollution are non-SUS admissions, this analysis will understate the relationship. However, such admissions are likely to be SUS admissions; in similar contexts, the poor have been more vulnerable to emissions from biomass burning and non-SUS admissions are rare in poor areas Jayachandran (2009).

It's possible that excluding non-SUS admissions upwardly biases the relationship between air pollution and hospital admissions. As a hypothetical, imagine that incomes are rising faster in cane-growing areas and private healthcare is a normal good. The income effect will decrease observed admissions in cane-growing areas independent of any benefit from reduced pre-harvest burning. Unfortunately, given the limited availability of non-SUS data, it is difficult to identify substitution from public to private healthcare.¹³

¹²Data from the IBGE 2000 demographic census.

¹³The fraction of non-SUS admissions increases by about three percentage points from 2008 to 2009 and falls by one percentage point from 2009 to 2010. There are no apparent seasonal patterns over that three year period.

2.4 Estimation

2.4.1 Birth Outcomes

To understand the effects of air pollution on birth outcomes, we regress birth outcomes on exposure, a variety of controls, and fixed effects:

$$y_{i,m,g,t} = \alpha + \rho M_i + \eta Z_t + \phi X_m + \beta AAI_{m,g,t} + \varepsilon \quad (2.1)$$

y is birth weight or an indicator for pre-term

M is a vector of mother and pregnancy characteristics

Z is a vector of quarter and year FEs

X is a vector of municipio FEs

AAI mean of daily AAI over pregnancy

where i indexes a mother, m indexes municipalities, g indexes gestational age, and t indexes birth date. The vector M includes mother's race, mother's education level, mother's age and its square, mother's marital status, number of clinic visits, birth order, delivery type, pregnancy type, and child gender. In some specifications, M also includes an indicator for pre-term births. The coefficient of interest is β .

We rely on several features of the data to generate precise measures of maternal exposure to air pollution during the pregnancy. We measure air pollution for each child individually based on date of conception and mother's residence. Recall that our data provide daily measurements of air pollution. Since the natality records include both birth date and gestational age, we are able to estimate date of conception. We construct the variable AQ as the mean of daily AAI over all days of the pregnancy in the mother's municipality of residence. As an alternative measure of exposure, we also estimate Equation (2.1) using the fraction of days during a pregnancy with AAIs higher than the 75th percentile of AAI in the entire sample. Since pollution may effect fetuses differently depending on their stage of development, it is also common in the literature to separate exposure by trimester. Although we do not present the results here, separating exposure by trimester yields qualitatively similar results.

Exposure to air pollution is certainly measured with error, but this approach has two major benefits: i) it is feasible for any location with reliable vital statistics, ii) it does not require spatial interpolation often used in this literature. Error comes from imprecision in the data. We assume exact gestational ages. We assume the mother spends the entire pregnancy in the municipality of residence given on the birth certificate. Satellite measurements provide somewhat-coarse measurements of ambient air pollution, not necessarily human exposure. Nevertheless, this approach can be applied to almost any location given its limited data requirements. Global, high-frequency, high-resolution satellite data are available for download. Vital

statistics, which are becoming better and more available, are the only other data required. Moreover, even when ground monitors are available, studies often rely on Kriging or other spatial prediction methods to assign exposure via interpolation from a network of ground monitors. But interpolation techniques can generate systematic errors due to terrain or weather patterns, whereas satellites offer direct measurements of air pollution.

We identify the effect of air pollution on health via a fixed effects strategy, employing fixed effects for municipality, year, and quarter. There is some evidence from the United States that quarter of birth is correlated with parental characteristics like education and income (see, e.g., (Bound et al., 1995)). With a few characteristics of the mother, we can only partly control for these factors given the available data. Additionally, in the study region, we meaningfully observe fluctuations in the number of births in each quarter (see Figure 2.8). The quarter fixed effects are intended to control for unobserved factors that vary by season. Year fixed effects will capture any region-wide trends in birth outcomes, while municipality fixed effects capture local time-invariant factors like prevailing weather patterns that might confound the relationship between air pollution and health.

The results are identified by within-municipality deviations from the quarterly and annual averages. Implicitly, the regression compares two children conceived in the same quarter, in the same year, and in the same municipality, but have different levels of exposure to air pollution because, for instance, one was conceived during pre-harvest burning and the other was conceived after. The results of estimating Equation (2.1) are presented in Table 2.4.

Table 2.4: The Effect of AAI on Birth Outcomes

	(1) Pre	(2) BW	(3) BW	(4) LBW	(5) LBW
Mean AAI	.0019 (.00607)	-26.4* (11)	-25* (10.2)	.0201*** (.00454)	.0192*** (.00469)
Pre-term=1			-768*** (9.04)		.469*** (.00501)
N	8,800,383	8,797,877	8,797,877	8,797,877	8,797,877
FE	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni
\bar{y}	.08	3,162	3,162	.085	.085
σ_y	.27	522	522	.28	.28

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The main exposure variable, here listed as Mean AAI, is the mean AAI over all days in each pregnancy in the mother's municipality of residence.

See Table 2.19 for a complete set of covariates and coefficients.

Standard errors are clustered by microregion, contiguous administrative groups that may contain as few as one municipality and as many as forty one. The median microregion is composed of eight municipalities. We believe this is a conservative approach that at least partially accounts for spatial correlation.

The measure of air quality, AAI, is unitless, so we use order statistics from the distribution of mean AAI to aid in interpretation of the estimates. Table 2.5 shows various percentiles from the distribution of mean AAI across all measured pregnancies. We use the 90th percentile pregnancy, i.e. the pregnancy with pollution exposure higher than 90 percent of pregnancies in the population, as the baseline for our comparisons. This value represents some of the worst levels of pollution exposure in Brazil, while removing extreme values. We use the 10th percentile of exposure to imagine what would happen if air pollution was significantly reduced but not completely eliminated. So we use a decline of $0.512 - 0.285 = 0.227$ units to discuss the magnitudes of our estimates. We acknowledge these are discretionary choices and invite the reader to make other comparisons using Table 2.5.

Table 2.5: Distribution of Mean AAI Across All Measured Pregnancies

	10 th pctile	25 th pctile	50 th pctile	75 th pctile	90 th pctile
Mean AAI	.285	.316	.374	.448	.512

Air pollution does not predict premature birth, and has extremely modest effects on birth weight. In column (1), the outcome is an indicator for pre-term birth, i.e. gestational age is indicated as less than 37 weeks. Exposure is not associated with significant or substantive changes to the likelihood of pre-term birth. Even at the extreme of the confidence interval, the 90-10 change in pollution exposure reduces the likelihood of pre-term birth by just three tenths of a percentage point, or 3.9 percent of the population mean.¹⁴

In columns (2) and (3), the outcome is birth weight and we see that high AAI has a significant, negative association with birth weight. Air pollution might affect birth weight directly or indirectly by causing pre-term deliveries; alternatively, air pollution might affect unmeasured aspect of infant health that itself drives both birth weight and gestation. Column (3) controls for pre-term birth while column (2) does not. Conditioning on pre-term birth does not meaningfully change the estimate. Consistent with the results from column (1), exposure seems to act through some other channel other than pre-term birth. Consistent with the effects on average birth weight, increases in AAI predict increases in the likelihood of low birth weight. Here also, the magnitude and significance of the estimate are not materially altered by controlling for pre-term births.

While exposure is associated with poorer health outcomes, the magnitude of the relationship is tiny. The estimates imply that moving from the 90th percentile of pollution exposure to the 10th percentile of pollution exposure would increase birth weight by about 6 grams.¹⁵ Although the estimate is statistically significant, it hardly seems meaningful considering that birth weight has a mean of almost 3,200 grams and a standard deviation greater than 500. The probability of low birth weight would fall by half of a percentage point.¹⁶ Again, compare this estimate to a mean of 0.0846 and a standard deviation of 0.278.

¹⁴ $(0.0019 + 1.96 \times 0.0061) \times (0.285 - 0.512) \approx -0.0031$; $-0.0031/0.08 \approx -0.039$.

¹⁵ $(0.285 - 0.512) \times -26.4 \approx 6$.

¹⁶ $(0.285 - 0.512) \times 0.020 \approx -0.005$.

Taken together, air pollution does not appear to drive important changes in birth weight.

The magnitude, sign, and significance of these results are robust to a range of alternative specifications. If damage to health is convex in air pollution, extreme exposures would be more important than the mean in determining health outcomes. To measure extreme exposures, for each birth we calculate the fraction of gestational days when AAI was above the 75th percentile of all municipality-day observations. Using this exposure variable does not materially alter the results, as shown in Table 2.6. Interacting the year and quarter fixed effects allows seasonal variations to change over time but does not change the results. Neither does interacting year fixed effects with indicator variables for each state, which flexibly controls for state-specific time trends. The results are nearly identical when using monthly, rather than quarterly, fixed effects to account for seasonal variation.

Table 2.6: The Effect of Extreme AAI on Birth Outcomes

	(1) Pre	(2) BW	(3) BW	(4) LBW	(5) LBW
Hi. AAI	.00639 (.00698)	-40.1*** (8.82)	-35.2*** (8.68)	.0288*** (.00498)	.0257*** (.00453)
Pre-term=1			-768*** (9.04)		.469*** (.00501)
N	8,800,383	8,797,877	8,797,877	8,797,877	8,797,877
FE	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni
\bar{y}	.08	3,162	3,162	.085	.085
σ_y	.27	522	522	.28	.28

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Municipality fixed effects appear to capture important unobservable correlates of air pollution and health. As shown in Table 2.7, dropping municipality fixed effects gives insignificant estimates with counterintuitive signs. Thus, within a given municipality, AAI is associated with slightly worse birth outcomes but, across municipalities, AAI is associated with slightly better birth outcomes. One possibility is that richer municipalities experience higher AAI because of local production and transportation, but children are healthier thanks to better nutrition and medical care. Since our goal is to estimate the causal effect of air pollution on health outcomes, including municipality fixed effects is appropriate in this example, as they proxy for omitted variables like income or access to health services.

The effects of exposure on health outcomes are stronger outside major urban areas. In Table 2.8, we see estimates increase in magnitude and precision if we exclude births from large urban areas. As above, this change could be explained by omitted variables. Of course, fixed effects cannot completely capture time-varying unobservable differences. If income-generating, air-polluting economic activity increased more rapidly in the cities, then this is also a candidate explanation for the results of regressions excluding

Table 2.7: The Effect of AAI on Birth Outcomes without Muni FE

	(1) Pre	(2) BW	(3) BW	(4) LBW	(5) LBW
Mean AAI	.0171 (.0152)	67.5 (55.9)	80.7 (46.3)	.000426 (.013)	-.00761 (.00708)
Pre-term=1			-771*** (8.62)		.47*** (.00499)
N	8,800,383	8,797,877	8,797,877	8,797,877	8,797,877
FE	Y,Q	Y,Q	Y,Q	Y,Q	Y,Q
\bar{y}	.08	3,162	3,162	.085	.085
σ_y	.27	522	522	.28	.28

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

major metropolitan areas. To put it another way, the variation in AAI outside of the major metropolitan areas may be less confounding than the variation in AAI within the major metropolitan areas. Another potential explanation is differences in exposure; Table 2.9 shows that Mean AAI tends to be higher and more extreme in major metropolitan areas. Perhaps the damage to health is, in fact, concave. In any case, even the twice-as-high estimate we recover outside of major metropolitan areas is quite small in magnitude.

Table 2.8: The Effect of AAI on Birth Outcomes without Major Metros

	(1) Pre	(2) BW	(3) BW	(4) LBW	(5) LBW
Mean AAI	.00961 (.00895)	-53.8*** (9.09)	-46.6*** (8.99)	.023*** (.00376)	.0186*** (.00414)
Pre-term=1			-746*** (4.99)		.457*** (.00395)
N	4,216,418	4,214,617	4,214,617	4,214,617	4,214,617
FE	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni
\bar{y}	.075	3,169	3,169	.081	.081
σ_y	.26	518	518	.27	.27

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Births are excluded if the mother resides in any major metropolitan area, as defined by the IBGE. The major metropolitan areas in the study region are: Belo Horizonte (MG), Vale do Aço (MG), Rio de Janeiro (RJ), São Paulo (SP), Baixada Santista (SP), Campinas (SP), Curitiba (PR), Londrina (PR), Maringá (PR), Goiânia (GO), Distrito Federal (GO, MG).


A small effect on the mean could be explained by a small effect on the entire distribution or a large effect

Table 2.9: Distribution of Mean AAI by Metro Region

	10 th pctile	25 th pctile	50 th pctile	75 th pctile	90 th pctile
Mean AAI inside Metro	.297	.344	.418	.473	.543
Mean AAI outside Metro	.278	.300	.332	.379	.443

on a small subset of children. To investigate this possibility, we run the same regression describe using birth weight as the outcome but omitting the air quality variable $AI_{m,g,t}$. Then we group the population into “treated” and “untreated” based on the air quality variable: the “treated” with high exposure to air pollution and the “untreated” with low exposure. Finally, we plot the cumulative distribution of the regression residuals for each group. If the treatment affects some individuals differently than others, we would expect the two CDFs to diverge. In the plot below, treatment is defined as being in the 95th percentile of pollution exposure. We see little evidence for treatment-effect heterogeneity as the CDFs to do not noticeably diverge. This result survives under several alternative definitions of treatment.

Figure 2.7: CDF of Regression Residuals by Pollution Exposure



These results suggest a causal relationship between air pollution and birth outcomes but we remain cautious in our conclusions. Our fixed-effects strategy cannot account for changes in income or labor supply that might be correlated with both birth outcomes and pollution-generating activities like sugarcane burning. For example, large-scale pre-harvest burning might be associated with an increase in the quantity of labor supplied by households. That increase could, in turn, generate physical or emotional stress in

mothers that affects birth outcomes. Attenuation bias associated with classical measurement error may partially explain the small magnitude of our results.

Compared to Rangel and Vogl (2016), our results show much smaller effects on birth weight.¹⁷ The comparison is complicated by the fact that Rangel and Vogl use the number of upwind fires as their measure of exposure to pollution. Taking a somewhat crude approximation of their data, we assume that the number of upwind fires observes a Poisson distribution with mean 2. In this case, the difference between the 90th percentile of exposure and the 10th percentile is roughly 4 fires.¹⁸ The authors' preferred estimates show that one less fire raises birth weight by 23 grams, so moving from the 90th percentile to the 10th implies an increase in birth weight of about 92 grams. For that same change in exposure, their estimates imply that the likelihood of low birth weight falls by about 3.3 percentage points. These estimates are an order of magnitude larger than our own.

The difference may be partly explained by the confounding effects of economic activity. The relationship between air pollution and health could have two channels: i) the direct effect of air pollution, ii) the effect of economic activity that is correlated with air pollution, e.g. earnings and employment. In the sugarcane growing region of Brazil, fires and exposure to air pollution occur during the harvest, when earnings and employment tend to be high. Higher earnings and employment might lead to better health outcomes. Rangel and Vogl (2016) address this issue by measuring the effects of downwind fires on health, where only the earnings and employment channel is active. They find that downwind fires indeed predict better health outcomes. However, the magnitude of the relationship between downwind fires and health is much smaller than one between upwind fires and health, suggesting that this confounding variation is an incomplete explanation for the differences between our estimates and those of Rangel and Vogl (2016). We discuss other potential explanations later in this chapter.

¹⁷As in our analysis, their results on gestational age generally do not reach conventional levels of statistical significance. However, our results on gestational age are also much smaller in magnitude than those of Rangel and Vogl (2016).

¹⁸See Table 1 of Rangel and Vogl (2016) for the mean and standard deviation of fire counts.

2.4.2 Infant Mortality

Following Jayachandran (2009), we estimate the following specification to investigate the relationship between air quality and infant mortality:

$$y_{m,t} = \alpha + \eta Z_t + \phi X_m + \beta_1 AAI_{m,t-1} + \beta_2 AAI_{m,t} + \beta_3 AAI_{m,t+1} + \varepsilon \quad (2.2)$$

y is all-cause, one-year infant mortality rate or the log of cohort size at one year

Z is a vector of quarter and year FEs

X is a vector of municipio FEs

AAI mean of daily AAI by month

where m indexes municipalities and t indexes the month-year of birth. Note that the timing of air pollution measurements is different between Equations (1) and (2). Here, we include the mean daily air pollution in the month before birth, the month of birth, and the month after birth, capturing pre-, peri- and post-natal exposure. As before, standard errors are clustered by microregion.

Because these outcomes are measured at the municipality level, rather than the individual level, air pollution is measured as average of the daily values in a particular municipality over a calendar month. To interpret the estimates, we again summarize the distribution of the air pollution variable, as shown in Table 2.10. The difference between a 90th percentile municipality-month and a 10th percentile municipality-month is, at most, 0.3 AAI units.

Table 2.10: Distribution of Lagged, Concurrent, and Lead AAI

	10 th pctile	25 th pctile	50 th pctile	75 th pctile	90 th pctile
Lag AAI	0.216	0.265	0.331	0.410	0.510
AAI	0.215	0.264	0.329	0.408	0.504
Lead AAI	0.215	0.263	0.328	0.407	0.503

The estimates for the effect of AAI on the mortality rate are imprecise. No individual coefficient is statistically significant, nor are the coefficients jointly insignificant, as a Wald test cannot reject the null hypothesis that all three AAI coefficients are zero. At the extreme of the confidence interval, moving from the 90th percentile of exposure to the 10th in the month before birth would lower the mortality rate by 0.627 per 1,000 births.¹⁹ In the month of birth, the same change in exposure would lower the mortality rate by 1.333, and in the month after birth by 0.775.^{20,21} Given the population mean mortality rate of 15.3 per 1,000 births, we cannot exclude the possibility that a decline in exposure in the months around birth

¹⁹ $(0.260 - 1.96 \times 1.222) \times (0.510 - 0.216) \approx -0.627$

²⁰ $(-1.578 - 1.96 \times 1.548) \times (0.504 - 0.215) \approx -1.333$

²¹ $(-0.080 - 1.96 \times 1.333) \times (0.503 - 0.215) \approx -0.775$

Table 2.11: Effect of AAI on Mortality Rate and Cohort Size

	M.Rate	C.Size
2005	-0.146 (0.986)	0.018* (0.008)
2006	0.158 (1.017)	-0.049*** (0.008)
2007	-0.890 (1.032)	-0.076*** (0.008)
2008	-2.166* (0.995)	-0.058*** (0.009)
2009	-2.438* (0.957)	-0.087*** (0.010)
2010	-2.957** (0.962)	-0.086*** (0.010)
2011	-3.119** (0.993)	-0.074*** (0.010)
Q2	-1.622*** (0.369)	0.022*** (0.003)
Q3	-1.378** (0.476)	-0.052*** (0.004)
Q4	-0.426 (0.376)	-0.118*** (0.004)
Lag AAI	0.260 (1.222)	-0.079*** (0.010)
AAI	-1.578 (1.548)	0.007 (0.013)
Lead AAI	-0.080 (1.333)	0.034** (0.013)
N	192,243	192,034
FE	Q,Y,Muni	Q,Y,Muni
Wald p-val	.743	0
\bar{y}	15.30	2.52
σ_y	55.64	1.38

Standard errors in parentheses

2004 is the omitted year; Q1 is the omitted quarter.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

would reduce child mortality rates by more than 10 percent. The results are similar using 3-month lags and leads and when controlling for state-specific time trends using state-by-month fixed effects.

Mortality rates can underestimate the effects of air pollution since they are usually measured as a fraction of live births, neglecting effects on miscarriages and stillbirths. As such, we also estimate Equation (2) with log cohort size as the outcome. High AAI in the month prior to birth has a negative, strongly significant association with log cohort size; moving from the 90th percentile of exposure to the 10th in that month would increase cohort size by about 2.3 percentage points.²² AAI in the month of birth does not have a large or significant impact on log cohort size but lower AAI in the month after, surprisingly, is associated with a decrease in cohort size. Decreasing exposure in the month after birth tends to increase cohort size by about 1 percent.²³ The same patterns emerge with year by quarter fixed effects, when excluding major metropolitan areas, and using 3-month leads and lags of air quality.²⁴

Echoing the results on birth outcomes, we do not find strong and consistent evidence that reductions in air pollution would lead to improvements in infant mortality. Rangel and Vogl (2016) have a similar pattern of results insofar as they find end-of-pregnancy air pollution reduces cohort size, but noisy estimates of its effect on mortality rates. However, as before, our estimates are smaller in magnitude. Rangel and Vogl find that one less downwind fire increases cohort size by 2.34 log points. Moving from the 90th percentile of exposure to the 10th implies an increase in cohort size by a substantial 10 percentage points. While their paper boasts credible identification from fluctuations in wind direction, it seems unlikely that a single downwind fire during the final trimester would cause births to fall by more than 2 percent. The mean number of fires in the data is between 2 and 3, and someone is always downwind, so their results suggest that cohorts in this region are decreased by 5 to 8 percentage points due to agricultural burning.

2.4.3 Hospital Admissions

Data on hospital admissions allow us to examine the relationship between air pollution and the health of children and adults. We consider the following outcomes: the weekly count of hospital admissions for respiratory causes and the weekly count of hospital admissions for all causes.

The two outcomes are counts characterized by a high conditional variance relative to the mean; therefore, we estimate a negative binomial model with municipality fixed effects. The final count is assumed to have a Poisson distribution where the Poisson parameter is itself a random variable that varies across municipalities and over time. The Poisson parameter has a Gamma distribution where the shape parameter is a function of covariates and the scale parameter is a municipality-specific fixed effect. Specifically, we

²² $-0.079 \times (0.510 - 0.216) \approx 0.023$

²³ $0.034 \times (0.510 - 0.216) \approx 0.010$

²⁴See Table 2.18.

estimate a model of the form

$$y_{i,t} \mid \gamma_{i,t} \sim \text{Poisson}(\gamma_{i,t}) \quad (2.3)$$

$$\gamma_{i,t} \mid \delta_i \sim \text{Gamma}(\lambda_{i,t}, \delta_i) \quad (2.4)$$

$$\lambda_{i,t} = \exp(x_{i,t} \beta), \quad (2.5)$$

where i indexes municipalities, t indexes weeks, and $x_{i,t}$ includes year and quarter FEs, mean of daily AAI over the current week, and mean of daily AAI over the previous week (lag). The results are presented below in Table 2.12. Standard errors are bootstrapped. In column (1) the outcome is total admissions while in column (2) the outcome is respiratory admissions.

Table 2.12: The Effect of AAI on Hospital Admissions

	(1)	(2)
	Adm	Resp
Mean AAI (wk)	0.00465 (0.00241)	-0.00389 (0.00613)
Mean AAI (lag)	0.00384 (0.00240)	0.00677 (0.00554)
Observations	1,104,164	1,104,164

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Consistent with the results of Rangel and Vogl (2016), there appears to be no significant, short-term relationship between air pollution and hospital admissions. This model is suited to the data but it is difficult to interpret the magnitude of the resulting estimates $\hat{\beta}$. That said, there is no statistically significant relationship between air quality and hospital admissions. The results are substantively the same using the fraction of high pollution days as the measure of exposure and in a linear regression where the outcome is the fraction of total admissions attributable to respiratory causes. Rather than total hospital admissions, Rangel and Vogl study the rate of hospital admissions during an infant's first full day of life. They also find no statistically significant results.

If a relationship between air pollution and hospital admissions exists, there are a variety of reasons we might not detect it. As before, exposure is measured with error. We are unable to observe usage of private hospitals. We omit weather-related variables like temperature and humidity that are known to influence respiratory admissions. Previous research has shown that air pollution can trigger cardiovascular events, an outcome that we have not included in our analysis. Finally, sensitive populations like asthmatics may take protective measures in response to pre-harvest burning.

2.5 Conclusion

The health effects of air pollution vary from place to place, but a lack of widespread pollution monitoring prevents researchers from developing locally-relevant estimates. The increasing quality and availability of remote sensing data provides an opportunity to fill this gap. We evaluate a method for estimating the health consequences of air pollution that can be implemented anywhere vital statistics data are available.

This method is appealing because the data are widely available and identification does not require the researcher to find a natural experiment. For the main analysis, we combine two data sources. The first is high-resolution satellite measurements of air pollution with global coverage and daily frequency. These data are freely available for download from various government agencies. The second data source is vital statistics. Recording births and deaths is a basic function of government, and the quality and availability is steadily improving. We use these data to estimate the health consequences of air pollution via a fixed effects strategy. Such a strategy may be subject to concerns about time-varying unobservable confounders, but it also allows us to study contexts without a natural experiment to provide an instrument for pollution exposure.

We apply this method in a policy-relevant context where we can compare our estimates to those from an approach that is typical in the environmental economics literature. South-Central Brazil is a region that experienced substantial air pollution from agricultural burning, ultimately provoking a policy response from state governments. Policymakers were forced to weigh potential costs to agriculture against the potential gains to health without good information about the gains to health. Years later, the effects of agricultural burning on health were studied by Rangel and Vogl (2016). They use an approach that is not uncommon in the literature: identify a plausible instrument for air pollution as measured by ground monitors, here wind direction, and draw outcomes from vital statistics data. We use this study as a benchmark.

As implemented here, our method produces estimates that are much smaller in magnitude than those in the benchmark study. For roughly comparable changes in pollution exposure, our method estimates changes in birth weight that are an order of magnitude smaller than those from Rangel and Vogl (2016). We find effects on one measure of infant mortality, cohort size, that are also meaningfully smaller. Both papers lack the precision to find a statistically significant relationship between air pollution and hospital admissions.

There are several possible explanations for these discrepancies. The first is unobserved confounders like earnings and employment that are correlated with polluting economic activity and health. Rangel and Vogl (2016) present some evidence supporting this possibility. Of course, one solution is expanding the set of control variables, e.g. adding earnings and employment data from Brazil's RAIS system, or by identifying a natural experiment as Rangel and Vogl do. Unfortunately, both of these solutions limit the

main appeal of this method, namely its feasibility in many contexts. A second potential explanation is attenuating measurement error. Satellites can measure air pollution in locations without ground-based monitors, but they still do not directly measure human exposure.

Some features of the benchmark study that may also be responsible for the discrepancies. The health effects in that study are identified as the difference between the health effects of upwind fires and downwind fires. This variation has obvious appeal. Who can say which way the wind will blow? But families have some choice about where to live, and families who live downwind from sugarcane growers may be different from those who do not.²⁵ Sugarcane growers can also choose when to burn and how much, and weather factors into those decisions. In short, Rangel and Vogl are essentially using an instrument, and that strategy yields a local average treatment effect (LATE) for some unknown population. Perhaps these issues are related to the authors' very large estimates of air pollution on cohort size. Finally, Rangel and Vogl have a narrower geographic and chronological scope.

Taken together, these results do not recommend the approach evaluated here, in spite of its promise. Perhaps it would work better at a smaller spatial or temporal scale, but here it is inconsistent with the broader literature and persuasive empirical work in the very same context. The results support economists' commitment to identifying valid instruments. Nonetheless, we believe there is a need for localized estimates of the effects of air pollution supported by improved measurement. Perhaps other solutions will emerge as air monitors grow cheaper and more mobile.

²⁵It's worth noting that in Table A6 Rangel and Vogl (2016) show that some observable characteristics of mothers and infants are *not* predicted by treatment.

2.6 Additional Tables and Figures

Table 2.13: Largest Sugarcane Growers

	Brazil		India		China		Rest of world	
	Area	Output	Area	Output	Area	Output	Area	Output
2012	9.7	721.1	5.1	347.9	1.8	124.0	11.3	763.6
2011	9.6	734.0	4.9	342.4	1.7	115.1	11.1	743.0
2010	9.1	717.5	4.2	292.3	1.7	111.5	10.5	698.1
2009	8.6	691.6	4.4	285.0	1.7	116.3	10.7	716.9
2008	8.1	645.3	5.1	348.2	1.8	124.9	10.8	741.5
2007	7.1	549.7	5.2	355.5	1.6	113.7	10.4	713.3
2006	6.4	477.4	4.2	281.2	1.4	93.3	10.0	663.3
2005	5.8	423.0	3.7	237.1	1.4	87.6	10.2	656.3
2004	5.6	415.2	3.9	233.9	1.4	91.0	10.6	691.9
2003	5.4	396.0	4.5	287.4	1.4	92.0	10.6	695.2
2002	5.1	364.4	4.4	297.2	1.4	92.2	10.8	673.1
2001	5.0	345.9	4.3	296.0	1.3	78.0	10.3	624.6
2000	4.8	327.7	4.2	299.3	1.2	69.3	10.4	630.5

Notes: area in millions of hectares, output in millions of tonnes.

Source: Food and Agriculture Organization of the United Nations, FAOSTAT database.

Table 2.14: Area under cultivation (millions of Ha)

Year	Rest of Brazil	Other major growers	Study region	Total
2000	0.6	0.8	3.5	4.9
2001	0.6	0.9	3.6	5.0
2002	0.6	0.8	3.8	5.2
2003	0.7	0.8	3.9	5.4
2004	0.7	0.8	4.2	5.6
2005	0.7	0.8	4.3	5.8
2006	0.7	0.7	4.9	6.4
2007	0.8	0.8	5.5	7.1
2008	0.8	0.8	6.6	8.2
2009	0.8	0.8	7.2	8.8
2010	0.8	0.8	7.6	9.2
2011	0.8	0.8	8.0	9.6
2012	0.9	0.7	8.1	9.8

Source: IBGE Produção Agrícola Municipal, Tabela 1612.

Table 2.15: Output (millions of Tonnes)

Year	Rest of Brazil	Other major growers	Study region	Total
2000	29.1	43.0	254.0	326.1
2001	31.4	44.7	268.2	344.3
2002	35.1	42.8	286.5	364.4
2003	40.4	45.7	309.9	396.0
2004	41.2	45.3	328.7	415.2
2005	39.5	40.8	342.6	423.0
2006	43.0	41.1	393.5	477.6
2007	47.2	44.6	457.9	549.7
2008	49.4	49.6	546.3	645.3
2009	49.3	46.2	596.0	691.6
2010	48.7	44.1	624.7	717.5
2011	50.5	48.6	634.9	734.0
2012	52.7	41.9	626.4	721.1
Total	557.5	578.4	5,669.8	6,805.7

Source: IBGE Produção Agrícola Municipal, Tabela 1612.

Table 2.16: Total revenue (billions of Reais - nominal)

Year	Rest of Brazil	Other major growers	Study region	Total
2000	0.9	1.1	4.7	6.7
2001	1.0	1.3	6.4	8.7
2002	1.9	1.4	8.2	11.5
2003	1.6	1.4	9.3	12.3
2004	1.6	1.5	9.0	12.2
2005	1.6	1.5	10.1	13.1
2006	1.9	1.7	14.1	17.7
2007	2.1	1.7	15.3	19.1
2008	2.3	2.0	16.3	20.7
2009	2.6	2.3	19.8	24.6
2010	2.7	2.5	23.1	28.3
2011	3.4	3.1	32.7	39.2
2012	3.5	2.5	34.4	40.5
Total	27.1	24.2	203.3	254.5

Source: IBGE Produção Agrícola Municipal, Tabela 1612.

Table 2.17: Median Município-level Yield (tons / Ha)

Year	Rest of Brazil	Other major growers	Study region	Total
2000	30.0	43.2	50.0	40.0
2001	31.0	45.0	50.0	40.0
2002	35.0	46.4	50.0	42.0
2003	35.0	52.0	57.4	45.0
2004	35.0	52.0	60.0	46.6
2005	34.0	50.0	60.0	45.0
2006	37.0	55.0	60.0	49.0
2007	40.0	57.6	60.0	50.0
2008	40.0	55.0	65.0	50.0
2009	38.0	58.7	65.0	50.0
2010	40.0	50.0	66.0	50.0
2011	40.0	58.0	65.3	50.0
2012	37.3	50.0	66.1	50.0
Total	35.0	50.0	60.0	48.0

Notes: Computed from IBGE Produção Agrícola Municipal, Tabela 1612.

Table 2.18: Robustness Checks for Mortality

	M.Rate	M.Rate	M.Rate	C.Size	C.Size	C.Size
2006	0.135 (1.014)	1.013 (1.056)	0.284 (0.488)	-0.049*** (0.008)	-0.053*** (0.009)	-0.068*** (0.005)
2007	-0.901 (1.035)	-0.155 (1.060)	-0.760 (0.546)	-0.076*** (0.008)	-0.083*** (0.008)	-0.096*** (0.006)
2008	-2.214* (0.990)	-1.493 (1.018)	-2.029*** (0.519)	-0.058*** (0.009)	-0.066*** (0.009)	-0.079*** (0.006)
2009	-2.458* (0.958)	-1.749 (0.974)	-2.272*** (0.487)	-0.087*** (0.010)	-0.095*** (0.010)	-0.104*** (0.007)
2010	-2.951** (0.961)	-2.507* (1.005)	-2.768*** (0.502)	-0.086*** (0.010)	-0.095*** (0.010)	-0.105*** (0.008)
2011	-3.118** (0.990)	-2.493* (1.034)	-2.924*** (0.484)	-0.074*** (0.010)	-0.084*** (0.010)	-0.077*** (0.008)
Q2	-1.610*** (0.360)	-1.693*** (0.399)	-1.570*** (0.387)	0.022*** (0.003)	0.023*** (0.003)	0.009** (0.003)
Q3	-1.337** (0.460)	-1.362* (0.525)	-1.314** (0.461)	-0.052*** (0.004)	-0.053*** (0.004)	-0.055*** (0.004)
Q4	-0.392 (0.362)	-0.273 (0.395)	-0.395 (0.398)	-0.118*** (0.004)	-0.120*** (0.004)	-0.099*** (0.004)
Lag AAI	0.039 (1.243)	0.308 (1.341)		-0.079*** (0.010)	-0.079*** (0.011)	
AAI	-1.489 (1.523)	-1.718 (1.680)	-1.556 (1.589)	0.007 (0.013)	0.008 (0.014)	0.008 (0.012)
Lead AAI	-0.234 (1.272)	0.232 (1.434)		0.035** (0.013)	0.033* (0.014)	
3-mo. pre AAI			-0.024 (1.804)			-0.176*** (0.016)
3-mo. post AAI			-0.592 (1.837)			0.097*** (0.019)
N	192,243	172,611	183,244	192,034	172,410	183,048
FE	Q,Y	Q,Y,Muni	Q,Y,Muni	Q,Y	Q,Y,Muni	Q,Y,Muni
Wald p-val	.679	.781	.678	0	0	0
\bar{y}	15.30	15.40	15.28	2.52	2.35	2.52
σ_y	55.64	57.39	55.50	1.38	1.23	1.38

Standard errors in parentheses

2004 is the omitted year; Q1 is the omitted quarter.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.19: The Effect of AAI on Birth outcomes (all coefficients)

	(1)	(2)	(3)	(4)	(5)
	Pre	BW	BW	LBW	LBW
Mother's age	-.00515*** (.000925)	15.1*** (2.37)	11.1*** (1.67)	-.00445*** (.000965)	-.00204*** (.000544)
Age squared	.000111*** (.000017)	-.293*** (.0433)	-.208*** (.0303)	.000105*** (.0000177)	.0000532*** (9.91e-06)
1st born	-.0167*** (.000664)	91.7*** (1.48)	78.9*** (.96)	-.028*** (.000704)	-.0202*** (.000412)
2nd born	-.0177*** (.00126)	108*** (2.73)	94.8*** (1.82)	-.0326*** (.00126)	-.0243*** (.000692)
3rd born	-.0168*** (.00166)	113*** (3.29)	100*** (2.06)	-.0319*** (.00173)	-.024*** (.00098)
4th born	-.0167*** (.00203)	115*** (3.39)	103*** (2.16)	-.0296*** (.00192)	-.0218*** (.00108)
5th+ born	-.0202*** (.00238)	129*** (4.51)	114*** (3.1)	-.0303*** (.00265)	-.0207*** (.00169)
Single	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
Married	.00401*** (.000603)	3.18 (2.2)	6.26*** (1.81)	-.00404*** (.000542)	-.00592*** (.000321)
Widow	.00387 (.00224)	-15.1* (6.78)	-12.2* (5.77)	.00396 (.00284)	.0022 (.00222)
Legally separated	.00897*** (.00103)	-16.2*** (2.53)	-9.26*** (2.37)	.00223** (.000833)	-.00199** (.000611)
None	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
1-3 years	-.0000565 (.00198)	25.6*** (3.7)	26*** (3.25)	-.0106*** (.00202)	-.0109*** (.00174)
4-7 years	.003 (.0018)	31.7*** (3.42)	34.5*** (3.22)	-.0128*** (.00187)	-.0145*** (.00172)

Table 2.19: The Effect of AAI on Birth outcomes (all coefficients)

	(1)	(2)	(3)	(4)	(5)
	Pre	BW	BW	LBW	LBW
8-11 years	.00763*** (.00207)	41.8*** (4.42)	48.2*** (3.89)	-.0176*** (.00203)	-.0215*** (.00178)
12 or more	.0107*** (.00239)	25.6*** (5.83)	34.3*** (4.85)	-.0203*** (.00209)	-.0257*** (.00163)
Single	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
Twins	.413*** (.0094)	-.848*** (4.5)	-.531*** (9.07)	.519*** (.00519)	.325*** (.00517)
Triplets+	.782*** (.018)	-1,387*** (24.6)	-.787*** (15.7)	.827*** (.0146)	.46*** (.00955)
Vaginal	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
C-section	.0119*** (.00201)	36.3*** (4.4)	45.4*** (3)	.00693** (.00245)	.00132 (.00167)
None	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
1-3 visits	-.00303 (.00348)	71.6*** (8.5)	69.3*** (6.13)	-.0368*** (.00575)	-.0354*** (.00431)
4-6 visits	-.0611*** (.00333)	194*** (5.33)	147*** (3.78)	-.0963*** (.00463)	-.0677*** (.00378)
7+ visits	-.139*** (.00499)	318*** (7.42)	211*** (5.18)	-.161*** (.00576)	-.0959*** (.00421)
Male	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
Female	-.00452*** (.000255)	-.110*** (.488)	-.113*** (.466)	.0167*** (.000305)	.0188*** (.000282)
Branca	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)

Table 2.19: The Effect of AAI on Birth outcomes (all coefficients)

	(1) Pre	(2) BW	(3) BW	(4) LBW	(5) LBW
Preta	-0.00248* (.0012)	-4.56* (1.85)	-6.45** (2)	.00483*** (.00106)	.00598*** (.00111)
Amarela	-0.00412* (.00201)	-20.5* (9.1)	-23.7* (9.23)	-0.00154 (.00157)	.000378 (.00134)
Parda	-0.00473*** (.000935)	7.96*** (1.28)	4.34*** (1.02)	.00161* (.000728)	.00382*** (.000461)
Indigena	-0.00983* (.00479)	13.5 (52.7)	5.64 (51.5)	-0.0065 (.00929)	-0.00167 (.00847)
2005	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
2006	-0.0105*** (.00121)	6.29*** (1.59)	-1.69 (1.43)	-0.0051*** (.000733)	-0.000221 (.000682)
2007	-0.00835*** (.0015)	-1.84 (1.74)	-8.14*** (1.83)	-0.00291*** (.000587)	.000942 (.000663)
2008	-0.00679*** (.00118)	2.05 (2.27)	-3.07 (2.05)	-0.003*** (.000876)	.00013 (.000739)
2009	-0.0055*** (.00121)	.29 (2.22)	-3.83* (1.94)	-0.00294*** (.000869)	-0.00042 (.000687)
2010	-0.00449** (.00151)	-4.71 (3.03)	-8.06** (2.48)	-0.00222 (.00128)	-0.000169 (.000879)
2011	.00648** (.00236)	-9.46* (4.08)	-4.39 (3.34)	-0.00188 (.00167)	-0.00498*** (.00142)
2012	.0371*** (.00328)	-3.03 (4.01)	25.6*** (3.25)	-0.00411** (.00153)	-0.0216*** (.00116)
Q1	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
Q2	.0000603 (.000382)	-4.64*** (.856)	-4.61*** (.771)	.000711 (.000409)	.000688 (.000381)

Table 2.19: The Effect of AAI on Birth outcomes (all coefficients)

	(1)	(2)	(3)	(4)	(5)
	Pre	BW	BW	LBW	LBW
Q3	-.00171* (.000689)	3.19* (1.33)	1.87 (1.11)	-.0012* (.000528)	-.000395 (.000442)
Q4	.0000569 (.000469)	-1.03 (1.17)	-.988 (.956)	.000372 (.00038)	.000346 (.000278)
Mean AAI	.0019 (.00607)	-26.4* (11)	-25* (10.2)	.0201*** (.00454)	.0192*** (.00469)
Pre-term=0			0 (.)		0 (.)
Pre-term=1			-768*** (9.04)		.469*** (.00501)
Constant	.237*** (.0144)	2,672*** (31.4)	2,853*** (21.7)	.271*** (.0138)	.16*** (.00818)
N	8,800,383	8,797,877	8,797,877	8,797,877	8,797,877
FE	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni	Y,Q,Muni
\bar{y}	.08	3,162	3,162	.085	.085
σ_y	.27	522	522	.28	.28

Standard errors (SE) in parentheses

Coefficients of 0 with missing SEs correspond to the omitted categories for categorical variables

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 2.8: Quarterly Births (2005–2012)




Figure 2.9: All-Cause, one-year Infant Mortality Rate




Figure 2.10: Output and Area Cultivated for Large Sugarcane Producers





Figure 2.11: Output and Area Cultivated for Sugarcane-Producing Regions of Brazil



CHAPTER 3

Worker Transitions in a Time of Technological Change

A long history of technological innovation is behind the jaw-dropping prosperity of the developed world. In spite of the long-term benefits, the arrival of new technology often raises concerns about job displacement. I use a large panel dataset to describe the transitions of workers as mechanical harvesting substituted for labor in Brazilian sugarcane. This episode yields several policy relevant lessons. I find that harvest workers are more likely to switch industries than occupations. Workers who previously struggled to remain employed or increase their earnings, like those with less experience or education, are even less likely to find success during the period of technological change. Finally, local labor market conditions matter greatly; the most important observable source of variation in employment outcomes is place of work.

3.1 Introduction

The development and adoption of labor-saving technology has long raised anxieties about job losses and lower wages. Queen Elizabeth I denied a patent for a knitting machine, concerned it would displace hand knitters (Conniff, 2011). Keynes fretted about “technological unemployment” in 1930, as did Leontief in 1952 (Keynes, 1963; Leontief, 1952). Recent predictions have been dire: for instance, that 45 percent or more of jobs in wealthy countries will be automated in the next two decades (Brynjolfsson and McAfee, 2014; World Bank Group, 2016; Frey and Osborne, 2017). Such concerns were at the fore as one of Brazil’s major industries rapidly switched from manual to automated production; sugarcane harvesting, which once employed hundreds of thousands of workers, was almost completely mechanized in just a few years (Osse (2002), Chapter 1). Focusing on the case of Brazilian sugarcane, this paper is among the first to use large-scale panel data to study the transitions of individual workers as technology substitutes

for their labor.

In the long run, labor reallocates to more productive pursuits as technology performs a greater number of tasks or trade sends tasks abroad. But individuals who performed these tasks may suffer lasting harm. They may find it costly to transition to new jobs, enduring lower earnings or long spells of unemployment. This paper identifies the individual characteristics that predict transitions to new, higher-earning employment, helping policymakers to devise solutions and target vulnerable populations. Specifically, I consider how demographics, firm size, previous work experience, and local labor market conditions affect the likelihood of a transition to a higher-earning job.

Sugarcane harvesting mechanized rapidly, leading to a decline in employment declined of almost 45 percent, or 115,000, even as output increased. Mechanization is evident in a dramatic shift in the occupational structure of sugarcane harvesting. Over 80 percent of sugarcane workers were engaged in manual occupations in 2007. By 2014, only about 40 percent of workers were engaged in manual occupations; the share and the count of workers operating or supporting machinery increased simultaneously. State governments in the sugarcane growing regions estimate that over 85 percent of sugarcane area was harvested mechanically in 2014.

The period of mechanization shows changes in wages and employment that imply leftward shifts in the supply of manual labor and rightward shifts in demand for other types of labor. In absolute terms, real earnings for manual sugarcane workers increased by about 20 percent between 2007 and 2014, even as employment fell by two thirds. Non-manual sugarcane workers experienced even larger increases in real earnings and an 11 percent increase in employment.

Firms of all sizes adopted mechanical harvesting, and mechanization did not bring major consolidation. In principle, the expense and lumpiness of a harvesting machine might present a barrier to adoption for smaller firms. But, as described in Chapter 1, there are rental markets and other sharing mechanisms for harvesting machines. The decline in manual harvesting is visible for even the smallest firms, although it is somewhat more pronounced for firms with more than 50 employees. It does not appear that mechanization is associated with consolidation; on the contrary, the share of workers at large firms decreased by about 10 percentage points between 2007 and 2014.

Although policymakers implemented a program to retrain sugarcane workers for new jobs in sugarcane, I find that manual sugarcane are more likely to switch industries than occupations. The sugarcane workforce became a shrinking core of experienced workers. During mechanization, the sugarcane workforce is increasingly composed of individuals with prior experience in sugarcane, and an increasing fraction of the workforce exits into other occupations and industries. I observe high occupational persistence; manual sugarcane workers are much more likely to come from or go to manual agricultural work in another industry than they are to switch occupations within sugarcane.

Mechanization sharpens a selection process that retains high-earning workers in stable jobs. Those with

previous experience in sugarcane tend to earn more and are more likely to have long-term employment, and mechanization makes these differences more pronounced. Those exiting sugarcane tend to have lower earnings, more jobs but fewer months of work, and, again, mechanization makes these differences more pronounced. Perhaps mechanization increases the incentive to shed workers who are less productive, less experienced, or less reliable.

Even though upward pressure on wages was an important factor in driving mechanization, mechanization is associated with higher earnings growth for individuals switching into manual sugarcane work and lower earnings growth for those switching out. Chapter 1 argues that increases in aggregate labor demand, essentially a more lucrative outside option, forced sugarcane growers to offer higher wages until mechanization became cheaper than manual harvesting. Workers with the lowest earnings growth were the first to leave: those switching out of manual sugarcane work during the period of mechanization had lower earnings growth than those who switched out before mechanization. This pattern may result from worker heterogeneity. For instance, the model of mechanization developed in Chapter 1 predicts that the highest-paid workers would be the first ones displaced by machinery. If these workers did not also have high-paying outside options, those displaced would suffer greater declines in earnings.

The longitudinal data employed in this analysis allow us to identify the characteristics of places and individuals that predict successful transitions in employment. Educated men are more likely to remain employed, to find employment outside of sugarcane, and to increase their earnings. This suggests women and the less educated as important populations for policymakers. One such policy might be a migration incentive. I find that those in larger labor markets are more successful at finding jobs, if less successful in securing a raise. Importantly, there are substantial differences across states in various outcomes that indicate a successful transition.


3.2 Sugarcane Cultivation and the Adoption of Mechanical Harvesting

Thanks to favorable climate and soil characteristics, a six-state region in South-Central Brazil accounts for a substantial and growing fraction of world sugarcane production. These states, which I refer to as the “study region”, accounted for 27 percent of cumulative world sugarcane production between 1998 and 2014.¹ As shown in Figure 3.1, both area harvested and output more than doubled during this period.

Between 2007 and 2014, sugarcane growers in these states transitioned from manual to mechanical harvesting. Since manual harvesting requires a much larger quantity of labor than mechanical harvesting,

¹The six states in the study region are Goiás, Minas Gerais, Paraná, Mato Grosso do Sul, Rio de Janeiro, São Paulo. Data from FAOSTAT and the Brazilian Census Bureau’s *Produção Agrícola Municipal (PAM)* survey.

Figure 3.1: Sugarcane Production and Area




I measure mechanization as the quantity of labor per hectare of sugarcane harvested. As shown in Figure 3.2, hours worked per hectare is stable through 2007, ranging from 60 to 70.² After 2007, labor intensity falls precipitously, dropping under 20 hours per hectare by 2013. This corresponds to the rapid, widespread adoption of mechanical harvesting.

I study the causes of the transition in Davis (2017), finding that rising wages played a primary role. It's not that this technology was new; harvesting machines had developed and refined in other countries, and they were available in Brazil even though growers chose not to use them. Since manual harvesting is preceded by burning the field, this practice was the target of environmental regulation which coincided with the adoption of mechanical harvesting, and government regulators claim that regulation was the causal driver of the transition in harvesting practices. However, various empirical tests undermine this claim, finding that the regulation had little effect on harvesting practices. Instead, Davis (2017) finds that strong wage growth for sugarcane harvest workers is sufficient to explain the adoption of mechanical harvesting.

²The numerator, hours worked, is drawn from the RAIS dataset, which is described in detail in the following section. The denominator, area harvested, is drawn from the Brazilian Census Bureau's Produção Agrícola Municipal (PAM) survey.

Figure 3.2: Labor intensity of sugarcane harvesting



3.3 Data

3.3.1 Collection and Coverage

The main data source for this study are the 1998–2014 editions of RAIS, a mandatory employer survey conducted annually by the Brazilian Ministry of Labor, recording all formal work in the country.³ Audits, together with punishments like fines and public shaming, provide strong incentives for complete and accurate reporting (Alvarez et al., 2017). Importantly, these data include identifiers for firms, establishments, and individuals, allowing me to track firms and individuals over time.

For historical reasons, most sugarcane workers are formal and, consequently, appear in RAIS. By construction, RAIS includes only formal employment, so, at first blush, it may not seem like an ideal source to study seasonal agricultural workers. But, in interviews, farmers and farm workers indicate that labor in the sugarcane sector is predominantly formal sector and unionized, and RAIS captures roughly 60 to 75 percent of the sugarcane employment recorded in household survey data. Direct comparisons between household survey data, the Brazilian Census Bureau’s PNAD, and RAIS are complicated for a several of reasons: i) the unit of analysis in each dataset is different, ii) quantity of labor is measured differently, and iii) many sugarcane workers are seasonal migrants. The figure I report here, 60 to 75 percent, is the national count of employment spells from RAIS divided by national count of individuals from PNAD. I use the national counts because RAIS records the place of work and PNAD records the place of residence.

³The acronym stands for *Relatório Anual de Informações Sociais*, which roughly translates to Annual Report on Social Information.

For seasonal migrants, these will not be the same so a single individual might appear in different places in each dataset. The economy as a whole is formalizing during this period; the formal sector employment share of adult men grew from 64 percent in 1996 to 74 percent in 2012 (Alvarez et al., 2017). Thus, while some trends in the data may result from increasing formalization, RAIS is nonetheless an excellent source.

3.3.2 Variable Definitions

RAIS records information about each job spell, along with some demographics and information regarding the establishment and firm.

Information about job spells includes tenure, date of separation, occupation code, average monthly earnings, and earnings for the month of December. I construct hire date based on tenure and separation date.⁴ Hours per week is also recorded, but there is no meaningful variation; virtually all employment spells indicate 44 hours per week. Since this variable does not appear to accurately measure hours, I use earnings rather than wages. Throughout this study, I report real earnings in 2014 reais (R\$) adjusted by the Brazilian Census Bureau's IPCA deflator. Given the lack of variation in hours, I measure the quantity of labor as months worked

Except where otherwise noted, I focus on workers in the sugarcane industry, often separating them into three categories based on their occupation: manual, mechanical, and other.⁵ Manual includes occupation codes for manual workers engaged in agricultural cultivation. Mechanical includes occupation codes for mechanics and operators of harvesters and related machinery. Other includes all other occupations, most commonly road-going truck drivers, managers, and office assistants.

Worker gender, age, and education are included. Education is recorded as a categorical attainment variable. When studying the transitions of individual workers, I drop those with inconsistent gender, age, or education.⁶

The data include detailed industry codes for each firm, and the firm identifier allows me to measure the characteristics of an individual firm's workforce.

⁴This process can be imprecise. Although tenure is reported to the tenth of a month, many years of data do only give the separation month, not the date. As such, I conduct all analysis by quarter or by year.

⁵There are two occupational coding schemes in the data, known by their Portuguese-language acronyms CBO 1994 and CBO 2002. These schemes are not compatible, as some individual occupations in CBO 1994 correspond to multiple occupations in CBO 2002 and vice versa. However, the coarse coding I use in the analysis, i.e. manual, mechanical, other, is consistent across years.

⁶I make exceptions for individuals under two conditions: i) the inconsistency is only in one variable, ii) at least two thirds of the individual's observations agree. Under these circumstances, I retain the individual and replace all observations of the inconsistent variable with the mode.

3.3.3 Sample

Except where otherwise noted, I study all individuals employed during the harvest season by a firm engaged in sugarcane cultivation. Specifically, I include all workers employed as of June 30 each year, which falls near the peak of the harvest. For computational ease, some analyses use a simple random sample of the population.

3.4 Changes in the Composition of Workers and Firms

3.4.1 How Did the Workforce Change During Mechanization?

Expansion and mechanization are evident from changes in employment by occupation and in demographics. Tables 3.1, 3.2, and 3.3 show employment information for manual workers, mechanical workers, and others in sugarcane cultivation. The quantity of manual labor peaks in 2007 at 1.5 million months worked, falling to just over 500,000 months worked by 2014. Meanwhile, the quantity of labor supplied by workers operating or supporting machinery increased steadily throughout the period. The average age of manual workers increased sharply after 2007, as did the fraction of female workers. This pattern is consistent with the elimination of the most physically strenuous jobs, manual cutting of sugarcane stalks.

Changes in employment and wages reveal shifts in labor demand and labor supply consistent with the argument in Chapter 1, namely increased demand for manual workers from other sectors bid up wages and led to the adoption of mechanical harvesting. Real earnings rose substantially in all occupation categories, throughout the period, and across the earnings distribution. Since employment was increasing for mechanical workers, there must have been an outward shift in the demand for mechanical workers. By contrast, after 2007, employment for manual workers fell even as earnings continued to rise, implying an inward shift in the supply of manual workers.

Table 3.1: Manual worker characteristics by year

Year	<i>N</i>	Mon. work	Pct. female	Age	P_{25} earn	P_{50} earn	P_{75} earn
1998	149,671	1,248,871	0.15	33	591	776	991
1999	143,439	1,156,139	0.13	33	585	780	995
2000	129,595	1,007,963	0.12	33	616	829	1,069
2001	139,758	1,151,934	0.12	33	631	833	1,073
2002	130,337	1,051,960	0.11	33	670	866	1,096
2003	135,104	1,117,520	0.11	33	705	900	1,129
2004	130,909	1,099,594	0.11	33	708	921	1,169
2005	149,827	1,250,515	0.11	33	745	958	1,208
2006	163,201	1,401,351	0.11	33	810	1,032	1,299
2007	176,193	1,500,179	0.12	33	838	1,062	1,326
2008	160,321	1,368,827	0.14	34	864	1,086	1,353
2009	136,305	1,220,582	0.14	34	887	1,107	1,363
2010	125,305	1,103,357	0.14	35	951	1,193	1,468
2011	104,404	918,288	0.16	36	963	1,213	1,502
2012	87,374	782,753	0.17	37	982	1,237	1,534
2013	75,114	694,737	0.18	37	1,005	1,269	1,571
2014	57,420	533,596	0.18	39	985	1,262	1,562

Table 3.2: Mechanical worker characteristics by year

Year	<i>N</i>	Mon. work	Pct. female	Age	P_{25} earn	P_{50} earn	P_{75} earn
1998	13,799	143,823	0.00	33	1,123	1,539	1,879
1999	12,585	129,998	0.00	34	1,127	1,468	1,765
2000	13,228	134,168	0.00	34	1,097	1,430	1,733
2001	14,093	145,652	0.00	34	1,159	1,522	1,847
2002	13,081	132,533	0.00	34	1,213	1,579	1,922
2003	15,039	153,291	0.00	35	1,219	1,533	1,911
2004	15,516	158,063	0.00	35	1,201	1,467	1,831
2005	18,139	184,733	0.00	34	1,256	1,527	1,862
2006	21,586	221,999	0.00	34	1,308	1,587	1,924
2007	25,544	261,020	0.00	35	1,344	1,646	2,008
2008	27,574	284,110	0.00	35	1,405	1,701	2,069
2009	26,724	279,387	0.01	35	1,455	1,773	2,115
2010	28,211	295,005	0.01	36	1,512	1,867	2,234
2011	29,356	299,936	0.01	35	1,563	1,890	2,297
2012	29,616	306,028	0.01	36	1,633	1,960	2,401
2013	32,747	343,945	0.02	36	1,690	2,036	2,471
2014	30,851	327,537	0.02	36	1,700	2,043	2,504

Table 3.3: Other worker characteristics by year

Year	<i>N</i>	Mon. work	Pct. female	Age	<i>P</i> ₂₅ earn	<i>P</i> ₅₀ earn	<i>P</i> ₇₅ earn
1998	43,945	430,152	0.07	35	851	1,319	1,854
1999	41,944	405,978	0.07	35	847	1,280	1,782
2000	42,422	395,190	0.07	35	818	1,245	1,740
2001	44,845	438,184	0.06	36	896	1,334	1,852
2002	38,302	376,179	0.06	36	969	1,399	1,916
2003	32,159	323,551	0.07	37	1,104	1,566	2,047
2004	34,491	339,320	0.07	37	1,011	1,452	1,992
2005	38,852	391,696	0.07	37	1,097	1,554	2,048
2006	45,015	457,948	0.09	37	1,101	1,605	2,111
2007	54,317	541,359	0.09	37	1,142	1,633	2,159
2008	57,827	584,679	0.08	37	1,255	1,726	2,228
2009	58,856	601,525	0.09	37	1,253	1,747	2,238
2010	55,753	577,448	0.09	38	1,386	1,908	2,385
2011	55,569	564,724	0.09	37	1,439	1,949	2,476
2012	54,682	558,830	0.10	38	1,509	2,039	2,585
2013	58,287	608,301	0.10	38	1,573	2,130	2,677
2014	52,734	563,677	0.10	39	1,625	2,183	2,738

3.4.2 Did Mechanization Encourage Consolidation? Which Firms Mechanized?

There is mixed evidence that mechanization is associated with increasing firm size or consolidation, as shown in Table 3.4. Lacking data on firm-level output, I use measure consolidation in two ways: the count of firms with more than 50 workers and employment at firms with more than 1,500 workers as a fraction of total sugarcane employment.⁷ First, we see that the count of firms with greater than 50 employees falls from 127 in 2008 to 104 in 2014. This change could reflect consolidation among large firms, or it could be that mechanizing firms slip below the 50 employee threshold. Indeed, the count of firms below 50 employees increases substantially, possibly because firms are getting smaller or due to formalization, e.g. among family farms. The second measure suggests declining consolidation. Firms larger than 1,500 workers account for 85 percent of sugarcane employment in 2007 but only 75 percent in 2014. Total employment at these firms fell by more than half. The fraction of sugarcane workers employed by firms between 500 and 1,500 employees grew by 7 percentage points, while total employment remained relatively stable.

Mechanization is evident in firms of all sizes after 2007. I measure the mechanization of firms as the

⁷Table 3.4 separates firms by a count of workers at each firm. The size categories are i) 50 or fewer workers, ii) 50 to 500 workers, iii) 500 to 1,500 workers, and iv) more than 1,500 workers. Separating firms by quantiles of employment gives qualitatively similar results.

fraction of workers who are in manual occupations. For firms with more than 50 employees, about 70 percent of workers are in manual occupations in 2007. By 2014, it was close to 40 percent. Even the smallest firms show evidence of mechanization; they had 40 to 50 percent of their workers in manual occupations from 1998 to 2007, but only 32 percent as of 2014.

Table 3.4: Firm-level workforce characteristics by firm employment

Year t	N_{firm}	$\sum X$	$X \leq 50$				$50 < X \leq 500$					
			$\frac{\sum X}{N_{firm}}$	Frac. of Tot. Emp.	Frac. Man.	P_{50} earn	N_{firm}	$\sum X$	$\frac{\sum X}{N_{firm}}$	Frac. of Tot. Emp.	Frac. Man.	P_{50} earn
1998	129	1,443	11	0.01	0.53	749	55	9,321	169	0.04	0.75	862
1999	147	1,653	11	0.01	0.52	797	44	8,382	191	0.04	0.78	846
2000	135	1,340	10	0.01	0.44	846	45	9,177	204	0.05	0.65	934
2001	129	1,257	10	0.01	0.44	896	48	10,266	214	0.05	0.72	858
2002	154	1,444	9	0.01	0.40	958	45	8,784	195	0.05	0.68	885
2003	197	1,879	10	0.01	0.46	954	46	7,877	171	0.04	0.69	928
2004	193	1,667	9	0.01	0.45	907	50	9,010	180	0.05	0.65	954
2005	230	2,321	10	0.01	0.48	970	41	7,806	190	0.04	0.68	836
2006	304	2,741	9	0.01	0.44	1,028	55	8,441	153	0.04	0.68	973
2007	445	3,324	7	0.01	0.45	991	67	10,663	159	0.04	0.70	1,047
2008	523	3,607	7	0.01	0.44	1,046	73	12,028	165	0.05	0.69	1,158
2009	569	3,500	6	0.02	0.40	1,127	64	10,724	168	0.05	0.60	1,230
2010	587	3,700	6	0.02	0.40	1,171	55	8,067	147	0.04	0.51	1,333
2011	564	3,633	6	0.02	0.41	1,204	57	8,040	141	0.04	0.37	1,449
2012	574	3,632	6	0.02	0.36	1,360	59	8,817	149	0.05	0.51	1,462
2013	606	3,757	6	0.02	0.36	1,338	55	7,525	137	0.05	0.43	1,560
2014	592	3,453	6	0.02	0.32	1,435	59	7,982	135	0.06	0.44	1,616

Table 3.4: Firm-level workforce characteristics by firm employment (*continued*)

Year t	N_{firm}	$\sum X$	$500 < X \leq 1500$				$1500 < X$					
			$\frac{\sum X}{N_{firm}}$	Frac. of Tot. Emp.	Frac. Man.	P_{50} earn	N_{firm}	$\sum X$	$\frac{\sum X}{N_{firm}}$	Frac. of Tot. Emp.	Frac. Man.	P_{50} earn
1998	42	38,952	927	0.19	0.73	943	34	157,714	4,639	0.76	0.72	841
1999	41	39,470	963	0.20	0.72	911	31	148,463	4,789	0.75	0.72	850
2000	39	35,360	907	0.19	0.71	962	32	139,368	4,355	0.75	0.70	898
2001	42	39,100	931	0.20	0.72	997	30	148,073	4,936	0.75	0.70	914
2002	37	36,384	983	0.20	0.73	983	24	135,108	5,630	0.74	0.72	956
2003	31	30,577	986	0.17	0.74	1,050	25	141,973	5,679	0.78	0.75	982
2004	31	31,552	1,018	0.17	0.77	1,052	20	138,687	6,934	0.77	0.72	1,014
2005	30	29,039	968	0.14	0.73	1,096	23	167,652	7,289	0.81	0.73	1,064
2006	27	27,925	1,034	0.12	0.73	1,130	22	190,695	8,668	0.83	0.71	1,145
2007	23	23,500	1,022	0.09	0.69	1,200	28	218,567	7,806	0.85	0.69	1,183
2008	26	25,311	974	0.10	0.64	1,227	28	204,777	7,313	0.83	0.65	1,246
2009	28	27,132	969	0.12	0.61	1,276	28	180,529	6,447	0.81	0.62	1,278
2010	26	26,517	1,020	0.13	0.62	1,357	32	170,985	5,343	0.82	0.60	1,395
2011	25	22,618	905	0.12	0.58	1,454	31	155,038	5,001	0.82	0.56	1,462
2012	28	25,499	911	0.15	0.53	1,493	25	133,730	5,349	0.78	0.51	1,556
2013	29	27,548	950	0.17	0.50	1,600	23	127,327	5,536	0.77	0.45	1,667
2014	26	24,163	929	0.17	0.44	1,609	19	105,420	5,548	0.75	0.40	1,754

3.5 Mobility of Workers in the Sugarcane Industry

3.5.1 Employment, Fluidity, and Tenure

Historically, sugarcane growers relied on seasonal labor during the May to October harvest; as machines replaced harvest labor, the seasonality in employment declined substantially and the remaining jobs in sugarcane became more long-term. Seasonal patterns are clearly visible in Figure 3.3. Employment rises through the first two quarters of each year, plateaus during the third quarter, and then crashes during the fourth quarter. These patterns are especially pronounced for manual workers; in many years, more than half of manual work is seasonal. Figure 3.3 also makes clear how mechanization is reflected in the level and composition of employment. Manual work declined precipitously even as non-manual work showed meaningful increases. Seasonality declined for two reasons. First, manual, which had stronger seasonal patterns, declined as a fraction of sugarcane employment.

The reduced seasonality associated with mechanization is primarily attributable to reduced hiring rates in the first half of the year and reduced separation rates in the last quarter. To make an analogy to demography, Figure 3.3 plots the population, while Figure 3.5 shows birth rates and death rates; specifically, Figure 3.5 shows hiring and separation rates separately by quarter from 1998 to 2014. Employment rises when the hiring rate (black line) is above the separation rate (red line), and falls when the hiring rate is below the separation rate. The seasonal pattern is clear. As sugarcane firms prepare for the harvest during the first half of the year, hiring rates can be two to three times as large as separation rates. Hiring and separation rates are both low during the third quarter, the peak of the harvest. Separation rates spike dramatically in the fourth quarter as sugarcane growers shed a substantial fraction of their workforce. The period of mechanization is characterized by steep declines in hiring rates during the first two quarters, and steep declines in separation rates in the fourth quarter.

This reduction in seasonal labor is easily visible in the distribution of job tenure; the remaining jobs in sugarcane tend to be more permanent. Figure 3.4 plots the CDF of tenure at the end of Q2 in 1998, 2007, and 2014. In 1998 and 2007, about 60 percent of workers had less than a year of tenure, but by 2014 only 30 percent of workers had less than a year of tenure.

Figure 3.3: Aggregate Employment in Sugarcane by Quarter




Figure 3.4: Distribution of Job Tenure Over Time







Figure 3.5: Hire and Separation Rates by Quarter




(a) Q1



(b) Q2



(c) Q3



(d) Q4

3.5.2 Where Do Sugarcane Workers Come From? Go?

The longitudinal nature of the RAIS data allow us to identify the industries and occupations that are sources and destinations for sugarcane workers, and how these flows might have changed during the period of mechanization. During the period of mechanization, the sugarcane workforce was increasingly composed of workers with recent experience as manual sugarcane workers, rather than those transitioning from other industries and occupations.

Around 80 percent of manual sugarcane workers in year t were, in year $t - 1$, manual sugarcane workers or outside the formal labor force. A similar fraction go on to be manual sugarcane workers or to exit the formal labor force in period $t + 1$. Table 3.5 describes the flows of individuals into and out of manual sugarcane work, prior to (1998–2006) and during the period of mechanization (2007–2014). The table is constructed by taking all manual sugarcane workers in period t and calculating the fraction of those workers who came from several industry-occupation categories in year $t - 1$ (left panel) or went to the same categories in year $t + 1$ (right panel). For example, the first row indicates that, from 1998–2006, 52.5 percent of manual sugarcane workers in year t were also doing manual sugarcane work in the prior year. Only 1.5 percent worked in some other sugarcane occupation, while 14.3 percent worked in other industries, and 31.7 percent had no *formal sector* job. Of the manual sugarcane workers in year t , 53.1 would do manual sugarcane work the following year, 2.4 percent would take a different occupation in sugarcane, and 16.9 percent would move on to jobs in other industries. Some 27.5 percent leave the formal workforce the year after manual sugarcane work.⁸

These flows change during the period of migration; fewer manual workers are drawn from other industries or occupations, and fewer manual workers continue in the following year. Manual sugarcane workers are less likely to come from other occupations in sugarcane, other industries, or from outside the formal labor force. Manual sugarcane workers are 5.5 percentage points less likely to remain in that job in the following year, with more of them moving to other occupations in sugarcane, to other industries, and outside the formal labor force.

Admittedly these changes are small compared to the magnitudes of the flows; even during mechanization, nearly half of the workforce in any year is drawn from other industries or occupations, and nearly half of the workforce continues manual sugarcane work the following year. We might have expected mechanization to reduce the inflows and increase outflows more sharply.⁹ These figures may overstate

⁸For compactness, Table 3.5 averages each column across years. For the year-by-year rates, see Table ??.

⁹Indeed, the inflows of about 45.0 percent and outflows of 52.4 percent during the period of mechanization scarcely seem able to explain the decline in employment shown in Table 3.1. There are two reasons that this apparently small difference in flows accounts for a decrease in employment of about 67 percent. First, and most obviously, the outflows exceed inflows for many years. Second, the inflows appear larger during periods of declining employment because they are calculated as a fraction of employment in year t . For example, consider the change in employment from 2007 to 2008. Employment in 2008 equals employment in 2007, less outflows from 2007, plus inflows from 2008. If employment is declining from 2007 to 2008, the inflows will appear larger since they are calculated as a fraction of the smaller 2008 employment, whereas the outflows are calculated as a fraction of the larger 2007 employment.

Table 3.5: Occupation and Industry Transition Rates for Manual Sugarcane Workers (period means)

Year t	$t - 1$				$t + 1$			
	Cane Ind.		Oth. Ind.	No job	Cane Ind.		Oth. Ind.	No job
	Man. Occ.	Oth. Occ.	Oth. Occ.		Man. Occ.	Oth. Occ.	Ind: Oth.	
1998–2006	0.525	0.015	0.143	0.317	0.531	0.024	0.169	0.275
2007–2014	0.550	0.012	0.139	0.299	0.476	0.031	0.182	0.312

the case if those with no formal sector jobs in $t - 1$ or $t + 1$ were working informally as manual sugarcane workers, but, to join or exit the formal labor force, those workers must have switched firms. With its persistently high inflows and outflows, manual sugarcane work remained very fluid during the period of mechanization.

Tables 3.6–3.9 disaggregate the information in Table 3.5 by industry and occupation.¹⁰ As in Table 3.5, each cell indicates the fraction of manual sugarcane workers in year t who worked in an industry or occupation in either $t - 1$ or $t + 1$.¹¹ Tables 3.6–3.7 cover industry and occupation for the prior year of work ($t - 1$), while Tables 3.8–3.9 cover the following year ($t + 1$). The rows list the most commonly observed industries for manual sugarcane workers, and the columns list the most commonly observed occupations.¹² Blank cells in the tables correspond to combinations of industry and occupation are never observed as destinations or origins for manual sugarcane workers.

Policymakers devoted resources to the idea that manual sugarcane workers would adapt to mechanization by becoming machine operators, but manual sugarcane workers are more likely to switch industries than occupations.¹³ The fraction of manual sugarcane workers who go to or come from the manual agriculture occupation is always higher, by about 10 percentage points, than fraction who go to or come from the sugarcane cultivation industry. The inequality holds even if we group sugarcane cultivation with the

¹⁰Table 3.5 does not precisely match Tables 3.6–3.9. For example, the first column, first row of Table 3.5 measures the same quantity as the first row, first column of Table 3.6, but the values differ by 20 basis points. Differences arise because Table 3.5 presents the averages across years. Rather than constructing annual values and taking an average, Tables 3.6–3.9 use aggregate counts from the entire period, so they are effectively weighted averages where the weight is annual employment.

¹¹Tables 3.6–3.7 do not sum to one; the remaining fraction of workers had no formal sector job in $t - 1$ or $t + 1$.

¹²The industries are i) sugarcane cultivation, ii) sugarcane milling & ethanol production, iii) agricultural services, iv) construction, v) citrus cultivation, vi) cattle ranching, vii) coffee cultivation, viii) ground transportation, ix) temporary or contract labor, x) business services, and xi) other. The occupations are i) manual agricultural worker, ii) agricultural machine operator, iii) demolition and construction, iv) truck driver, v) production line feeder, and vi) other. These are the most commonly observed industries for individuals in the year before and the year after manual sugarcane work.

¹³For instance, the Inter-American Development Bank and the Brazilian Sugarcane Industry Association (known by its Portuguese-language acronym UNICA) partnered on a retraining program that launched in 2009. See <http://unica.com.br/news/706475192031019628/iadb-approves-early-results-of-sugarcane-industry-por-centoE2-por-cento80-por-cento99s-worker-requalification-program/>.

Table 3.6: Industries and Occupations in $t - 1$ for Manual Cane Workers in t ($t \in [1999, 2007]$)

	Manual ag.	Mach. op.	Demo. / cons.	Truck driver	Prod. line	Other	Ind. Total
Cane cult.	0.523	0.001	0.000	0.001	0.000	0.011	0.536
Cane mill & Ethanol	0.043	0.000	0.000	0.000	0.000	0.003	0.047
Ag. services	0.018	0.000	0.000		0.000	0.001	0.019
Construction		0.000	0.001	0.000		0.000	0.002
Citrus	0.004	0.000				0.000	0.005
Beef	0.003	0.000	0.000	0.000	0.000	0.001	0.004
Coffee	0.005	0.000				0.000	0.005
Ground transport	0.007	0.000	0.000	0.000		0.001	0.007
Contract labor	0.003		0.000	0.000	0.000	0.001	0.004
Bus. services	0.004	0.000	0.000	0.000		0.001	0.006
Other	0.026	0.000	0.001	0.000	0.000	0.018	0.047
Occ. Total	0.635	0.002	0.003	0.001	0.001	0.038	0.680

closely-related industry of sugarcane milling and ethanol production. The transition rate from manual sugarcane worker to machine operator increases by about 50 percent during the period of mechanization, it peaks at only 1.3 percent.

Changes in disaggregated inflows and outflows are generally small and diffuse. The primary inflow industries for manual sugarcane workers are sugarcane cultivation and sugarcane milling & ethanol production. The primary inflow occupation is manual agricultural work. During the period of mechanization, workers are several percentage points more likely to have come from those industries or that occupation. The only other notable change is the 3 percentage point decrease in the inflow of workers from outside the formal labor force. The primary outflow industries are the same, as is the primary outflow occupation. The period of mechanization brings a decrease in these outflows of about 5-7 percentage points, along with a 3.6 percentage point increase in the fraction of workers leaving the formal labor force. The period of mechanization also shows a 1 percentage point uptick in the outflow to demolition and construction occupations.

Table 3.7: Industries and Occupations in $t - 1$ for Manual Cane Workers in t ($t \in [2008, 2014]$)

	Manual ag.	Mach. op.	Demo. / cons.	Truck driver	Prod. line	Other	Ind. Total
Cane cult.	0.563	0.001	0.000	0.000	0.000	0.006	0.571
Cane mill & Ethanol	0.064	0.000	0.000	0.000	0.000	0.003	0.068
Ag. services	0.012	0.000	0.000	0.000	0.000	0.000	0.013
Construction	0.000		0.002			0.001	0.003
Citrus	0.005	0.000				0.000	0.005
Beef	0.005	0.000			0.000	0.001	0.006
Coffee	0.004	0.000				0.000	0.004
Ground transport	0.001	0.000	0.000	0.000	0.000	0.001	0.002
Contract labor	0.000		0.000		0.000	0.000	0.001
Bus. services	0.001		0.000	0.000		0.000	0.001
Other	0.011	0.000	0.003	0.000	0.002	0.020	0.036
Occ. Total	0.666	0.003	0.006	0.001	0.003	0.032	0.710

Table 3.8: Industries and Occupations in $t + 1$ for Manual Cane Workers in t ($t \in [1998, 2006]$)

	Manual ag.	Mach. op.	Demo. / cons.	Truck driver	Prod. line	Other	Ind. Total
Cane cult.	0.532	0.006	0.000	0.002	0.000	0.015	0.556
Cane mill & Ethanol	0.072	0.001	0.000	0.000	0.000	0.006	0.079
Ag. services	0.016	0.000	0.000	0.000		0.001	0.017
Construction	0.000		0.001	0.000	0.000	0.001	0.002
Citrus	0.003	0.000		0.000		0.000	0.004
Beef	0.003	0.000	0.000	0.000	0.000	0.000	0.003
Coffee	0.004	0.000				0.000	0.004
Ground transport	0.006	0.000	0.000	0.000		0.001	0.007
Contract labor	0.002		0.000		0.000	0.001	0.003
Bus. services	0.003	0.000	0.000	0.000	0.000	0.001	0.005
Other	0.021	0.000	0.001	0.000	0.001	0.022	0.046
Occ. Total	0.662	0.008	0.003	0.003	0.002	0.048	0.726

Table 3.9: Industries and Occupations in $t + 1$ for Manual Cane Workers in t ($t \in [2007, 2013]$)

	Manual ag.	Mach. op.	Demo. / cons.	Truck driver	Prod. line	Other	Ind. Total
Cane cult.	0.479	0.010	0.000	0.002	0.001	0.015	0.507
Cane mill & Ethanol	0.074	0.002	0.000	0.001	0.001	0.006	0.084
Ag. services	0.015	0.001	0.000	0.000	0.000	0.001	0.017
Construction	0.000	0.000	0.005		0.000	0.002	0.008
Citrus	0.005	0.000		0.000	0.000	0.000	0.006
Beef	0.004	0.000	0.000	0.000	0.000	0.001	0.005
Coffee	0.003	0.000				0.000	0.004
Ground transport	0.001	0.000	0.000	0.000	0.000	0.001	0.003
Contract labor	0.000		0.000		0.000	0.001	0.001
Bus. services	0.001		0.000		0.000	0.001	0.001
Other	0.011	0.001	0.007	0.000	0.003	0.032	0.055
Occ. Total	0.594	0.013	0.013	0.004	0.006	0.059	0.690

3.5.3 How Did the Characteristics and Outcomes of Workers Change?

The endogenous choices of workers and firms mean that a two-sided selection process operates as workers enter and leave manual sugarcane employment. In this section, I analyze how the outcomes and characteristics of workers differ according to their previous work experience, i.e. selection into manual sugarcane employment, and by their subsequent work experience, i.e. selection out of manual sugarcane employment. Specifically, I estimate a series of regressions with the form:

$$\begin{aligned}
 y_{i,t} = & \beta_0 + \beta_1 \mathbf{1}[Ind_{i,t-1} = Cane, Occ_{i,t-1} = Oth.] + \beta_2 \mathbf{1}[Ind_{i,t-1} = Not\ cane] + \beta_3 \mathbf{1}[No\ job_{t-1}] \\
 & + \beta_4 \mathbf{1}[t > 2007] + \mathbf{1}[t > 2007] \times \left(\beta_5 \mathbf{1}[Ind_{i,t-1} = Cane, Job_{i,t-1} = Oth.] \right. \\
 & \left. + \beta_6 \mathbf{1}[Ind_{i,t-1} = Not\ cane] + \beta_7 \mathbf{1}[No\ job_{t-1}] \right) + \varepsilon_{i,t}.
 \end{aligned}
 \tag{3.1}$$

To clearly reveal how the selection process changed during the period of mechanization, this regression is constructed like a difference-in-difference specification. The sample includes all individuals with a manual occupation in the sugarcane industry as of year t , and the outcomes y are measured in year t . The “treatment” analog is work experience in period $t - 1$. As in Table 3.5, there are four categories: i) manual occupation in sugarcane, ii) other occupation in sugarcane, iii) other industry, and no formal job in $t - 1$. The omitted category is i), manual sugarcane workers. I consider the following outcomes: log monthly earnings in sugarcane, log total earnings in year t , months worked, number of jobs held, age,

and gender.¹⁴

Coefficients β_1 , β_2 , and β_3 indicate how sugarcane workers differ prior to mechanization depending on their most recent work experience. For example, referring to Table 3.10, the estimate in the first column and first row indicates that workers from non-manual occupations in sugarcane are positively selected compared to those from manual occupations in sugarcane. Specifically, log earnings in t were 0.089 higher for workers from non-manual occupations in $t - 1$, compared to those who were manual workers in $t - 1$. The analog for “post period” or “treatment period” is the period of mechanization, i.e. years after 2007. Coefficients β_5 , β_6 , and β_7 will indicate if selection patterns changed during the period of mechanization.

Those with previous experience in sugarcane earn more and are more likely to have stable, long-term employment, and mechanization makes these differences more pronounced. As you can see in the first two columns of Table 3.10, those coming from outside sugarcane (β_2) or the formal labor force (β_3) earn less, and this gap widens during the period of mechanization. The middle two columns show that those coming from outside industries or the formal labor force are employed fewer months of the year, and are more likely to work multiple jobs. There are compositional differences as well: those with previous experience in sugarcane are older and more likely to be female. As with the employment outcomes, these differences become more pronounced during the period of mechanization.

There are a number of reasons benefits might accrue to workers with previous sugarcane experience. Experience may make workers more productive. Seniority might be rewarded per se, as predicted in a range of bargaining and contracting models.¹⁵ Some types of workers may be more productive than others, and it may take experience for workers or firms learn a worker’s type. It is possible that young males both earn less and switch industries more. Finally, sugarcane cultivation was expanding beyond its historical range during this period; these newly cultivated areas might have been less productive and faced a limited supply of experienced sugarcane workers.

As shown in Table 3.11, those who exit sugarcane tend to have lower earnings, more jobs but fewer months of work, and, again, mechanization makes these differences more pronounced. They also tend to be younger and more male. As before, a number of forces may be at work. Unproductive types might be revealed and move on to other work. Young men might be more mobile but command lower earnings. Low productivity in certain areas might induce workers to switch industries, in part because of the limited earnings potential in sugarcane.

¹⁴I divide total earnings by 12 to make the units comparable to monthly earnings in sugarcane, but total earnings captures both differences in monthly earnings and differences in months worked.

¹⁵Topel (1991) summarizes various theories predicting returns to seniority.

Table 3.10: Characteristics of Manual Sugarcane Workers in t by Employment in $t - 1$

	Log Earn Cane	Log Earn All	Mon. Work	No. Jobs	Age	Female
Cane, Not manual (β_1)	0.089 (0.008)	0.068 (0.010)	-0.318 (0.048)	0.047 (0.010)	-0.856 (0.218)	-0.063 (0.007)
Other industry (β_2)	-0.024 (0.003)	-0.160 (0.004)	-1.142 (0.016)	0.201 (0.004)	-2.659 (0.075)	-0.034 (0.002)
No job (β_3)	-0.124 (0.002)	-0.355 (0.003)	-1.737 (0.012)	0.093 (0.003)	-4.429 (0.056)	-0.005 (0.002)
Post (β_4)	0.252 (0.002)	0.315 (0.003)	0.611 (0.012)	-0.064 (0.003)	3.407 (0.057)	0.042 (0.002)
Post \times (Cane, Not) (β_5)	0.061 (0.010)	0.034 (0.014)	-0.316 (0.063)	0.052 (0.014)	-1.167 (0.289)	-0.033 (0.009)
Post \times (Oth. Ind.) (β_6)	-0.004 (0.003)	-0.190 (0.005)	-1.648 (0.022)	0.288 (0.005)	-5.474 (0.099)	-0.065 (0.003)
Post \times No job (β_7)	-0.085 (0.003)	-0.330 (0.004)	-2.015 (0.017)	0.133 (0.004)	-6.818 (0.076)	-0.017 (0.002)
N	303,060	303,686	308,948	308,948	308,948	308,948
\bar{Y}_t cane, man.	6.85	6.57	9.54	1.20	34.92	0.13

Standard errors in parentheses

Table 3.11: Characteristics of Manual Sugarcane Workers in t by Employment in $t + 1$

	Log Earn Cane	Log Earn All	Mon. Work	No. Jobs	Age	Female
Cane, Not manual (β_1)	0.094 (0.006)	0.141 (0.008)	0.295 (0.037)	0.015 (0.008)	-3.042 (0.168)	-0.074 (0.005)
Other industry (β_2)	-0.035 (0.002)	-0.151 (0.003)	-1.012 (0.016)	0.158 (0.003)	-2.555 (0.072)	-0.034 (0.002)
No job (β_3)	-0.160 (0.002)	-0.387 (0.003)	-1.583 (0.013)	0.036 (0.003)	-1.456 (0.060)	0.009 (0.002)
Post (β_4)	0.256 (0.002)	0.335 (0.003)	0.764 (0.013)	-0.057 (0.003)	3.396 (0.057)	0.042 (0.002)
Post \times (Cane, Not) (β_5)	0.081 (0.006)	0.097 (0.009)	0.011 (0.040)	0.116 (0.009)	-5.520 (0.182)	-0.104 (0.005)
Post \times (Oth. Ind.) (β_6)	0.007 (0.003)	-0.159 (0.004)	-1.441 (0.019)	0.230 (0.004)	-5.473 (0.086)	-0.069 (0.003)
Post \times No job (β_7)	-0.072 (0.002)	-0.305 (0.003)	-1.822 (0.016)	0.094 (0.003)	-4.194 (0.072)	-0.020 (0.002)
N	316,919	317,574	322,636	322,636	322,636	322,636
\bar{Y}_t cane, man.	6.81	6.52	9.35	1.22	33.92	0.13

Standard errors in parentheses

Equation (3.1) estimates how sugarcane workers in year t compare to each other based on their employment status in year $t - 1$ or $t + 1$. We can also estimate the change in individual outcomes based on employment status. For example, do earnings increase for individuals who switch into sugarcane? To do so, I estimate the following modification of Equation (3.1):

$$y_{i,t} - y_{i,t-1} = \beta_0 + \beta_1 \mathbf{1}[Ind_{i,t-1} = Cane, Occ_{i,t-1} = Oth.] + \beta_2 \mathbf{1}[Ind_{i,t-1} = Not\ cane] + \beta_4 \mathbf{1}[t > 2007] + \mathbf{1}[t > 2007] \times \left(\beta_5 \mathbf{1}[Ind_{i,t-1} = Cane, Job_{i,t-1} = Oth.] + \beta_6 \mathbf{1}[Ind_{i,t-1} = Not\ cane] + \right) + \varepsilon_{i,t}. \quad (3.2)$$

The key difference is that the outcome is now an individual-level change. Using an individual-level change restricts the sample since we must observe earnings in consecutive periods for the same individual. Therefore, this regression excludes any workers outside the formal labor force in $t - 1$ or $t + 1$, and the coefficients β_3 and β_7 do not appear. The other change is to the sugarcane earnings variable. The sample is still defined as manual sugarcane workers in year t , but workers who enter or exit the sugarcane industry will not have earnings from sugarcane in $t - 1$ or $t + 1$. Therefore, I use the difference in log earnings from the longest-tenured active job in Q2 of each year, which is manual sugarcane employment in t .

It would not be surprising if manual sugarcane workers enjoyed a compensating differential, given the physical risks and demands of the job. Indeed, those who switch in to manual sugarcane enjoy substantial gains in earnings and employment. Those exiting to other industries experience lower earnings growth and lower employment. As shown in Table 3.12, workers joining sugarcane from another industry experience Q2 earnings growth about 0.13 log points, a substantial 0.091 log points higher than workers previously in manual sugarcane work. Those joining sugarcane also show a larger increase in months worked and, consequently, annual earnings growth of 0.203 log points. Table 3.13 shows that workers leaving sugarcane experience wage growth of 0.015 log points, which is 0.025 log points below the wage growth enjoyed by workers remaining in sugarcane in $t + 1$.

Evidence suggests upward pressure on wages was the primary explanation for mechanization, and yet the earnings and employment advantages of manual sugar work grow larger during the period of mechanization (see Chapter 1). Those entering manual sugarcane work have even higher earnings growth during the period of mechanization, per the estimate of β_6 in Table 3.12. The estimate of β_6 in Table 3.13 shows that those leaving the sugarcane industry suffer additional earnings losses during the period of mechanization. Changes in months worked and the number of jobs follow similar patterns. Table 3.13 shows that moving into another occupation is rewarded even more during the period of mechanization, as some workers transition to higher earning jobs as machine operators or managers. But we know from Table 3.9 that this is a small flow. During the period of mechanization, only about 2 percent of manual sugarcane workers in year t go on to other sugarcane jobs in $t + 1$.

There are several candidate explanations. The first is union power. Many manual sugarcane workers are unionized, and union-driven restrictions in labor supply would result in above-market returns to sugarcane work, driving up earnings and possibly stimulating mechanization. However, with mechanization a widely-used substitute, this behavior would seem myopic. Moreover, given that half of manual sugarcane workers in any year came from other industries or from outside the formal labor force, it would be surprising if unions were able to restrict supply.

The second candidate explanation is mismeasurement. Workers who exit the formal labor market may, in fact, enjoy higher earnings growth during the period of mechanization. Recall that this outflow is relatively large; from 2007–2014, roughly 30 percent of manual sugarcane workers leave the formal labor force the following year. Third, workers may demand a higher wage if they perceive that the trend towards mechanization limits their future employment opportunities in manual sugarcane work. Why learn the skills and develop a network if you expect to be replaced by a machine? Finally, we might expect the highest-paid workers to be the first ones to be displaced by mechanization. If all manual sugarcane workers have similar outside options, those displaced by mechanization would experience larger earnings losses.

Table 3.12: Change in Characteristics of Manual Cane Workers $Y_t - Y_{t-1}$ by Employment in $t - 1$

	Δ Log Earn Q2	Δ Log Earn All	Δ Mon. Work	Δ No. Jobs
Cane, Not manual (β_1)	-0.002 (0.006)	0.028 (0.009)	0.073 (0.045)	0.044 (0.012)
Other industry (β_2)	0.091 (0.002)	0.143 (0.003)	0.270 (0.015)	0.024 (0.004)
Post (β_4)	-0.006 (0.002)	-0.022 (0.002)	-0.159 (0.012)	-0.008 (0.003)
Post \times (Cane, Not) (β_5)	-0.017 (0.008)	-0.017 (0.012)	-0.053 (0.059)	0.032 (0.016)
Post \times (Oth. Ind.) (β_6)	0.075 (0.003)	0.100 (0.004)	0.120 (0.020)	0.051 (0.006)
N	206,887	207,481	211,922	211,922
$\Delta \bar{Y}$ cane, man.	0.04	0.06	0.20	-0.01

Standard errors in parentheses

Table 3.13: Change in Characteristics of Manual Cane Workers $Y_{t+1} - Y_t$ by Employment in $t + 1$

	Δ Log Earn Q2	Δ Log Earn All	Δ Mon. Work	Δ No. Jobs
Cane, Not manual (β_1)	0.112 (0.004)	0.120 (0.007)	0.165 (0.035)	0.013 (0.010)
Other industry (β_2)	-0.025 (0.002)	-0.051 (0.003)	-0.069 (0.015)	0.086 (0.004)
Post (β_4)	-0.006 (0.002)	-0.021 (0.002)	-0.155 (0.012)	-0.007 (0.003)
Post \times (Cane, Not) (β_5)	0.137 (0.005)	0.169 (0.007)	0.334 (0.038)	-0.094 (0.011)
Post \times (Oth. Ind.) (β_6)	-0.023 (0.002)	-0.042 (0.003)	-0.106 (0.018)	0.092 (0.005)
N	223,603	224,247	228,829	228,829
$\Delta \bar{Y}$ cane, man.	0.04	0.05	0.18	-0.02

Standard errors in parentheses

3.5.4 Who Successfully Transitions Out of Manual Sugarcane?

There are many potential causes of worker displacement, from labor-saving technology to international trade. While such changes might be positive for the aggregate economy, displacement can have negative consequences for the affected individuals. Citizens and policymakers should be concerned about their welfare on moral grounds, and because easing transitions might reduce political resistance to changes that improve aggregate welfare. Thus, for both the design and targeting of policies directed towards displaced workers, it is important to understand the characteristics of individuals and labor markets associated with successful transitions. To do so, I estimate the following linear probability model on a sample of all manual sugarcane workers between the ages of 18 and 45:

$$y_{i,m,t+1} = \beta_0 + \beta_1 X_{i,t} + \beta_2 Z_{m,t} + \gamma_t + \omega_s + \varepsilon_{i,t+1} \quad (3.3)$$

where i indexes individuals, m indexes municipalities, t indexes year, and s indexes states. I include year and state fixed effects γ and ω . Separately for 1998–2006 and 2007–2014, I regress a four employment-transition outcomes on individual characteristics and characteristics of the local labor market.

The outcomes y are indicators that assume a value of 1 if sugarcane worker i in year t : i) has a formal-sector job any time in $t + 1$, ii) has a formal-sector job outside the sugarcane industry any time in $t + 1$, iii) has a formal-sector job any time in $t + 1$ and earned more in $t + 1$ than in t , or iv) worked more than 9 months in $t + 1$.¹⁶ Employment, especially employment in a new industry, with higher earnings, or for a longer portion of the year, might all be considered desirable outcomes for workers in an industry facing displacement of labor.

The individual characteristics X are: i) a quadratic in age, ii) an indicator for female, and iii) indicator for completion of elementary education.¹⁷ Older workers may have better developed professional networks, making it more easy to switch jobs. Conversely, younger works might be more adaptable. Cultural factors, especially surrounding family and childbearing, may affect the relative mobility of men and women. It seems likely that higher levels of education are associated with improved employment outcomes.

The local labor market characteristics Z are: i) industry concentration, as measured by the Herfindahl-Hirschman Index (HHI) of employment by one-digit industry classification, ii) occupation concentration, as measured by the HHI of employment by one-digit occupation classification, and iii) the size of the local labor market, as measured by the log of employment in each municipality.¹⁸ The diversity of local employment may provide opportunities when a particular industry or occupation is affected by a

¹⁶Nine months is the median months worked for manual sugarcane workers and appears to be a threshold for seasonal vs. full-time work.

¹⁷The data have a more disaggregated education measure, but it can be inconsistent within adults over time. This coarser measure has improved consistency while still capturing most of the variation within sugarcane workers.

¹⁸HHI is defined as $\sum_i \left(\frac{x_i}{x}\right)^2$ where i denotes industry or occupation, x_i denotes employment in that industry or occupation, and $x = \sum_i x_i$ denotes total employment (Hirschman, 1964).

displacement-inducing change. A thick local labor market may also improve the likelihood of a quick, profitable transition.

Table 1 shows the results for the pre-mechanization period $t \in [1998, 2006]$ while Table 2 shows the results for the mechanization period $t \in [2007, 2013]$. All standard errors are clustered at the municipality level.

Educated men are more likely to remain employed, to find employment outside of sugarcane, to increase their earnings. By contrast, those who less educated or female are less likely to enjoy these outcomes. Women are about 5 percentage points less likely to be employed or employed outside of sugarcane. Those with at least an elementary education show improvements of 3–4 percentage points across all outcomes. The effects of age are generally modest; in the first three columns of Table 3.14, the difference in outcomes between an 18 year old and a 45 year old is no more than 2 percentage points. Column 4 is an exception. Compared to 18 year olds, a 45 year old is about 30 percentage points more likely to be seasonal in $t + 1$.

Place also matters. Although not printed in the tables, the difference between states is at least 10 percentage points for every outcome. Larger labor markets make it easier to find a job, but harder to get a raise. Doubling the size of the labor market makes increases the likelihood of non-cane employment by about 1.25 percentage points, while lowering the likelihood of getting a raise by about 0.6 percentage points. Doubling the size of the labor market increases the likelihood of finding non-seasonal work by about 3.5 percentage points. In interpreting the coefficients on industry and occupation HHI, note that HHI varies from 0 to 1, and it is increasing in concentration. Taking the signs at face value, areas with high industry concentrations make it harder to find a job, but easier to get a raise. Conversely, areas with high occupation concentrations make it easier to find a job, but harder to get a raise. The estimates are imprecise, but these results are consistent with an earlier finding in this paper that workers find it easier to switch industries than occupations.

The period of mechanization brought changes for women, while location remains important. As shown in Table 3.15, women are still less likely than men to be employed in $t + 1$, but that effect is smaller than in the earlier period. By contrast, the period of mechanization brings a decrease in the likelihood that women find jobs outside of sugarcane. As before, the differences between states are at least 10 percentage points on every outcome, and large labor markets seem to make finding a non-cane job easier.

Consistent with experimental evidence, these results suggest that migration assistance might help workers successfully transition to new work (Bryan et al., 2014). The largest differences in outcomes is explained by state of residence, and individuals in larger labor markets. Being female and lacking education was also negatively correlated with subsequent employment and increased earnings.

Table 3.14: Transition Probabilities by Individual and Municipality Characteristics (1998–2006)

	Job _{t+1} = 1	Not-Cane Job _{t+1} = 1	Job _{t+1} = 1, Δ log Earn > 0	Not-seasonal _{t+1} = 1
Age _{i,t}	-0.0013 (0.0013)	-0.0036 (0.0016)	-0.0119 (0.0016)	0.0192 (0.0027)
Age _{i,t} ²	0.0000 (0.0000)	0.0000 (0.0000)	0.0002 (0.0000)	-0.0002 (0.0000)
Female _{i,t}	-0.0453 (0.0055)	-0.0506 (0.0062)	-0.0190 (0.0061)	0.0158 (0.0134)
Ele. Edu. _{i,t}	0.0301 (0.0050)	0.0356 (0.0053)	0.0318 (0.0067)	0.0375 (0.0171)
Ind. HHI _{m,t}	-0.0398 (0.0554)	-0.1032 (0.0820)	0.1022 (0.0629)	0.1615 (0.1666)
Occ. HHI _{m,t}	0.0191 (0.0854)	0.0540 (0.1149)	-0.2104 (0.0896)	0.1864 (0.2426)
Log Emp _{m,t}	0.0034 (0.0038)	0.0182 (0.0067)	-0.0078 (0.0041)	0.0495 (0.0121)
<i>N</i>	161,184	161,184	159,528	161,184
\bar{Y}	0.837	0.235	0.442	0.434
<i>R</i> ²	0.016	0.021	0.032	0.044

Standard errors in parentheses

Table 3.15: Transition Probabilities by Individual and Municipality Characteristics (2007–2013)

	Job _{t+1} = 1	Not-Cane Job _{t+1} = 1	Job _{t+1} = 1, Δ log Earn > 0	Not-seasonal _{t+1} = 1
Age _{i,t}	-0.0005 (0.0022)	-0.0002 (0.0022)	-0.0131 (0.0022)	0.0099 (0.0035)
Age _{i,t} ²	0.0001 (0.0000)	-0.0001 (0.0000)	0.0002 (0.0000)	-0.0000 (0.0001)
Female _{i,t}	-0.0230 (0.0061)	-0.0810 (0.0066)	-0.0101 (0.0075)	0.0545 (0.0113)
Ele. Edu. _{i,t}	0.0416 (0.0059)	0.0291 (0.0069)	0.0470 (0.0065)	0.0218 (0.0187)
Ind. HHI _{m,t}	-0.0145 (0.0750)	0.0829 (0.1159)	0.0183 (0.0809)	0.1981 (0.1676)
Occ. HHI _{m,t}	0.0160 (0.1315)	-0.1963 (0.1794)	0.0508 (0.1909)	-0.0300 (0.3442)
Log Emp _{m,t}	-0.0006 (0.0056)	0.0103 (0.0071)	0.0041 (0.0050)	0.0037 (0.0141)
<i>N</i>	104,657	104,657	102,521	104,657
\bar{Y}	0.822	0.289	0.405	0.521
<i>R</i> ²	0.011	0.021	0.010	0.049

Standard errors in parentheses

3.6 Conclusion

Theory suggests several mechanisms by which labor-saving technology ultimately benefits aggregate employment and wages. Casual empirics support this optimism. After centuries of technological innovation, unemployment is low and wages are high in most developed countries. Nonetheless, labor-saving technology is often greeted with fears about the consequences for workers. The Brazilian case highlights an important conceptual distinction: labor-saving technology may be adopted if the cost of technology falls or if the cost of labor rises. Fears of job displacement assume a fall in the cost of technology. In Brazilian sugarcane, adoption was caused by a rise in the cost of labor. Nonetheless, individuals could suffer if they have difficulties finding a new job with similar earnings.

This episode offers a several important lessons. First, workers transition between industries more frequently than between occupations. Relatedly, overall employment is likely to fall as labor saving technology is adopted. These observations recommend against retraining programs that teach workers to operate the new technology. Placing workers in other industries where they can use their existing skills seems like a more profitable approach.

Second, it provides suggestions about which groups are most vulnerable during periods of technology adoption. The mechanization of sugarcane sharpened a selection process that favored high earning, experienced workers. Workers with limited experience, or with weaker attachments to the labor force, find their difficulties amplified by the rise of technology. Women, and those with less education, tended to be less successful in finding new employment during the period of technology adoption.

Finally, local labor market conditions matter greatly, and mobility could dramatically improve outcomes. In Brazil, individuals in some states are vastly more likely to remain employed, to find a job outside of sugarcane, to earn a wage, and to find permanent work. This is consistent with a body of observational and experimental evidence that barriers to mobility harm households and aggregate growth (see, e.g., Chyn (2017); Hsieh and Moretti (2015); Bryan et al. (2014)). The mechanisms are various. In the developed world, restrictions on housing supply appear to be a large factor. In the developing world, migration may be too big a risk to take. But moving to opportunity may be one of the most important ways to mitigate technology-induced displacement.

BIBLIOGRAPHY

- IADB Approves Early Results of Sugarcane Industry's Worker Requalification Program. <http://unica.com.br/news/706475192031019628/iadb-approves-early-results-of-sugarcane-industry-por-centoE2-por-cento80-por-cento99s-worker-requalification-program/>. Accessed: 2017-08-01.
- Alvarez, J., Benguria, F., Engbom, N., and Moser, C. (2017). Firms and the Decline in Earnings Inequality in Brazil. Research Paper 17-47, Columbia Business School.
- Anand, G. (2016). Farmers Unchecked Crop Burning Fuels Indias Air Pollution. *The New York Times*.
- Andreae, M. O. et al. (1991). Biomass burning: its history, use, and distribution and its impact on environmental quality and global climate. *Global biomass burning: Atmospheric, climatic and biospheric implications*, pages 3–21.
- Arbex, M. A., Böhm, G. M., Saldiva, P. H., Conceição, G. M., III, A. C. P., and Braga, A. L. (2000). Assessment of the effects of sugar cane plantation burning on daily counts of inhalation therapy. *Journal of the Air & Waste Management Association*, 50(10):1745–1749.
- Arbex, M. A., Martins, L. C. a., de Oliveira, R. C., Pereira, L. A. A., Arbex, F. F., Cançado, J. E. D., Saldiva, P. H. N., and Braga, A. L. F. (2007). Air Pollution from Biomass Burning and Asthma Hospital Admissions in a Sugar Cane Plantation Area in Brazil. *Journal of Epidemiology and Community Health*, 61(5):395–400.
- Arbex, M. A., Saldiva, P. H. N., Pereira, L. A. A., and Braga, A. L. F. (2010). Impact of Outdoor Biomass Air Pollution on Hypertension Hospital Admissions. *Journal of Epidemiology and Community Health*, 64(7):573–579.
- Athey, S. and Imbens, G. W. (2006). Identification and Inference in Nonlinear Difference-in-Differences Models. *Econometrica*, 74(2):431–497.
- Bound, J., Jaeger, D., and Baker, R. (1995). Problems with Instrumental Variables Estimation when the Correlation Between the Instruments and the Endogenous Explanatory Variables is Weak. *Journal of the American Statistical Association*, 90:435–540.
- Bryan, G., Chowdhury, S., and Mobarak, A. M. (2014). Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh. *Econometrica*, 82(5):1671–1748.
- Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company, 1st edition.
- Calonico, S., Cattaneo, M. D., and Farrell, M. H. (Forthcoming 2016). On the Effect of Bias Estimation on Coverage Accuracy in Nonparametric Inference. *Journal of the American Statistical Association*.

- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295–2326.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2015). Optimal Data-Driven Regression Discontinuity Plots. *Journal of the American Statistical Association*, 110(512):1753–1769.
- Cançado, J. E. D., Saldiva, P. H. N., Pereira, L. A. A., Lara, L. B. L. S., Artaxo, P., Martinelli, L. A., Arbex, M. A., Zanobetti, A., and Braga, A. L. F. (2006). The Impact of Sugar Cane-Burning Emissions on the Respiratory System of Children and the Elderly. *Environmental Health Perspectives*, 114(9):725–729.
- Cattaneo, M. D., Jansson, M., and Ma, X. (2016). Simple Local Regression Distribution Estimators with an Application to Manipulation Testing. Working paper.
- Chay, K. Y. and Greenstone, M. (2003a). Air Quality, Infant Mortality, and the Clean Air Act of 1970. Working Paper 10053, National Bureau of Economic Research.
- Chay, K. Y. and Greenstone, M. (2003b). The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession. *The Quarterly Journal of Economics*, 118(3):1121–1167.
- Chyn, E. (2017). Moved to Opportunity: The Long-Run Effect of Public Housing Demolition on Labor Market Outcomes of Children. Working paper.
- Coissi, J. (2008). Usinas lideram ranking da poluição ambiental na região. *Folha de S. Paulo*.
- Conley, T. G. and Udry, C. R. (2010). Learning about a New Technology: Pineapple in Ghana. *The American Economic Review*, 100(1):pp. 35–69.
- Conniff, R. (2011). What the Luddites Really Fought Against. *Smithsonian Magazine*.
- Credendio, J. E. (2008). Usinas lideram ranking de muitas por poluição em SP. *Folha de S. Paulo*.
- Currie, J. and Neidell, M. (2005). Air Pollution and Infant Health: What Can We Learn from California’s Recent Experience? *The Quarterly Journal of Economics*, 120(3):1003–1030.
- Currie, J., Neidell, M., and Schmieder, J. F. (2009). Air pollution and infant health: Lessons from New Jersey. *Journal of Health Economics*, 28(3):688 – 703.
- Currie, J. and Walker, R. (2011). Traffic Congestion and Infant Health: Evidence from E-ZPass. *American Economic Journal: Applied Economics*, 3(1):65–90.
- David, P. A. (1966). The Mechanization of Reaping in the Ante-bellum Midwest. *Industrialization in Two Systems*, 30.
- David, P. A. (1969). *A Contribution to the Theory of Diffusion*. Memorandum: Center for Research in Economic Growth. Research Center in Economic Growth, Stanford University.
- Davis, C. A. (2017). Why Did Sugarcane Growers Suddenly Adopt Existing Technology? Working Paper.
- Davis, L. W. (2008). The Effect of Driving Restrictions on Air Quality in Mexico City. *Journal of Political Economy*, 116(1):pp. 38–81.

- den Bulte, C. V. (2000). New Product Diffusion Acceleration: Measurement and Analysis. *Marketing Science*, 19(4):366–380.
- Dube, O. and Vargas, J. F. (2013). Commodity Price Shocks and Civil Conflict: Evidence from Colombia. *The Review of Economic Studies*, 80(4):1384–1421.
- Duflo, E., Kremer, M., and Robinson, J. (2011). Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. *American Economic Review*, 101(6):2350–90.
- Duncan, B. N., Bey, I., Chin, M., Mickley, L. J., Fairlie, T. D., Martin, R. V., and Matsueda, H. (2003). Indonesian wildfires of 1997: Impact on tropospheric chemistry. *Journal of Geophysical Research: Atmospheres*, 108(D15).
- Folha de S. Paulo (2008). Usina diz que recorreu de multa ambiental.
- Foster, A. D. and Rosenzweig, M. R. (1995). Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *Journal of Political Economy*, pages 1176–1209.
- Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254 – 280.
- Gazeta Mercantil (2002). O fim da queima da palha da cana.
- Gleason, J. F., Hsu, N. C., and Torres, O. (1998). Biomass burning smoke measured using backscattered ultraviolet radiation: SCAR-B and Brazilian smoke interannual variability. *Journal of Geophysical Research: Atmospheres*, 103(D24):31969–31978.
- Gollin, D. (2010). Chapter 73 Agricultural Productivity and Economic Growth. volume 4 of *Handbook of Agricultural Economics*, pages 3825 – 3866. Elsevier.
- Greenstone, M. and Hanna, R. (2014). Environmental Regulations, Air and Water Pollution, and Infant Mortality in India. *American Economic Review*, 104(10):3038–72.
- Henrique, B. (2008). Cetesb multa produtores de cana na regio de Ribeiro. *O Estado de S. Paulo*.
- Herman, J. R., Bhartia, P. K., Torres, O., Hsu, C., Seftor, C., and Celarier, E. (1997). Global distribution of UV-absorbing aerosols from Nimbus 7/TOMS data. *Journal of Geophysical Research: Atmospheres*, 102(D14):16911–16922.
- Hirschman, A. O. (1964). The Paternity of an Index. *The American Economic Review*, 54(5):761–762.
- Hsieh, C.-T. and Moretti, E. (2015). Housing Constraints and Spatial Misallocation. Working Paper 21154, National Bureau of Economic Research.
- Hsu, A., Wong, D., and Rosengarten, C. (2016). A New Air Pollution Database Is Good, but Imperfect. *Scientific American*.
- Hsu, N. C., Herman, J. R., Bhartia, P. K., Seftor, C. J., Torres, O., Thompson, A. M., Gleason, J. F., Eck, T. F., and Holben, B. N. (1996). Detection of biomass burning smoke from toms measurements. *Geophysical Research Letters*, 23(7):745–748.
- Jayachandran, S. (2009). Air Quality and Early-Life Mortality: Evidence from Indonesia’s Wildfires. *The Journal of Human Resources*, 44(4):916–954.

- Johnston, F. H., Henderson, S. B., Chen, Y., Randerson, J. T., Marlier, M., DeFries, R. S., Kinney, P., Bowman, D. M., and Brauer, M. (2012). Estimated global mortality attributable to smoke from landscape fires. *Environmental health perspectives*, 120(5):695.
- Keynes, J. (1963). *Essays in Persuasion*. Norton Library. W. W. Norton & Company.
- Lara, L., Artaxo, P., Martinelli, L., Camargo, P., Victoria, R., and Ferraz, E. (2005). Properties of Aerosols from Sugar-cane Burning Emissions in Southeastern Brazil. *Atmospheric Environment*, 39(26):4627 – 4637.
- Leontief, W. (1952). Machines and Man. *Impact of Science on Society*, 3(4):280.
- Mansfield, E. (1961). Technical change and the rate of imitation. *Econometrica*, 29(4):741–766.
- Manuelli, R. E. and Seshadri, A. (2014). Frictionless Technology Diffusion: The Case of Tractors. *American Economic Review*, 104(4):1368–91.
- Mathiesen, K. (2016). The global air pollution 'blindspot' affecting 1 billion people. *The Guardian*.
- Mazzoli-Rocha, F., aes, C. B. M., Malm, O., Saldiva, P. H. N., Zin, W. A., and Faffe, D. S. (2008). Comparative Respiratory Toxicity of Particles Produced by Traffic and Sugar Cane Burning. *Environmental Research*, 108(1):35 – 41.
- McCrary, J. (2008). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test . *Journal of Econometrics*, 142(2):698 – 714.
- Mlhave, L., Kjrgaard, S. K., Sigsgaard, T., and Lebowitz, M. (2005). Interaction between ozone and airborne particulate matter in office air. *Indoor Air*, 15(6):383–392.
- Morris, M., Kelly, V. A., Kopicki, R. J., and Byerlee, D. (2007). *Fertilizer Use in African Agriculture*. The World Bank.
- O Globo G1 (2016). Polcia Ambiental usa imagens de satlite para identificar queimadas.
- Olmstead, A. L. and Rhode, P. W. (1993). Induced Innovation in American Agriculture: A Reconsideration. *Journal of Political Economy*, pages 100–118.
- Olmstead, A. L. and Rhode, P. W. (1995). Beyond the threshold: An analysis of the characteristics and behavior of early reaper adopters. *The Journal of Economic History*, 55:27–57.
- Osse, J. S. (2002). Colheita mecânica de cana-de-açúcar fecha 300 mil vagas. *Folha de S. Paulo*.
- Pope III, C. A. and Dockery, D. W. (2006). Health Effects of Fine Particulate Air Pollution: Lines That Connect. *Journal of the Air & Waste Management Association*, 56(6):709–742.
- Rangel, M. A. and Vogl, T. (2016). Agricultural Fires and Infant Health. Working Paper 22955, National Bureau of Economic Research.
- Samora, R. (2006). Usinas so multadas em R\$ 1,6 mi por queima irregular de cana. *Reuters*.
- Spinelli, E. (2001). Usinas mantm queimadas na regio. *Folha de S. Paulo*.

- Tamás, G., Weschler, C. J., Toftum, J., and Fanger, P. O. (2006). Influence of ozone-limonene reactions on perceived air quality. *Indoor Air*, 16(3):168–178.
- Tomazela, J. M. (2016). Polcia Ambiental adota critrio que anistia queima de cana em SP. *O Estado de S. Paulo*.
- Topel, R. (1991). Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority. *Journal of Political Economy*, 99(1):145–176.
- Torres, O., Bhartia, P. K., Herman, J. R., Ahmad, Z., and Gleason, J. (1998). Derivation of aerosol properties from satellite measurements of backscattered ultraviolet radiation: Theoretical basis. *Journal of Geophysical Research: Atmospheres*, 103(D14):17099–17110.
- Torres, O., Bhartia, P. K., Sinyuk, A., Welton, E. J., and Holben, B. (2005). Total Ozone Mapping Spectrometer measurements of aerosol absorption from space: Comparison to SAFARI 2000 ground-based observations. *Journal of Geophysical Research: Atmospheres*, 110(D10).
- Tsao, C.-C., Campbell, J. E., Mena-Carrasco, M., Spak, S. N., Carmichael, G. R., and Chen, Y. (2011). Increased estimates of air-pollution emissions from Brazilian sugar-cane ethanol. *Nature Climate Change*, 2(1):53–57.
- US Environmental Protection Agency. Particulate Matter. <http://www.epa.gov/airquality/particulatematter/>. Accessed July 7, 2014.
- Vanos, J. K., Cakmak, S., Kalkstein, L. S., and Yagouti, A. (2015). Association of weather and air pollution interactions on daily mortality in 12 canadian cities. *Air Quality, Atmosphere & Health*, 8(3):307–320.
- White, W. J. Economic History of Tractors in the United States. <https://eh.net/encyclopedia/economic-history-of-tractors-in-the-united-states/>. Accessed: 2016-06-16.
- WHO (2013). Review of Evidence on Health Aspects of Air Pollution. Technical report, World Health Organization Regional Office for Europe.
- World Bank Group (2016). Digital Dividends. World Development Report.