Rising Wages Drive Green Technology Adoption in Agriculture

C. Austin Davis

Abstract

Brazilian sugarcane growers relied on manual harvesting as late as 2007 even though mechanical harvesters had been available for many years. Manual harvesting required pre-harvest burning that generated severe air pollution. By 2013, almost all sugarcane was harvested mechanically, largely eliminating burning. State governments attributed this transition to environmental regulation, but I find that rising factor prices were the primary driver. Broad-based economic growth drove real wages up by roughly two thirds for harvest workers. A tipping-point model shows how continuously rising wages can produce abrupt, widespread mechanization. Instrumental variables estimates imply that the observed wage increase was, by itself, large enough to account for the adoption of green technology. By contrast, regression discontinuity estimates bound regulation's contribution at roughly a quarter of the change, and difference-in-differences estimates detect no effect. These results suggest that economic development can drive environmental improvements in agriculture.

Agricultural burning is practiced across the Americas, Asia, and Africa, harming ecosystems and human health[1]. In Brazil, as late as 2007, sugarcane growers burned an area the size of New Jersey every year to prepare their fields for manual harvesting[2], producing particularly dangerous aerosols[3]. Harvest-season burning in the cane districts of São Paulo is linked to hospital admissions for respiratory illness[4–7] and hypertension[8], impaired respiratory function in exposed workers[9], and worse infant health[10]. By 2013 most of the burning had stopped. This paper assesses the roles of market forces and regulation in that transition.

I find that rising wages were the primary driver of changing agricultural practices, and that the contribution of regulation was modest at most. Sugarcane can be harvested manually or mechanically. Burning substantially improves the productivity of manual harvesting but has minimal effect on mechanical harvesting. Thus, burning is standard practice with manual harvesting and rarely used with mechanical harvesting. From 1999 to 2013, broad-based economic growth in Brazil drove real wages up by roughly two thirds for harvest workers. As labor became more expensive, growers abruptly adopted mechanical harvesting, and burning largely disappeared. A tipping-point model reconciles that abruptness with the gradual rise in wages: parcels differ in the wage at which machines become cheaper, and steady wage growth can cross many thresholds in a short interval.

To test wages as a causal explanation for mechanization, I estimate the wage elasticity of labor demand in sugarcane using instruments, similar to those of Dube and Vargas [11], that capture shifts in county-level agricultural labor supply driven by conditions in the markets for other crops. The estimated elasticity is large, suggesting that the observed increase in wages is sufficient to explain the mechanization of sugarcane harvesting.

State governments claimed credit for mechanization, citing environmental regulations that gradually banned the pre-harvest burning associated with manual harvesting. Beginning in 2002, São Paulo and neighboring states passed increasingly stringent restrictions. Although enforcement intensity varied over time and across states[12], these regulations were salient, enforced, and coincided with the period of rapid mechanization[13–22].

However, two complementary analyses find no evidence that regulation drove the transition. Regression discontinuity estimates, exploiting an area threshold that exempted small growers, cannot distinguish the effect of regulation from zero, and the most negative value their confidence interval admits corresponds to about a quarter of the change in harvesting practices. A difference-in-differences approach, exploiting cross-state and over-time variation in regulatory stringency, detects no effect of regulation across 96 specifications and two heterogeneity-robust estimators, although its confidence sets are too wide to bound the contribution tightly. These findings identify an unusual channel. Environmental improvement

that accompanies rising incomes is often attributed to increased demand for environmental quality, expressed through regulation or through markets[23]. That channel may have operated here, in the form of regulation prompted by concerns about pollution, but not on a scale that could drive mechanization. Instead, the transition ran primarily through factor markets. Growers did not mechanize because anyone wanted cleaner air. Rising wages made the cleaner technology the cheaper one, and the reduction in burning followed as a byproduct.

The general lesson is not that rising wages reduce pollution. It is that relative factor prices determine which technologies are adopted, and the environmental consequence follows from which technology happens to be cleaner. Here the labor-intensive technique was the polluting one. Where the capital-intensive technique is dirtier, the same mechanism runs in reverse. Development and employment policies may therefore be effective against agricultural burning in settings where, as here, the polluting technique is the labor-intensive one.

This paper contributes to three literatures. The first concerns factor prices and mechanization. Studies of antebellum American agriculture emphasize relative prices in the adoption of machines[24–26], and labor scarcity has since been shown to induce mechanization directly, whether it arises from immigration restriction, out-migration, or population aging[27–29]. Rising farm wages likewise drove the substitution of tractors for horses[30], and relative factor prices govern whether technology develops in a clean or a dirty direction more generally[31]. These transitions usually play out over decades, which makes their aggregate consequences hard to attribute. Sugarcane mechanized in six years, using machines that had been available and unused for far longer.

The second is the literature on technology adoption in developing country agriculture, where learning and network effects are central[32–36]. Those frictions matter less here. Brazilian sugarcane growers are sophisticated and well capitalized, extension services are widely available, and harvesting machines are rented more often than owned. The question is therefore not why growers failed to learn about an available technology, but why a technology they already knew about stayed unprofitable for so long. The answer is a price rather than a friction.

The third concerns agricultural productivity and structural transformation, the process by which labor leaves a sector that grows more productive as it employs fewer people[37–39]. Most of this work asks how agricultural productivity growth releases labor to other sectors[40]. Here the causation runs the other way: growth elsewhere in the economy raised the price of labor, pulled it out of agriculture, and led growers to adopt a technology that raised productivity. The environmental consequences of that process have only recently drawn attention[41], and they matter, because prior work finds polluting practices hard

to end in poor countries, where both regulation and NGO-backed technological fixes have failed[42–44]. Here the reduction in pollution was achieved through development itself, as embodied by higher wages for some of Brazil's poorest workers.

Results

Wages and the tipping-point mechanism

As political and economic conditions stabilized in the late 1990s, a period of rapid, sustained growth took hold in Brazil. From 2002 to 2013, real wages rose substantially in all occupations and industries while hours worked increased almost everywhere but agriculture (Fig. 1a,b; Extended Data Fig. ED1). Wage growth was substantial in construction, manufacturing, and services, the sectors likely to offer alternative employment to manual harvest laborers. In sugarcane, real wages rose by roughly two thirds between 1999 and 2013, and by about 25 percent over the 2007 to 2013 transition alone. Hours worked rose through 2007 alongside a near doubling of area harvested, then fell from 2008 to 2013 while wages continued to rise, consistent with contraction in the labor supply facing growers (Fig. 1c,d).

However, the timing does not obviously support wages as a driver of mechanization. Wages rose continuously from 1999, but labor intensity held between 60 and 70 hours per hectare through 2007 before falling by 47 hours per hectare over the next six years (Fig. 2a). Area harvested green overtook area harvested burned in 2008, and burned area fell from 2.1 to 0.8 million hectares by 2013 (Fig. 2b). I develop a tipping-point model that reconciles gradual movement in relative factor prices with abrupt adoption.

In the model, growers are profit maximizers and price takers. A grower's landholdings may be diverse in their characteristics, so the grower divides them into homogenous parcels and chooses a harvesting technique parcel by parcel. The two technologies, manual and mechanical harvesting, have separate production functions, inputs, and factor prices (Methods). Growers face no fixed cost of acquiring a machine, which is usually rented rather than owned (Methods).

Abrupt adoption follows from four assumptions, the last of which can be checked against data. First, output on a parcel is the sum of the two technology-specific production functions. Second, each production function has constant or increasing returns to scale. Together these imply a threshold: for each parcel there is a manual wage, given the price of machine services, above which mechanical harvesting is cheaper. Because costs are linear, the grower harvests the whole parcel one way or the other.

Third, the threshold varies across parcels with characteristics of the land. Steeper parcels

require machines to move more slowly, raising it. Burning approximately triples the productivity of manual harvesting, so regulation that restricts burning lowers it. Fourth, those characteristics are distributed in an S shape, and the thresholds inherit it.

The fourth assumption determines the aggregate response. Where the distribution of thresholds is flat, wage increases do little because most parcels strongly favor one technique. Where it is steep, many parcels are nearly indifferent and a small wage increase induces widespread mechanization (Extended Data Fig. ED2). Thresholds are not observed, but parcel steepness and area, the characteristics that most directly govern them, have this shape (Extended Data Fig. ED3).

Causal effect of wages on mechanization

To estimate the causal effect of wages on mechanization, I construct instrumental variables from exogenous fluctuations in the markets for other crops, which shift the supply of labor facing sugarcane growers. The study region's other major crops are maize, soybeans, coffee, and oranges, each employing between 0.5 and 2.5 million workers. For each, I interact the output of the top three producers other than Brazil with historical county-level acreage of that crop.

The first component provides time-varying shifts in labor supply; the second scales those shifts by local exposure to the crop. For instance, unfavorable rains in Indonesia lower coffee output and raise the international price. Brazilian coffee farmers hire more labor, contracting the supply facing sugarcane growers. The size of that contraction depends on how much coffee is grown locally: a county with extensive coffee cultivation experiences a large shift, while a county that grows no coffee experiences none. The design is conceptually similar to Dube and Vargas [11]. Methods gives full details, including lag structure and identifying assumptions.

Table 1 presents the second-stage estimates. Columns (1)–(4) each use instruments from a single crop, relying on different sources of cross-sectional and time series variation. In three of the four crops the estimated elasticities are negative, significant, and similar in magnitude: -2.94 (s.e. 0.99) for maize, -2.20 (0.55) for coffee, and -3.20 (0.73) for oranges. The soybean estimate is imprecise and positive, 1.42 (1.48). Column (5) combines the instruments from all four crops and yields an elasticity of -2.73 (0.42). Column (6) adds a squared wage term. It makes similar predictions over the observed range of wages, but estimates the curvature implied by the model (Extended Data Fig. ED4).

First-stage strength varies across these columns and bears on how the estimates should be read. The maize and coffee first stages are strong ($F = 19.9$ and 15.1) and soybeans

is adequate ($F = 13.9$), but oranges is weak ($F = 5.1$) and the pooled specification in Column (5) has a first-stage F of 9.5 across twelve excluded instruments. Conventional confidence intervals are undersized when instruments are weak, so the interval on Column (5) should be read as optimistic[45, 46]. The substantive basis for the result is therefore not the precision of the pooled column but the agreement of the two strongest and non-overlapping sources of variation: the maize and coffee elasticities lie within about one standard error of each other and of the pooled estimate.

I use the estimated elasticity to predict the change in labor intensity attributable to wage changes alone. Substituting the estimates from Column (5), the observed change in log area harvested, and the observed change in log median real wages:

$$\widehat{\Delta \log L^s} = 0.84 \times 0.7 + (-2.7) \times 0.52 = -0.82 \text{ with } 95\% \text{ CI } [-1.20; -0.44].$$

The observed change in the log of aggregate quantity of labor was -0.49. While the confidence interval contains -0.49, the point estimate is meaningfully larger. Thus, wage increases were sufficient to explain the transition, and the over-prediction suggests that some force held labor use above what the wage and area changes alone imply. The rental price of harvesting services may have increased as demand grew faster than the stock of machines and trained operators, which raises the threshold wage on every parcel and slows adoption. Sugarcane area also nearly doubled. If newly planted area requires more labor than established stands, an expansion would increase labor use for a time. Both act on labor demand, and because the elasticity is identified from shifts in labor supply, neither biases the estimate unless correlated with the instruments.

Bounding the contribution of regulation

Sugarcane fields are burned in preparation for manual harvesting but not mechanical harvesting. São Paulo passed a gradual ban of pre-harvest burning in 2002[47], and the five neighboring states restricted burning later, through instruments that differ in form and legal standing (Methods). The regulations carried meaningful incentives: violations could be punished by large fines, and failure to comply jeopardized mandatory state environmental licenses. As described above, these regulations were salient, enforced, and coincided with the period of rapid mechanization for which regulators took credit[48, 49].

I evaluate these claims using two complementary identification strategies, and I state at the outset what they can and cannot establish. Neither design rejects the hypothesis that regulation had no effect, and neither is precise enough to rule out an economically meaningful one. The regression discontinuity supports an upper bound: effects large enough to account

for the bulk of the transition lie outside the range its data admit. The difference-in-differences design finds no detectable effect but is too imprecise to bound one.

First, the São Paulo regulation was much less stringent for growers with less than 150 hectares, providing a regression discontinuity design. Using the 2006 Census of Agriculture, I compare harvesting techniques between unregulated establishments just below the area threshold and regulated establishments just above. Growers could in principle understate area to avoid regulation, but reported area shows no evidence of such manipulation (Supplementary Information). Establishments above the threshold are 4.1 percentage points less likely to harvest manually only, with a 95% confidence interval of $[-13.9, 4.7]$ percentage points (Fig. 3). The estimate is not distinguishable from zero. More than 60 percent of regulated establishments still reported manual harvesting only, although the law permitted none of them to do so, and mechanical harvesting covered about 84 percent of São Paulo sugarcane area by 2013 (Fig. 2b). Even the most negative value admitted by the confidence interval therefore corresponds to about a quarter of the observed decline in exclusive manual harvesting (Methods). Regulated growers also show no changes in the complementary inputs that mechanization would require, including fuel expenditure, machine counts, and contract services (Supplementary Information).

This design bounds the effect of regulation as it stood in São Paulo in 2006, and it cannot be repeated for another census year or another state (Methods). The outcome is exclusive manual harvesting, which a positive burning restriction makes impracticable; a complying establishment records a zero in 2006 and under the stricter requirements that followed. But enforcement may have intensified alongside stringency, and the difference-in-differences design that follows covers those years directly.

Second, I exploit cross-state and over-time variation in regulatory stringency in a difference-in-differences framework covering 1999 to 2013, a window that spans the transition. Using administrative data covering the universe of formal-sector employment (RAIS), I measure municipality-level labor intensity as hours worked by manual laborers in sugarcane per hectare harvested. I estimate how changes in regulatory stringency affect log labor intensity, controlling for municipality and year fixed effects. Because stringency varies across only six states, inference is by randomization over the states' stringency paths (Methods). The estimated effect is small and not distinguishable from zero: at 2013 stringency it accounts for 0.72 hours per hectare, 1.2 percent of the observed decline of 60 hours per hectare ($p = 0.87$; Table 2). None of 96 alternative specifications rejects the null of no effect (Fig. 4).

The headline confidence set rejects contributions larger than about half of the observed decline, but most alternative specifications admit the entire decline, so this design does not bound regulation's contribution tightly (Methods). Because stringency is continuous and

adopted at different times across states, two-way fixed effects can misweight heterogeneous effects[50–52]; heterogeneity-robust estimators also yield small point estimates that differ in sign (Extended Data Table ED1; Methods).

Compared to the regression discontinuity, the difference-in-differences analysis estimates a conceptually different parameter from separate data. The two designs are complementary in their weaknesses: the first is cleanly identified but confined to 2006 and to establishments near the threshold, while the second covers the full transition and is sensitive to the behavior of large growers but is imprecise. Neither analysis supports the claim that regulation drove the transition. The regression discontinuity rules out a contribution larger than about one quarter of the change in the share of establishments engaged in exclusively manual harvesting. The difference-in-differences design finds no detectable effect on municipality-level labor intensity, with small point estimates, but its confidence sets are too wide to provide a comparable bound.

Discussion

Agricultural burning carries significant environmental and health consequences. It persists in Brazil, Southeast Asia, and sub-Saharan Africa. Understanding the rapid reductions in Brazilian sugarcane burning matters beyond one crop in one country. The answer lies in factor prices rather than in environmental policy. Exploiting variation in agricultural labor supply to identify the wage elasticity of labor demand, I find that economic development, in the form of higher wages for some of the poorest workers in Brazil, was sufficient by itself to account for the transition.

The limited estimated effect of regulation deserves careful interpretation. The results establish an upper bound, not the absence of an effect. One design bounds the effect well below what the transition would require, and a second, with different data and a different source of variation, detects no effect. Still, regulation may well have complemented market forces, by coordinating expectations or by accelerating adoption on the margin. Indeed, the model suggests it is a more potent lever the closer producers sit to their adoption thresholds.

The most important alternative to the wage attribution is a fall in the rental price of harvesters. The tipping-point mechanism runs through the gap between the two factor prices, so cheaper machinery lowers threshold wages on every parcel and would generate the same abrupt adoption through the same S-shaped distribution. Brazil expanded subsidized agricultural machinery credit substantially over this period, which may have lowered costs for purchasing or renting[53]. The wage elasticity is identified from shifts in agricultural labor supply, and the sufficiency calculation asks only whether the observed change in wages could

by itself have produced the observed adoption, whether or not machine costs also fell. But the specification absorbs no national trend in machine costs, and I am not aware of a series that could be included as a control or used to apportion the estimated effect. Distinguishing the wage channel from the machinery-cost channel is a valuable extension of this work.

These results have implications for sustainability policy in two domains. First, the lesson is not that rising wages reduce pollution. It is that relative factor prices determine which technology is adopted, and the environmental consequence follows from which technology happens to be cleaner. Here the labor-intensive technique was the polluting one, so development and employment policies may be effective tools against agricultural burning. Where the capital-intensive technique is the dirtier one, the same mechanism runs in reverse. Governments seeking to address polluting practices need not rely solely on environmental regulation, but neither can they count on growth to deliver environmental gains. Acting on that recommendation requires knowing what raised wages, which this paper does not establish. The growth was broad based, which suggests large increases in demand for low-skilled labor, but its sources remain open.

Second, the findings challenge the assumption that environmental protection comes at the cost of economic growth, and the converse assumption that growth comes at the cost of the environment. As lawmakers contemplated regulation, economists predicted that mechanization would dramatically depress employment[14]. 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. Major development markers, wages for the poor and agricultural productivity, improved alongside environmental outcomes.

The tipping-point mechanism identified in this paper may be relevant beyond sugarcane. The key insight is that a continuous, predictable price change can generate discontinuous outcomes at the aggregate level, even in the absence of information frictions, fixed costs, or improvements to the technology itself. This mechanism may help explain other episodes of rapid adoption where technologies were available long before they were widely used. Although Brazilian sugarcane growers generally do not face information or credit constraints common in smallholder agriculture, such constraints could reduce the responsiveness of producers to factor price changes. The same logic applies to research policy. As with wage changes, directed research could shift relative factor prices in favor of the cleaner technology. But the model implies that returns to directed research may be uneven: large when the prevailing wage sits in the dense part of the threshold distribution and near zero elsewhere.

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.

Methods

Study region

Although sugarcane is grown in two separate regions of Brazil, I focus on a group of six contiguous states in South-Central Brazil: Goiás, Minas Gerais, Paraná, Mato Grosso do Sul, Rio de Janeiro, and São Paulo. These states accounted for about 80 percent of Brazil's sugarcane tonnage between 1990 and 2010[54], and they experienced the changes in harvesting practices that motivate this paper[55] (Fig. 2). By contrast, the smaller growing region in the Northeast has long used a more rudimentary form of mechanized harvesting. The study region is home to about half of Brazil's population[56] 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[54, 57]. Extended Data Fig. ED5 is a map of Brazil indicating the study region.

Sugarcane cultivation and harvesting

Except where another source is named, the descriptive figures in this subsection come from industry interviews (see Data).

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. The capital-to-labor ratio is higher for mechanical harvesting; growers and machine manufacturers report that harvesting the same area requires 70 to 90 percent less labor using the mechanical technology.

The productivity of manual harvesting is greatly enhanced by pre-harvest burning. In a burned field, workers can clear 8 to 10 tons of sugarcane per day. In an unburned field, slowed by heavy dense vegetation, workers clear around 3 tons per day. Burning also raises the rate of mechanical harvesting, but by much less. A field trial in Paraná records an effective field capacity of 1.02 hectares per hour on burnt cane against 0.71 on green, a gain of about 44 percent, against the roughly threefold gain burning gives manual cutting[58]. Green harvesting also produces smaller losses and less damage to the cut cane. Mechanical harvesting is therefore predominantly practiced without burning. Restrictions on burning

effectively decrease total factor productivity of the manual technology relative to the mechanical technology.

The sharing of harvesting machines is extensive. Contract harvest services, machine rental, and land rental are all common. 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 each harvesting technique depends on characteristics of the land. The size, shape, and layout of a parcel affect the rate of mechanical harvesting: machines cut as little as 450 tons per day on poorly prepared fields and as much as 1,000 tons per day under ideal conditions, against typical performance of 500 to 650. Manual harvesting productivity is much less sensitive to these characteristics.

Background on regulation

Following constituent complaints about air pollution and health concerns, the state of São Paulo passed a gradual ban of pre-harvest burning in 2002[47]. Under the regulations, property owners were permitted to burn only a fraction of each property, as recorded in the local property registry, with scheduled reductions: 80 percent in 2002, 70 percent in 2006, 50 percent in 2011, 20 percent in 2016, 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[59]. The 2007 revision was a voluntary agreement between growers and the environmental regulator rather than a change in the statute. The regulator held leverage through environmental licensing, other regulations, and legislators' threats of stricter law, and the two sides agreed to more aggressive restrictions while avoiding an uncertain and costly legislative process. The other five states restricted burning later and through instruments that differ in form and legal standing. Goiás legislated in 2006 and Mato Grosso do Sul in 2007, Rio de Janeiro in 2011, while Minas Gerais and Paraná acted by administrative norm rather than statute, in 2009 and 2010[60–65]. Their schedules differ from São Paulo's and from each other's, and Methods sets out how each is measured.

The regulations included meaningful incentives. Violations of the São Paulo regulation could be punished by large fines, levied for each hectare burned in excess of the allowed fraction; in 2002, the fine amounted to 13 percent of average per-hectare revenue, and it was revised upward each year, roughly tracking inflation. Failure to comply jeopardized mandatory state environmental licenses. The São Paulo government partnered with the Brazilian space agency INPE to monitor harvesting practices via satellite imagery[55, 66]. National media publicized both the regulations and enforcement efforts, and fines were levied

on growers of all sizes[13, 16–21], though enforcement intensity varied over time and across states[12]. Independent analysis of satellite fire detections finds a sustained decline in fire occurrence over the regulated period[22].

Crucially, the law's schedule applied only to what it termed mechanizable areas, defined as plantings on terrain above 150 hectares with a grade of 12 percent or less, and it classified the areas of each rural property separately[47]. Land below the area threshold was not required to be mechanized before 2010, and steeper land was exempt on the same terms. I exploit the area threshold in the regression discontinuity design.

Data

Administrative labor data (RAIS)

The primary labor data come from confidential administrative micro data maintained by the Brazilian Ministry of Labor. Known by its Portuguese-language acronym RAIS (Relação Anual de Informações Sociais), these data are compiled annually and comprise the universe of formal-sector employment from 1998 to 2014. Each record corresponds to an employment spell, providing information about the employee and the employer, including the municipality of work, wages, duration of employment, hours worked, and detailed industry and job classifications. In interviews, farmers and farm workers indicate that labor in the sugarcane sector is predominantly formal-sector and unionized, and those reports find support in the data. To gauge how much of that labor RAIS observes, I divide the national count of employment spells in RAIS by the national count of individuals in PNAD, the household survey, for sugarcane and coffee together. The counts are national because RAIS records place of work and PNAD records place of residence, and some sugarcane workers are seasonal migrants, so one individual may appear in different places in the two sources. Consistent with interview reports, that ratio is roughly 60 to 75 percent.

The ratio is not a rate of coverage, and may overstate one. Spells and individuals are different objects, a seasonal worker generates several spells, and the PNAD count is drawn from a single reference week. Because the labor quantities and the wages used throughout are formal-sector measures, any change over the sample period in the share of sugarcane labor that is formal would appear in them. Whether that share moved, and in which direction, cannot be established from these data.

Census of Agriculture

The 2006 Census of Agriculture records a variety of information about every agricultural establishment in Brazil, including location, size, crops grown, and harvesting technique.

The Census Bureau endeavors to survey every agricultural establishment in the country. The unit of observation corresponds to the decision-making unit. Given the level of detail and disaggregation, these data are confidential.

Agricultural survey and international data

The PAM survey (Produção Agrícola Municipal), conducted by the Brazilian Census Bureau, contains annual, municipality-level area harvested for all crops. Household survey data from the PNAD (Pesquisa Nacional por Amostra de Domicílios) provide information on wages and hours worked. International crop output data are drawn from the UN FAO's FAOSTAT database.

Industry interviews

Several descriptive features of sugarcane harvesting reported above come from eleven semi-structured interviews I conducted in São Paulo state between 26 September and 28 October 2014. Respondents included mill managers and agricultural staff, independent growers ranging from small to large, harvester manufacturers and dealers, researchers at a sugarcane technology center and at universities, a producers' association, union representatives and harvest workers, and staff at the state environmental agency and the agricultural price agency. The interviews establish how harvesting is organized and what governs its cost. They inform the assumptions of the model and the descriptive account of the industry; no estimate reported in this paper is derived from them.

Measuring regulatory stringency

Pct is the fraction of existing cane area on which pre-harvest burning is prohibited in a given state and year. Restricting the measure to existing plantings is deliberate. Several states imposed immediate requirements on newly planted area or on newly licensed mills, alongside the gradual schedules that governed the standing crop. Measuring those provisions would require the share of each municipality's area that is newly planted, which is unsatisfactory on two counts: that share is constructed from the same area series that forms the denominator of labor intensity, and the decision to expand is itself a grower response the requirement may influence. The consequence is that Pct understates stringency in São Paulo, Goiás and Minas Gerais, in the same direction in each.

Three conventions turn a published schedule into a number. Where an instrument sets separate timetables for mechanizable and non-mechanizable area, Pct follows the mechanizable schedule, and no average is taken across the two. This is the schedule the São Paulo

target attaches to, and following it in the other five states keeps them on the same footing. A deadline counts in the year it names, so a requirement of 20 percent by 2012 enters as 0.20 in 2012. Values carry forward between milestones and are capped at one.

In São Paulo, Lei 11.241/2002 requires burning to be eliminated on 20 percent of mechanizable area in 2002, 30 percent in 2006, 50 percent in 2011, 80 percent in 2016 and 100 percent in 2021[47]. In June 2007 the state and the mill association signed the Protocolo Agroambiental, a voluntary agreement advancing complete elimination to 2014 and requiring green harvest on 70 percent of mechanizable area from 2010[59]. I use the statutory schedule through 2009 and the Protocolo thereafter, since the Protocolo set no intermediate target before 2010. Adherence was voluntary but carried licensing benefits: Resolução SMA 88/2008 gives mills in conformity up to fifteen years to implement their environmental adequacy plans and extends license renewal periods by one third[67]. Signatories accounted for between 54 and 77 percent of state cane area over the sample period. The 70 percent target applied to mills; the parallel commitment signed with the cane suppliers' association in 2008 was 60 percent above 150 hectares and 20 percent below, and Pct uses the mills' figure.

In Goiás, Lei 15.834/2006 set a schedule beginning at 10 percent of harvested area in 2008, but the state court declared that schedule unconstitutional, striking Article 1 and each of its five subsections[60, 68]. The judgment was handed down in April 2009, so the first milestone year fell while the challenge was pending. Nothing was administered under the schedule in the meantime: the implementing decree Article 9 required within sixty days was never issued, and the state's own account of its environmental actions for 2008 describes no percentage requirement[69]. I therefore record no requirement on existing plantings in any year. Goiás did restrict burning through licensing, requiring environmental licenses for newly installed mills to prohibit burning of all the cane they process, but that reaches new capacity rather than the established sector[70].

Minas Gerais has no statute restricting pre-harvest burning. Deliberação Normativa COPAM 133/2009 prohibits burning outright on land newly planted to cane below 12 percent grade and limits operations established from 2008 to burning 20 percent of newly planted area, but for operations established earlier it sets only an endpoint, requiring burning to cease by 2014 on schedules agreed individually with the environmental agency[64]. Existing plantings therefore carry no measurable requirement before 2014. Rotational replanting falls under the regime for established operations rather than counting as new planting.

In Mato Grosso do Sul, Lei 3.357/2007, as amended by Lei 3.404/2007, requires burning to be eliminated on mechanizable area within six years beginning in 2010, at a rate of at least 16.75 percent a year[61, 62]. I count the six increments inclusively, so the first falls in 2010 and the last in 2015. Two features of the drafting support this. Six increments of 16.75

percent sum to 100.5, so the rate and the period state one six-step schedule twice, which requires the stated starting year to carry the first increment. And the provision this one replaced required elimination over twenty years from 2006 at 5 percent a year, where the rate reaches exactly 100 percent only if 2006 is itself counted; the same construction appears, with no period to interpret, in a paragraph of the earlier law that was never revoked and whose reductions begin in 2010. A reader taking Article 3 in isolation might instead place the first increment in 2011, an alternative I note because it moves six state-years.

In Rio de Janeiro, Lei 5.990/2011 requires elimination on 20 percent of mechanizable area by 2012, 50 percent by 2014, 80 percent by 2018 and 100 percent by 2020, exempting plantings of 100 hectares or less[63].

In Paraná, Resolução SEMA 76/2010, re-enacted without change as Decreto 10.068/2014, requires elimination on 20 percent of mechanizable area by the end of 2015, 60 percent by 2020 and 100 percent by 2025[65, 71]. No requirement binds before the end of the estimation window.

Two further features of the regulations are not reflected in Pct. Each instrument applies its schedule only to areas it classifies as mechanizable, and the definitions differ: São Paulo, Goiás and Paraná require more than 150 hectares together with a grade at or below 12 percent, while Mato Grosso do Sul, Minas Gerais and Rio de Janeiro use the grade criterion alone. And the measure does not capture restrictions imposed below the state level, of which the most consequential is a federal court order suspending burning authorizations across 18 municipalities in Mato Grosso do Sul from 2010. Throughout, Pct records what each instrument required, not what was enforced.

Constructing labor intensity

I combine RAIS and PAM to construct a continuous, size-weighted measure of mechanization at the municipality-year level. Labor intensity is defined as hours of labor supplied by manual laborers in the sugarcane industry divided by total area devoted to sugarcane cultivation. Hours are calculated from employment spell length and contracted hours per week. Using detailed occupation and industry codes, I restrict the analysis to manual laborers in the sugarcane industry.

The panel is unbalanced. Municipalities enter as sugarcane cultivation spreads into new areas, from roughly 400 in the early years to almost 800 by the end of the sample. Entering municipalities are new to the crop and need not resemble established ones in labor intensity or in the pace at which they mechanize, so the composition of the sample changes over the same window in which the transition occurs. The difference-in-differences specifications include

municipality fixed effects, which absorb level differences between entering and established municipalities but not differences in trend, and the instrumental variables specification carries no municipality fixed effects at all. The specification curve for the difference-in-differences design varies the sample among nested alternatives, and no alternative rejects the null of no effect (see Difference-in-differences design).

For the difference-in-differences estimates I restrict the sample to municipalities that harvested sugarcane in at least one year between 1999 and 2001, before any state was treated. An alternative would be to require presence throughout the panel, but whether a municipality is still growing cane in a later year is itself an outcome that stringency could affect, so a balanced panel would condition the sample on the dependent variable. The restriction removes 552 of 8,262 municipality-years and retains 89 percent of the sample's harvested cane area in 2013.

Regression discontinuity design

I estimate the effect of the São Paulo regulation on harvesting practices by comparing establishments above and below the 150-hectare regulatory threshold. The primary outcome is an indicator for exclusive use of manual harvesting (value of one if manual only, zero if any mechanical harvesting). The assignment variable is reported establishment area.

Using the method of Calonico et al. [72], I estimate the treatment effect as

$$\tau = \lim_{x \to 150^+} \Big(\mathbb{E}[\, Y_i \mid X_i = x \,] \Big) - \lim_{x \to 150^-} \Big(\mathbb{E}[\, Y_i \mid X_i = x \,] \Big).$$

Separate local linear regressions are estimated above and below the threshold. Bandwidths for the point estimate are selected using the MSE-optimal method of Calonico et al. [72]; bandwidths for the bias correction are selected using the CER-optimal method of Calonico et al. [73]. I use the triangular kernel, estimate local linear regressions ($p = 1$), and estimate variance using a nearest-neighbor approach clustered by municipality. I use separate bandwidths on either side of the cutoff because the density of the assignment variable is decreasing in this region. I report 95% confidence intervals recentered to account for misspecification bias.

This design recovers the effect of the regulatory regime in force in São Paulo as of the 2006 Census of Agriculture. At that time, establishments in other states did not yet face binding regulation. Although it would be informative to repeat the analysis in prior or subsequent Censuses, no regulation had been enacted as of the preceding Census, and the area exemption had expired before the next.

Compliance looks like a zero above the threshold. The regulation prohibits burning on a share of the property, and hand cutting an unburned field recovers only about a third of the burned-field rate, which roughly triples the labor cost per ton. Complying without machinery is therefore prohibitively expensive, and in practice an establishment complies by mechanizing at least part of its area. That remains true under subsequent increases in stringency. However, enforcement may have intensified alongside increases in stringency, and the 2006 cross section cannot exclude that possibility. Using the most negative value the confidence interval admits is a hedge against it.

Deriving the upper bound from the regression discontinuity

The estimated discontinuity in exclusive manual harvesting is $\hat{\tau} = -0.041$ with a 95% confidence interval of $[-0.139, 0.047]$. To convert this into a bound on regulation's contribution, I take the most negative value the interval admits, 0.139, as a hypothetical upper limit on the share of regulated establishments moved away from exclusive manual harvesting. The benchmark against which this is measured is the change that actually occurred: more than 60 percent of regulated establishments reported exclusive manual harvesting in 2006, and I take the benchmark to be roughly 60 percentage points on the assumption that few were still doing so by 2013. The ratio 0.139/0.60 is 23 percent: regulation can account for at most about one quarter of the observed decline in the share of establishments harvesting manually only.

The same comparison can be read the other way. No establishment could satisfy the requirement while harvesting exclusively by hand, so a regulation that bound and was enforced would have driven exclusive manual harvesting close to zero above the threshold, a discontinuity of about 60 percentage points. The estimate is -0.041, and even the most negative value the interval admits is less than a quarter of that. The estimates are imprecise, but they are not consistent with a regulation that bound and was enforced at the threshold in 2006.

This is an extensive-margin calculation throughout. The outcome registers an establishment that mechanizes 1 percent of its area identically to one that mechanizes all of it, so it is silent on whether regulation raised the share of area mechanized among establishments that already used some machinery. That margin is addressed indirectly through the continuously measured inputs reported below. Within the extensive margin, assigning regulation the most negative value the confidence interval admits credits it with the largest effect these data will tolerate.

The 2006 share is measured directly in the Census. The endpoint is not, and had to be inferred from aggregate data. Satellite monitoring puts the mechanically harvested share of

São Paulo sugarcane area at about 84 percent in 2013, against under half in 2007 (Fig. 2b). Exclusive manual harvesting falls faster than this, because an establishment leaves the category as soon as it mechanizes any part of its area, and partial mechanizers continue to burn the remainder. How much faster cannot be determined from the available data. If 10 percent of regulated establishments were still harvesting manually only in 2013, the observed change is 50 rather than 60 percentage points and the upper bound rises from 23 to 28 percent of that change. If 20 percent of regulated establishments were still doing so, the observed change is 40 percentage points and the bound rises to 35 percent.

The density of reported area is discontinuous at 150 hectares, but this reflects heaping at round numbers throughout the distribution rather than sorting around the threshold. The degree of rounding at 150 is not unusual relative to other round values, and I find no evidence of strategic manipulation of the assignment variable (see Supplementary Information). While the Census does not record a continuous measure of mechanization, several continuously measured inputs serve as secondary outcomes in the Supplementary Information. I consider five machine-related outcomes: expenditure on contracting services, fuel expenditure, the number of harvesting machines, machine rental expenditure, and 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. A shift of area from manual to mechanical harvesting on the scale the regulation required would move these inputs. Of the eight, one is significant at the 5 percent level, the total number of workers, which is higher among regulated establishments; fuel expenditure and the count of harvesting machines are lower. All three run opposite to the direction mechanization implies, although the estimates are imprecise and their confidence intervals cannot rule out large effects.

Interpretation and limitations of the regression discontinuity design

Four limitations bear on how these estimates should be read. First, the design 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, while the benchmark against which it is measured, the change in exclusive manual harvesting across all regulated establishments, is a population-wide quantity. The bound therefore requires that the effect at the threshold be representative of the effect elsewhere in the size distribution.

Second, and bearing directly on that requirement, these estimates could miss an effect of the regulation under certain circumstances. 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. The difference-in-differences design that follows speaks to this concern, because it uses a continuous, size-weighted measure of mechanization and an identification strategy that does not rely on a comparison of establishments near the threshold.

Third, assignment at the threshold is not sharp. The law exempts plantings of 150 hectares or less, but it also exempts larger plantings on land with a grade above 12 percent, so an establishment just above the area threshold is regulated only if its land is also flat enough. Crossing the threshold therefore raises the probability of regulation by less than one, and what is reported here is an intention-to-treat effect that understates the effect on the establishments actually subject to the regulation. This would widen the bound rather than narrow it. Land qualifying for the grade exemption at this size is rare, however. A square of 150 hectares is about 1.2 kilometers on a side, so an average grade of 12 percent across it implies roughly 150 meters of elevation change from one edge to the other, a sustained slope steeper than any on the interstate highway system, where 6 to 7 percent is the practical maximum. Because the criterion is an average over the area rather than a maximum, meeting it requires a hillside on that scale and not a field with a steep corner. The frequency cannot be measured directly in these data: grade is observed only for satellite-delineated parcels, which are smaller than establishments, do not follow legal property boundaries, and are delineated in part by harvesting technique, so their distribution of steepness does not map onto the Census establishments used here.

Fourth, the direction of any remaining manipulation is favorable to the bound. 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 least likely to comply, it would lead to an overestimate of the effect of regulation, so the estimates reported here would if anything be an upper bound.

Difference-in-differences design

Cross-sectional and time-series variation in regulatory stringency enable a difference-in-differences evaluation. The required fraction of area harvested without burning varies across states and over time (Extended Data Fig. ED6). I estimate

$$\ln(\, L/T \,)_{j,s,t} = \mu_j + \gamma_t + \beta \, \mathrm{Pct}_{s,t} + \varepsilon_{j,s,t}, \tag{1}$$

where j indexes municipality, s indexes state, t indexes year, L/T is labor intensity, μ and γ are municipality and year fixed effects, and Pct is the fraction of existing cane area on which pre-harvest burning is prohibited, constructed as described in Data. The coefficient of interest β captures how changing regulatory stringency affects labor intensity, in proportion to its level. The counterfactual of interest throughout is a world without regulation, in which stringency is zero in every state and year. The headline specification is estimated without weights on municipalities that harvested sugarcane in at least one year between 1999 and 2001, 7,710 municipality-years in 844 municipalities. Each choice in it is one level of an axis in the specification curve described below, and the headline was named before that curve was estimated.

Fixed effects are for municipalities rather than states because the panel is unbalanced. Municipalities entering and leaving the sample shift state-year means of labor intensity, and because Pct varies at the state-year level, those compositional changes could confound its estimated effect under state fixed effects. Municipality fixed effects absorb differences in level between entering and established municipalities, though not differences in trend.

Inference

Stringency varies at the state level and the study region contains six states, so neither available clustering supports conventional inference, and the two fail for different reasons. Clustering on municipality assumes too much independence: stringency is constant within a state and year, so treating municipalities within a state-year as independent inflates the effective sample size and the standard errors are too small. Clustering on state has the right independence structure but leaves six clusters, and cluster-robust inference is consistent only as the number of clusters grows.

I therefore report no conventional standard error and use randomization inference throughout. The test permutes the six states' realized stringency paths among the six states. Goiás, Minas Gerais and Paraná share an identically zero path within the estimation window, so the 720 permutations collapse to 120 distinct assignments. In each of the 120 assignments, I calculate a test statistic as $\hat{\beta}/\widehat{\mathrm{SE}}_{\hat{\beta}}$, where $\widehat{\mathrm{SE}}_{\hat{\beta}}$ is the municipality-clustered standard error. That standard error is not directly used for inference; it serves only to scale estimates with different precision. The permutation reassigns stringency paths only, so each specification is estimated on the same municipality-years under every assignment. Under the null that stringency had no effect, the value of the test statistic is exchangeable across the 120 assignments. If the test statistic from the observed assignment ranks sixth or more extreme, the null is rejected in an exact test at the 5% level. I report the rank alongside the p-value because 120 assignments put a floor of 0.008 under any p-value the design can produce.

Outcome and functional form

The object is labor intensity rather than hours because cultivated sugarcane area changes meaningfully over the sample period. The functional form matters: parallel trends can hold in levels or in logs, not both. For the headline specification, I use logs because the Winsorized level fails the pre-treatment trend test described below and the log does not, while the unwinsorized level is too dominated by extreme values for that test to be informative. The mechanism also points to the log, since a restriction that shifts a share of area from manual to mechanical harvest changes labor use in proportion to baseline use rather than by a fixed number of hours per hectare.

Pre-treatment trends

I test for differential trends before any state is treated by regressing the outcome on municipality and year fixed effects and the interaction of a linear time trend with an indicator for states whose stringency ever changes. Two windows are available. The period 1999 to 2001 covers only three years but all six states are untreated. The longer window runs from 1999 to 2009 and excludes São Paulo, which is treated from 2002; Mato Grosso do Sul is first treated in 2010 and Rio de Janeiro in 2012, so both are untreated alongside the three never-treated states throughout. Inference uses the randomization procedure described above, applied to the states in each window with the sample held fixed. In the longer window it permutes which two of the five states carry a nonzero path, giving 10 assignments and a floor of 0.10 under any p-value, so I report the rank of the observed assignment.

Pre-treatment trends look more parallel in logs than in the Winsorized level. On the headline sample (municipalities that harvested sugarcane in at least one year between 1999 and 2001), the Winsorized level ranks first of 10 in the longer window, the most extreme value any assignment produces, and the log ranks eighth. The same ordering holds on every sample in the specification curve and when the indicator is replaced by each state's terminal stringency. In the short window, which includes São Paulo, the log ranks 20th of 20.

The test has clear limits. It detects linear departures only, and the long window cannot reach São Paulo, the state that supplies almost all of the treated observations.

Constructing the confidence set

The randomization test yields a p-value for the null of no effect. To show the range of effects consistent with the data, and to compare the difference-in-differences estimates with the bound from the regression discontinuity, I invert the test to obtain a 95% confidence set, i.e. the set of treatment effects b_0 that the test does *not* reject.

I start with a hypothesized treatment effect b_0 and subtract it from the observed outcome, $\tilde{y} = y - b_0 \times \text{Pct}$. If the hypothesis is correct, $\tilde{y}$ is log labor intensity in a world without regulation and should not depend on which state followed which stringency path. I then re-estimate the specification with $\tilde{y}$ as the outcome under each of the 120 assignments and rank the test statistic from the observed assignment against the others. If it ranks sixth or more extreme, the data reject b_0 at the 5% level. If it does not, the data are consistent with b_0. For $\tilde{y}$ to be defined, the hypothesis has to specify the effect on every observation, which is why b_0 is a response proportional to stringency rather than an average effect.

I repeat this over a range of hypothesized effects, and the 95% confidence set is the collection of values of b_0 that the test does not reject. Candidates are evaluated from -30 to 30 in steps of 0.05, and each endpoint of the accepted set is then located in steps of 0.005; no accepted set reaches either end of the range.

Because $\tilde{y}$ subtracts the hypothesized effect using the stringency paths states actually followed, changing b_0 changes the outcome in every permuted regression, not only in the observed one. The ranking is therefore redone from scratch at each candidate, so the set need not be centered on the point estimate or form a single interval. In the headline specification it is a single interval.

The hypothesized effect also holds the response constant in proportion across municipalities. If responses are heterogeneous, no single value is true and the set covers a constant-effect parameter rather than an average. Coverage rests on the same as-if-random assignment of stringency paths that the p-value already rests on, so hypothesizing an effect in order to invert the test adds no assumption beyond the one already made.

Specifications with a lag and a lead

Specifications with a lag and a lead of stringency estimate

$$\ln(\ L/T\ )_{j,s,t} = \mu_j + \gamma_t + \beta\,\text{Pct}_{s,t} + \beta_L\,\text{Pct}_{s,t-1} + \beta_F\,\text{Pct}_{s,t+1} + \varepsilon_{j,s,t}. \tag{2}$$

The estimate of interest is the sum $\hat{\beta}+\hat{\beta}_L+\hat{\beta}_F$. Stringency affects labor intensity in the year it applies, in the year after, and in anticipation the year before, and a world without regulation removes all three effects together. The set is built for that sum, with a hypothesized effect of the same form: stringency Pct moved log labor intensity by $s_0 \times \text{Pct}$ in the period in which it applied. Subtracting that quantity from the outcome lowers the estimated sum by exactly s_0 and leaves the residuals unchanged, so the test proceeds as in the single-coefficient case.

The hypothesized effect is more restrictive than the specification it is applied to. The estimator allows the response to be spread across three periods, while the assumption holds

that it arrives in the period that stringency applies. The restriction is imposed because a permutation test requires a hypothesis that determines every observation's untreated outcome, and a statement about the sum alone does not, since it leaves open how the effect is distributed across periods. Assigning three separate values instead would define a three-dimensional region of hypotheses consistent with the data, and reducing that to a set for the sum requires either treating the lag and lead as nuisance parameters to be profiled out, which is no longer exact, or a projection conservative enough to be uninformative. At $s_0 = 0$ every version of the assumption coincides, so the restriction does not affect the test of no effect. It bears only on the width of the set.

Relating the estimates to the observed decline

To compare an estimate with the change it might explain, I express it in hours per hectare in 2013. Holding each municipality's harvested area and every term other than stringency at its 2013 value, moving to a world without regulation changes municipality j's log labor intensity by $-x_j$, where $x_j = \hat{\beta}\,\mathrm{Pct}_{s,2013}$, so its counterfactual hours are $h_j e^{-x_j}$. The contribution of stringency to aggregate labor intensity in 2013 is

$$\Delta = \frac{\sum_j h_j - \sum_j h_j e^{-x_j}}{\sum_j a_j},$$

where h_j is observed hours, a_j is harvested area, and the sums run over the 2013 observations of the municipalities in the estimation sample. Observed hours are used rather than fitted hours because exponentiating a fitted log understates the level, while the actual and counterfactual terms built from observed hours carry the same residual. The counterfactual moves hours per hectare and not hectares, which is consistent with excluding provisions on newly planted area from Pct: area is a margin that stringency could itself affect. The base period for the decline has to fall between 1999 and 2001, the only years in which every state's stringency is zero. Written in full, the contribution is the difference between actual and counterfactual labor intensity at both ends of the window, and the base-period terms cancel only because stringency is zero in both the actual and counterfactual worlds.

Aggregate labor intensity fell from 82.27 hours per hectare over 1999 to 2001 to 21.83 in 2013, a decline of 60.45. It is measured over the municipalities observed in each period, 411 in the base and 617 in 2013, so the change in labor intensity arises from the combination of compositional differences in entering municipalities and trends in municipalities observed in both periods. Restricting both endpoints to the 354 municipalities observed in both gives a decline of 54.3 rather than 60.5 hours per hectare. I report the regional aggregate because the claim concerns the study region rather than a fixed panel, and because holding the set

fixed would condition the denominator on whether a municipality is still growing cane in 2013, the same post-treatment selection avoided in the estimation sample. Both figures are reported.

At the headline point estimate stringency accounts for 0.72 hours per hectare, 1.2 percent of the observed decline. The lower end of its 95% set, $\beta = -1.56$, corresponds to 50.6 percent. The headline specification therefore rejects any contribution larger than about half of the decline and rejects no smaller one, including zero. The specification curve below tests whether this holds beyond the headline specification.

Specification curve

To test the robustness of these results, I estimate the model under every combination of five choices, fixed together with the headline before estimation: (i) the sample, which is all municipalities, those harvesting cane in at least one year of 1999 to 2001, those harvesting in all three, or a balanced panel; (ii) weighting, either none or by harvested area; (iii) the upper tail of the log outcome, either unwinsorized or Winsorized at the first and 99th percentiles; (iv) dynamics, either contemporaneous, contemporaneous on the sample that supports a lag and a lead, or with a lag and a lead; and (v) the form of stringency, entered linearly or as an indicator for any positive stringency. A separate indicator for each level of stringency is not identified as a dose response with three states whose stringency changes, so it is not among the forms. All 96 resulting specifications are estimated, tested and displayed in Fig. 4.

Each displayed estimate is Δ expressed in log points, $\ln\left(\sum_j h_j / \sum_j h_j e^{-x_j}\right)$, so that the two forms share one axis. For specifications with a lag and a lead, $x_j = (\hat{\beta}+\hat{\beta}_L+\hat{\beta}_F)\,\mathrm{Pct}_{s,2013}$, the effect of stringency held permanently at its 2013 level. The lead is not evaluated at 2014 stringency, because the lag-and-lead sample ends in 2012 and no stringency later than 2013 enters its estimation. Each specification is evaluated over its own estimation sample, so across samples the curve varies the population as well as the specification. The reference line marks the full observed decline, $\ln(21.83/82.27) = -1.33$; the axis is not linear in the share of that decline, which is why shares are reported in hours per hectare in the text.

The rule for reading the curve was also fixed in advance. The headline specification's p-value is the result, and the grid is summarized by the distribution of its p-values rather than by the smallest. No specification rejects the null of no effect: none of the 96 on the contemporaneous coefficient, and none of the 32 lag-and-lead specifications on the sum of coefficients, on the lag, or on the lead. The headline p-value is 0.867, a rank of 104 of 120. All 96 sets contain zero. The point estimates range from -0.79 to 0.20 log points, a narrower range than the headline specification's own set of $[-0.88, 1.21]$. The linear and indicator forms differ in sign in the headline sample, with coefficients of -0.064 and 0.197. Area

weighting concentrates the estimate on a small number of large municipalities and widens the sets: 24 of the 48 area-weighted sets extend below the plotted range, reaching −8.20, and 7 are not a single interval. Unwinsorized and Winsorized specifications interleave along the curve.

The two forms of stringency can disagree because, under either form, two-way fixed effects need not recover a positively weighted average of municipality-level effects. With continuous stringency, the coefficient compares municipalities by how far their stringency lies from the mean, and it can combine negatively weighted effects with selection bias across doses[52]. With a binary indicator adopted at different times, municipalities already treated serve as comparisons for those treated later, and effects that change over time can receive negative weight[51]. The heterogeneity-robust estimators described below address this problem.

The rejection of large contributions in the headline specification does not extend across the curve. Of the 96 sets, 51 contain the full observed decline, 9 of 48 among unweighted specifications and 42 of 48 among area-weighted ones. The wide confidence sets do not provide the kind of tight bound calculated from the regression discontinuity results; instead, the difference-in-differences results consistently find no detectable effect of regulation, with point estimates that are mostly small.

Heterogeneity-robust estimators

Because stringency is continuous and states adopt at different times, the two-way fixed effects estimator above recovers a weighted average of the underlying treatment effects in which individual weights can be negative when those effects are heterogeneous[50–52]. A large literature now offers estimators robust to that problem[52, 74–81], and I re-estimate the headline specification with two of them (Extended Data Table ED1).

Most of these estimators are built for a binary treatment that stays on once adopted[74, 75, 78], and the leading estimator for a continuous treatment holds each unit's dose fixed after adoption[52]. Stringency here changes more than once after a state is first treated, rising twice in São Paulo and every year after 2010 in Mato Grosso do Sul, so none of these applies without redefining the treatment. The estimators I use are designed for treatments like this one.

The first, and the closer fit, is an event-study estimator for a non-binary treatment that changes more than once and may affect the outcome with lags[80]. It compares each municipality's outcome in successive years from the first change in its stringency with municipalities at the same baseline stringency whose stringency has not yet changed, and I normalize each effect by the cumulative change in stringency so that it is expressed per unit. Without that normalization, and because every state starts at zero, the estimator is equivalent to applying

the estimator of Callaway and Sant'Anna to a binary indicator for whether a state's stringency has ever changed[74, 80], which measures the effect of exposure averaged over increases that differ in size and timing across states rather than a response per unit of stringency. Every state starts at zero, so the comparison involves no extrapolation across baseline levels, and the comparison group consists of the three never-treated states and, until they are first treated, Mato Grosso do Sul and Rio de Janeiro. Its placebo estimators, which make the same comparisons before the first change, provide a further check on parallel trends.

The second compares municipalities whose stringency changes between consecutive years with municipalities whose stringency does not, and recovers the average across the former of the change in log labor intensity per unit change in stringency[81]. It has an unweighted version and a version that weights each municipality by the size of its change. The two coincide if and only if the responses are uncorrelated with the size of the change, so the comparison between them bears directly on the heterogeneity that motivates the exercise. This estimator assumes that only current stringency affects the outcome, which is harder to defend for a requirement that ratchets up. It also compares municipalities at the same baseline stringency, but no municipality with unchanged stringency shares a positive baseline stringency with a changing one in the same year, so for municipalities already regulated the comparison rests on a regression of unchanged municipalities' outcome changes on their baseline stringency. Because it uses only consecutive-year comparisons, it draws on 904 observations with changing stringency and 2,166 without, fewer than the 7,710 in the headline sample. For these reasons I treat the event-study estimates as the more credible of the two.

Extended Data Table ED1 reports the estimates, and three points stand out. First, none of the summary estimates is large relative to the observed decline. The event-study estimator gives an average total effect per unit of stringency of 0.111, and the most negative summary estimate, a weighted average slope of -0.190, corresponds to about 0.1 log points at the headline sample's area-weighted mean stringency in 2013, against an observed decline of 1.33. Second, the estimates do not agree in sign. The unweighted average slope is 0.107, which suggests that how the underlying responses are weighted may matter for the direction of the estimate, although without inference the difference cannot be distinguished from sampling variation. Third, the pre-period estimates are large relative to the treatment estimates and are not centered on zero; the first placebo of the weighted average slope is -0.557. These placebos cover at most three years before a change, and the longer-window evidence on pre-treatment trends is the test described above.

I report these as point estimates and draw no inference from them. The estimators change which weighted average of treatment effects is recovered; they do not change that stringency varies at the state-year level across six states, of which three ever change. The confidence

intervals and tests they report are computed by clustering on state, the approach rejected above. Two of those tests reject at the 5% level, equality of the two average slopes ($p = 0.023$) and joint nullity of the event-study effects ($p = 0.036$), and both rest on the six-cluster approximation. The 120-assignment permutation could be applied to either. Inference for the difference-in-differences design remains the randomization test applied to the specification above.

Interpretation and limitations of the difference-in-differences design

Four limitations bear on how these estimates should be read. First, the design detects no effect of stringency and, under its own inference, cannot bound one. Three states ever change their stringency, São Paulo supplies 98 percent of treated municipality-years, and point estimates from different estimands differ in sign.

Second, identification rests on a parallel trends assumption: absent any change in stringency, labor intensity in states whose stringency changed would have followed the same path as in states whose stringency did not. Because treatment is continuous, reading the coefficient as a dose response requires a stronger form of that assumption, i.e. it holds at every level of stringency rather than only between regulated and unregulated states. Parallel trends is also specific to the functional form, which is why that choice is disclosed above. São Paulo regulated first and also mechanized first, so early in the sample the only treatment variation comes from one state which may have begun mechanizing for other reasons. Mato Grosso do Sul and Rio de Janeiro add variation from 2010 and 2012, but only 71 treated municipality-years between them. The pre-treatment trend tests do not reach São Paulo beyond 1999 to 2001. The lead of stringency does, and the randomization test does not reject it in any of the 32 specifications that include it, although it carries the same imprecision that limits the design as a whole.

Third, the design assumes that regulation in one state does not affect harvesting practice in another. The six states are contiguous and share a single industry, so a change in stringency in one can reach its neighbors through the regional markets for harvest labor and machine services, or through the diffusion of practice among growers and contractors. Spillovers of this kind make the comparison states partly treated and bias the estimate toward zero, understating the effect of regulation. A related channel operates on the whole region at once. If regulation raised expected demand for harvesting machinery, suppliers may have expanded capacity and lowered the price of machine services everywhere, and an effect common to all states is absorbed by the year fixed effects rather than attributed to regulation. Rising wages could have prompted the same anticipation, so the channel is not specific to regulation, but no design identified from cross-state differences can detect it.

Fourth, the headline estimate targets a different average of municipality-level responses than the aggregate it is compared with. The contribution to aggregate labor intensity weights each municipality's response by its hours and its stringency, while the unweighted regression gives each municipality-year roughly equal weight. The two coincide if responses are unrelated to municipality size. Weighting by harvested area approximates the aggregation weight, but the exact weight is hours, the numerator of the outcome, and in this design area weighting concentrates the estimate on a small number of large municipalities and widens the sets severalfold. Whether responses vary with size cannot be tested here: Pct is constant within a state and year, so any comparison of large and small municipalities is also a comparison across states. I report the unweighted estimate and disclose the mismatch.

Instrumental variables strategy

I estimate the wage elasticity of labor demand in sugarcane using instrumental variables. For each of four crops grown in the sugarcane region (maize, soybeans, coffee, oranges), I interact the logged output from the top 3 producers of that crop (excluding Brazil) with the historical municipality-level area cultivated of that crop (1994–1998 average).

I include lags to account for the possibility that labor demand does not respond only to contemporaneous price shocks. For the temporary crops maize and soybeans, 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

$$\begin{aligned} Z_{j,t} \equiv \{ & \text{Int}_j^{maize} \times Y_t^{maize},\ \text{Int}_j^{maize} \times Y_{t-1}^{maize}, \\ & \text{Int}_j^{soy} \times Y_t^{soy},\ \text{Int}_j^{soy} \times Y_{t-1}^{soy}, \\ & \text{Int}_j^{cof} \times Y_t^{cof},\ \text{Int}_j^{cof} \times Y_{t-1}^{cof},\ ...,\ \text{Int}_j^{cof} \times Y_{t-4}^{cof}, \\ & \text{Int}_j^{orng} \times Y_t^{orng},\ \text{Int}_j^{orng} \times Y_{t-1}^{orng},\ ...,\ \text{Int}_j^{orng} \times Y_{t-4}^{orng} \}, \end{aligned} \tag{3}$$

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

The first-stage regression is

$$\log w_{j,t}^S = \gamma_0 + \gamma_1 X_{j,t} + \Gamma Z_{j,t} + u, \tag{4}$$

where w^S is the median real monthly wage for manual laborers in sugarcane, j indexes municipality, t indexes year, and X includes the log of sugarcane area harvested to control for increases in labor demand associated with expanding sugarcane cultivation.

The second-stage regression is

$$\log L^S_{j,t} = \beta_0 + \beta_1 X_{j,t} + \beta_2 \widehat{\log w^S_{j,t}} + \varepsilon, \tag{5}$$

where L^S is hours worked by manual laborers employed by sugarcane growers. The sample includes the years 1999–2013, with approximately 400 municipalities in the early years increasing to almost 800 by the end of the sample. Standard errors are clustered by municipality.

Identifying assumptions and specification choices

The instruments take the standard shift-share form: a municipality-specific share, the historical area of crop c, interacted with a common shock, the output of the top three foreign producers of crop c. Validity can be argued from either component[82, 83]. Here the shares are fixed at their 1994–1998 values, before the sample begins and before the sugarcane expansion, and the shocks are constructed to exclude Brazilian output. The latter choice is deliberate: this region grows 20 to 30 percent of the world's coffee, oranges, soybeans, and sugarcane, so a shock to Brazilian output would move international prices and sugarcane labor demand simultaneously and could not satisfy the exclusion restriction. Using foreign output does not eliminate the concern entirely, since large foreign producers may adjust output in anticipation of conditions in Brazil, but the consistency of the estimates across crops with different producers and different growing calendars limits how much of the result such behavior could explain. Maize is the least susceptible: Brazil accounts for less than 10 percent of world production while the United States and China together account for more than three quarters, so anticipatory responses by maize producers are unlikely to be organized around Brazilian conditions.

The specification includes neither municipality nor year fixed effects, and this choice deserves explicit statement because it is consequential. Both components of the instrument are close to collinear with such effects: the shares vary only across municipalities and the shocks vary only over time, so a specification saturated with both sets of fixed effects retains identifying variation only through the interaction and, in practice, has a first stage too weak to be informative. The cost is that the estimates rely on cross-sectional comparisons between municipalities with different historical crop mixes and on aggregate movements in foreign crop output over the sample period. The identifying assumption is correspondingly stronger: historical crop area must be unrelated to the evolution of sugarcane labor demand except through the labor-supply channel, and the foreign output series must be unrelated to Brazilian aggregate trends that independently affect mechanization. A national decline

in the cost of harvesting machinery, for example, would not be absorbed. The single-crop columns provide a partial check, since they rely on distinct shares and distinct shocks and would not be expected to agree if a common aggregate trend were driving the result.

First-stage strength is reported in Table 1. Two columns fall below the conventional threshold of 10, oranges ($F = 5.1$) and the pooled specification ($F = 9.5$, with twelve excluded instruments), so their confidence intervals understate sampling uncertainty and should be interpreted with the corrections proposed in the weak-instrument literature in mind[45, 84]. The maize ($F = 19.9$) and coffee ($F = 15.1$) columns do not have this problem and are the basis for the substantive claim.

Tipping-point model

I develop a model showing that S-shaped adoption can emerge from factor prices alone, in a full-information environment with no fixed costs and no improvements to technology.

Each harvesting technology has a CES production function with constant returns to scale. Manual harvesting (denoted p) and mechanical harvesting (denoted m) have 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 with $T_m + T_p = 1$. Production functions are:

$$Y_m = A_m \left(\alpha K_m^{-\eta} + \beta L_m^{-\eta} + (1 - \alpha - \beta) T_m^{-\eta} \right)^{\frac{-1}{\eta}}, \tag{6}$$

$$Y_p = A_p \left(\gamma K_p^{-\zeta} + \delta L_p^{-\zeta} + (1 - \gamma - \delta) T_p^{-\zeta} \right)^{\frac{-1}{\zeta}}. \tag{7}$$

Growers solve:

$$\max_{L,K,T} Y_m + Y_p - w_m L_m - r_m K_m - w_p L_p - r_p K_p \tag{8}$$

$$\text{s.t. } T_p + T_m = 1. \tag{9}$$

I interpret A_m as suitability for mechanical harvesting (high for flat parcels with optimal row layout; low for steep parcels) and A_p as the productivity boost from burning (high for burned parcels; low for unburned parcels). Factor prices are exogenous to any individual parcel.

With constant returns to scale and linear costs, growers will generally allocate an entire parcel to only one harvesting technique. This leads to the following result:

Proposition 1. *There exists a threshold manual wage for each parcel, above which growers mechanize and below which they harvest manually. The threshold wage ϕ satisfies $\partial\phi/\partial A_p >$*

0 *and* $\partial\phi/\partial A_m < 0$.

Since parcels differ in steepness, layout, and other characteristics, productivities differ across parcels, giving rise to a distribution of threshold wages. If the cumulative distribution of ϕ is S-shaped, steady wage increases can produce abrupt changes in aggregate mechanization rates. Empirical CDFs of parcel steepness and area support this shape (Extended Data Fig. ED3).

The model is frictionless by construction. It contains no learning, no credit constraint, and no indivisibility, because harvesting machines are rented more often than bought and fewer than 10 percent of establishments using mechanical harvesting own one. S-shaped diffusion curves are often read as evidence of such frictions, but they can emerge from static profit maximization alone when producers differ in the price at which the new technology becomes cheaper[30].

The model holds the rental price of harvesting machinery, r_m, fixed for simplicity. What the mechanism requires is not a constant r_m but a rising ratio w/r_m, so the argument goes through if machinery rental prices fell, or rose more slowly than wages. A fall in r_m lowers threshold wages for every parcel and would generate abrupt aggregate adoption through the same S-shaped distribution, so the tipping-point mechanism is not specific to wages. The empirical analysis identifies movements along the labor-demand curve and does not separately identify the contribution of machinery costs; the Discussion returns to this point.

Data availability

The analysis uses four sources. The RAIS employment micro data are confidential and are maintained by the Brazilian Ministry of Labor; they were obtained under an agreement permitting storage and analysis at a secure facility, and access must be requested from the Ministry directly. The 2006 Census of Agriculture micro data are confidential and were analyzed on-site at an IBGE facility in Rio de Janeiro under supervision; access must be requested from IBGE. Neither confidential source may be redistributed by the author. The PAM and PNAD surveys are publicly available from IBGE, and international crop output data are publicly available from the FAO's FAOSTAT database. Aggregated data sufficient to reproduce every figure and table in this paper, at a level of aggregation cleared for release, are provided with the Supplementary Information. Notes from the industry interviews described in Methods are not released, because they identify participants.

Code availability

All estimation code is available from the author and is deposited with the Supplementary Information. The regression discontinuity estimates use the `rdrobust` package implementing Calonico et al. [72] and Calonico et al. [73]. The heterogeneity-robust difference-in-differences estimates use the `did_multiplegt_stat` and `did_multiplegt_dyn` packages implementing de Chaisemartin et al. [81] and de Chaisemartin and D'Haultfœ uille [80]; package versions are recorded in the replication log, because both update automatically.

Author contributions

C.A.D. is the sole author and is responsible for all aspects of the study.

Competing interests

The author declares no competing interests.

References

[1] Vijesh V. Krishna and Maxwell Mkondiwa. Economics of Crop Residue Management. *Annual Review of Resource Economics*, 15:19–39, 2023. doi: 10.1146/annurev-resource-101422-090019.

[2] María Uriarte, Charles B. Yackulic, Tamar Cooper, Dan Flynn, Marina Cortes, Tanja Crk, Georgina Cullman, Meghan McGinty, and Jason Sircely. Expansion of Sugarcane Production in São Paulo, Brazil: Implications for Fire Occurrence and Respiratory Health. *Agriculture, Ecosystems & Environment*, 132(1-2):48 – 56, 2009. ISSN 0167-8809. doi: 10.1016/j.agee.2009.02.018. URL `http://www.sciencedirect.com/science/article/pii/S0167880909000760`.

[3] Flavia Mazzoli-Rocha, Clarissa Bichara Magalh aes, Olaf Malm, Paulo Hilário Nascimento Saldiva, Walter Araujo Zin, and Débora Souza Faffe. Comparative Respiratory Toxicity of Particles Produced by Traffic and Sugar Cane Burning. *Environmental Research*, 108(1):35 – 41, 2008. ISSN 0013-9351. doi: 10.1016/j.envres.2008.05.004. URL `http://www.sciencedirect.com/science/article/pii/S0013935108001278`.

[4] Marcos A Arbex, Gyorgy M Böhm, Paulo HN Saldiva, Gleice MS Conceiçao, Arden C Pope III, and Alfesio LF Braga. 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, 2000.

[5] Marcos Abdo Arbex, Lourdes Conceição Martins, Regiani Carvalho de Oliveira, Luiz Alberto Amador Pereira, Flávio Ferlin Arbex, José Eduardo Delfini Cançado, Paulo Hilário Nascimento Saldiva, and Alfésio Luís Ferreira Braga. 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, 2007. doi: 10.1136/jech.2005.044743. URL `http://jech.bmj.com/content/61/5/395.abstract`.

[6] José E. D. Cançado, Paulo H. N. Saldiva, Luiz A. A. Pereira, Luciene B. L. S. Lara, Paulo Artaxo, Luiz A. Martinelli, Marcos A. Arbex, Antonella Zanobetti, and Alfesio L. F. Braga. The Impact of Sugar Cane-Burning Emissions on the Respiratory System of Children and the Elderly. *Environmental Health Perspectives*, 114(9):725–729, May 2006.

[7] Andres Manrique, Kayan Clarke, Sarah Bisesi, Farah A. Arosemena, Eric S. Coker, and Tara Sabo-Attwood. The Adverse Health Effects of Air Pollution from Sugarcane Burning: A Scoping Review of Observational and Experimental Evidence. *Environmental Health Perspectives*, 133(1):016002, 2025. doi: 10.1289/EHP14456.

[8] Marcos Abdo Arbex, Paulo Hilário Nascimento Saldiva, Luiz Alberto Amador Pereira, and Alfésio Luís Ferreira Braga. Impact of Outdoor Biomass Air Pollution on Hypertension Hospital Admissions. *Journal of Epidemiology and Community Health*, 64(7):573–579, 2010. doi: 10.1136/jech.2009.094342. URL `http://jech.bmj.com/content/64/7/573.abstract`.

[9] Danielle Miyuki Goto, Marina Lança, Camila Akemi Obuti, Cristiane Maria Galv ao Barbosa, Paulo Hilário Nascimento Saldiva, Dirce Maria Trevisan Zanetta, Geraldo Lorenzi-Filho, Ubiratan de Paula Santos, and Naomi Kondo Nakagawa. Effects of Biomass Burning on Nasal Mucociliary Clearance and Mucus Properties After Sugarcane Harvesting. *Environmental Research*, 111(5):664 – 669, 2011. ISSN 0013-9351. doi: 10.1016/j.envres.2011.03.006. URL `http://www.sciencedirect.com/science/article/pii/S0013935111000764`.

[10] Marcos A. Rangel and Tom Vogl. Agricultural Fires and Infant Health. Working Paper 22955, National Bureau of Economic Research, December 2016. URL `http://www.nber.org/papers/w22955`.

[11] Oeindrila Dube and Juan F. Vargas. Commodity Price Shocks and Civil Conflict: Evidence from Colombia. *The Review of Economic Studies*, 80(4):1384–1421, 2013. doi: 10.1093/restud/rdt009. URL `http://restud.oxfordjournals.org/content/80/4/1384.abstract`.

[12] José Maria Tomazela. Polícia Ambiental adota critério que anistia queima de cana em SP. *O Estado de S. Paulo*, May 16 2016.

[13] Evandro Spinelli. Usinas mantêm queimadas na região. *Folha de S. Paulo*, June 21 2001.

[14] José Sergio Osse. Colheita mecânica de cana-de-açúcar fecha 300 mil vagas. *Folha de S. Paulo*, July 2 2002.

[15] Gazeta Mercantil. O fim da queima da palha da cana. November 12 2002.

[16] Roberto Samora. Usinas são multadas em R$ 1,6 mi por queima irregular de cana. *Reuters*, August 18 2006.

[17] José Ernesto Credendio. Usinas lideram ranking de multas por poluição em SP. *Folha de S. Paulo*, June 8 2008.

[18] Folha de S. Paulo. Usina diz que recorreu de multa ambiental. June 1 2008.

[19] Brás Henrique. Cetesb multa produtores de cana na região de Ribeirão. *O Estado de S. Paulo*, August 4 2008.

[20] Juliana Coissi. Usinas lideram ranking da poluição ambiental na região. *Folha de S. Paulo*, December 16 2008.

[21] O Globo G1. Polícia Ambiental usa imagens de satélite para identificar queimadas, July 13 2016.

[22] Fernando Valente and Márcio Laurini. Pre-harvest sugarcane burning: A statistical analysis of the environmental impacts of a regulatory change in the energy sector. *Cleaner Engineering and Technology*, 4:100255, 2021. doi: 10.1016/j.clet.2021.100255.

[23] Seema Jayachandran. How Economic Development Influences the Environment. *Annual Review of Economics*, 14:229–252, 2022. doi: 10.1146/annurev-economics-082321-123803.

[24] Paul A David. The Mechanization of Reaping in the Ante-bellum Midwest. *Industrialization in Two Systems*, 30, 1966.

[25] Alan L. Olmstead and Paul W. Rhode. Induced Innovation in American Agiculture: A Reconsideration. *Journal of Political Economy*, pages 100–118, 1993.

[26] Alan L. Olmstead and Paul W. Rhode. Beyond the threshold: An analysis of the characteristics and behavior of early reaper adopters. *The Journal of Economic History*, 55:27–57, 3 1995. ISSN 1471-6372.

[27] Michael A. Clemens, Ethan G. Lewis, and Hannah M. Postel. Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion. *American Economic Review*, 108(6):1468–1487, 2018. doi: 10.1257/aer.20170765.

[28] Richard Hornbeck and Suresh Naidu. When the Levee Breaks: Black Migration and Economic Development in the American South. *American Economic Review*, 104(3): 963–990, 2014. doi: 10.1257/aer.104.3.963.

[29] Daron Acemoglu and Pascual Restrepo. Demographics and Automation. *The Review of Economic Studies*, 89(1):1–44, 2022. doi: 10.1093/restud/rdab031.

[30] Rodolfo E. Manuelli and Ananth Seshadri. Frictionless Technology Diffusion: The Case of Tractors. *American Economic Review*, 104(4):1368–91, April 2014. doi: 10.1257/aer.104.4.1368. URL `http://www.aeaweb.org/articles?id=10.1257/aer.104.4.1368`.

[31] Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hémous. The Environment and Directed Technical Change. *American Economic Review*, 102(1):131–166, 2012. doi: 10.1257/aer.102.1.131.

[32] Andrew D Foster and Mark R Rosenzweig. Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *Journal of Political Economy*, pages 1176–1209, 1995.

[33] Timothy G. Conley and Christopher R. Udry. Learning about a New Technology: Pineapple in Ghana. *The American Economic Review*, 100(1):pp. 35–69, 2010. ISSN 00028282. URL `http://www.jstor.org/stable/27804921`.

[34] Esther Duflo, Michael Kremer, and Jonathan Robinson. Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. *American Economic Review*, 101(6):2350–2390, 2011.

[35] Tavneet Suri and Christopher Udry. Agricultural Technology in Africa. *Journal of Economic Perspectives*, 36(1):33–56, 2022. doi: 10.1257/jep.36.1.33.

[36] Michael Morris, Valerie A. Kelly, Ron J. Kopicki, and Derek Byerlee. *Fertilizer Use in African Agriculture*. The World Bank, 2007. doi: 10.1596/978-0-8213-6880-0. URL http://elibrary.worldbank.org/doi/abs/10.1596/978-0-8213-6880-0.

[37] Douglas Gollin. Chapter 73 Agricultural Productivity and Economic Growth. volume 4 of *Handbook of Agricultural Economics*, pages 3825 – 3866. Elsevier, 2010. doi: http://dx.doi.org/10.1016/S1574-0072(09)04073-0. URL http://www.sciencedirect.com/science/article/pii/S1574007209040730.

[38] Douglas Gollin, David Lagakos, and Michael E. Waugh. The Agricultural Productivity Gap. *The Quarterly Journal of Economics*, 129(2):939–993, 2014. doi: 10.1093/qje/qjt056.

[39] Klaus Deininger, Songqing Jin, and Meilin Ma. Structural Transformation of the Agricultural Sector in Low- and Middle-Income Economies. *Annual Review of Resource Economics*, 14:221–241, 2022. doi: 10.1146/annurev-resource-111820-033252.

[40] Paula Bustos, Bruno Caprettini, and Jacopo Ponticelli. Agricultural Productivity and Structural Transformation: Evidence from Brazil. *American Economic Review*, 106(6):1320–1365, 2016. doi: 10.1257/aer.20131061.

[41] Christopher B. Barrett, Ariel Ortiz-Bobea, and Trinh Pham. Structural Transformation, Agriculture, Climate, and the Environment. *Review of Environmental Economics and Policy*, 17(2):195–216, 2023. doi: 10.1086/725319.

[42] Lucas W. Davis. The Effect of Driving Restrictions on Air Quality in Mexico City. *Journal of Political Economy*, 116(1):pp. 38–81, 2008. ISSN 00223808. URL http://www.jstor.org/stable/10.1086/529398.

[43] Michael Greenstone and Rema Hanna. Environmental Regulations, Air and Water Pollution, and Infant Mortality in India. *American Economic Review*, 104(10):3038–72, 2014. doi: 10.1257/aer.104.10.3038. URL http://www.aeaweb.org/articles.php?doi=10.1257/aer.104.10.3038.

[44] Michael Greenstone and B. Kelsey Jack. Envirodevonomics: A Research Agenda for an Emerging Field. *Journal of Economic Literature*, 53(1):5–42, 2015. doi: 10.1257/jel.53.1.5.

[45] José Luis Montiel Olea and Carolin Pflueger. A Robust Test for Weak Instruments. *Journal of Business & Economic Statistics*, 31(3):358–369, 2013.

[46] Isaiah Andrews, James H. Stock, and Liyang Sun. Weak Instruments in Instrumental Variables Regression: Theory and Practice. *Annual Review of Economics*, 11:727–753, 2019. doi: 10.1146/annurev-economics-080218-025643.

[47] Governo do Estado do São Paulo. Lei Estadual No 11.241, de 19 de Setembro de 2002.

[48] Secretaria do Meio Ambiente. Protocolo Agroambiental – Diretivas. http://www.ambiente.sp.gov.br/etanolverde/protocolo-agroambiental/diretivas/. Accessed February 9, 2013.

[49] Brazilian Sugarcane Industry Association (UNICA). The Industry - Background. http://english.unica.com.br/. Accessed February 9, 2013.

[50] Clément de Chaisemartin and Xavier D'Haultfœ uille. Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9): 2964–2996, 2020.

[51] Andrew Goodman-Bacon. Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics*, 225(2):254–277, 2021. doi: 10.1016/j.jeconom.2021.03.014.

[52] Brantly Callaway, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. Difference-in-Differences with a Continuous Treatment. NBER Working Paper 32117, National Bureau of Economic Research, 2024.

[53] Priscila Souza, André Sant'Anna, Luciano Machado, Barbara Intropidi, and Pedro Vogt. Credit for Investments in Brazilian Agriculture and the Role of the Brazilian Development Bank. Technical report, Climate Policy Initiative/Pontifícia Universidade Católica do Rio de Janeiro, January 2022. URL `https://www.climatepolicyinitiative.org/publication/credit-for-investments-in-brazilian-agriculture-and-the-role-of-the-brazilian-devel`

[54] Instituto Brasileiro de Geografia e Estatística (IBGE). Produção Agrícola Municipal (PAM). https://www.ibge.gov.br/en/statistics/economic/agriculture-forestry-and-fishing/16773-municipal-agricultural-production-temporary-and-permanent-crops.html.

[55] Instituto Nacional de Pesquisas Espaciais (INPE). CANASAT - Sugarcane Crop Monitoring in Brazil by Earth Observing Satellite Images. http://www.dsr.inpe.br/laf/canasat/en/. Accessed February 10, 2013.

[56] Instituto Brasileiro de Geografia e Estatística (IBGE). Censo Demográfico 2010. https://www.ibge.gov.br/en/statistics/social/population/18391-2010-population-census.html, 2011.

[57] Food and Agriculture Organization of the United Nations (FAO). FAOSTAT Crops and Livestock Products Database. https://www.fao.org/faostat/en/#data/QCL.

[58] M. J. da Silva, L. O. Neves, M. H. F. Correa, and C. H. W. de Souza. Quality Indexes and Performance in Mechanized Harvesting of Sugarcane at a Burnt Cane and Green Cane. *Sugar Tech*, 23(3):499–507, 2021. doi: 10.1007/s12355-021-00957-9.

[59] Secretaria de Meio Ambiente do Estado de São Paulo. Protocolo Agroambiental do Setor Sucroenergético Paulista (Etanol Mais Verde). https://semil.sp.gov.br/sma/etanolverde/, 2007.

[60] Governo do Estado de Goiás. Lei nº 15.834, de 23 de Novembro de 2006: Dispõe sobre redução gradativa da queima da palha de cana-de-açúcar em áreas mecanizáveis. https://legisla.casacivil.go.gov.br/api/v2/pesquisa/legislacoes/79930/pdf, 2006.

[61] Governo do Estado de Mato Grosso do Sul. Lei nº 3.357, de 9 de Janeiro de 2007: Estabelece normas para a redução gradual da queima da palha da cana-de-açúcar. http://aacpdappls.net.ms.gov.br/appls/legislacao/secoge/govato.nsf, 2007.

[62] Governo do Estado de Mato Grosso do Sul. Lei nº 3.404, de 30 de Julho de 2007: Dispõe sobre a localização de estabelecimentos industriais para a produção de açúcar e álcool carburante. http://aacpdappls.net.ms.gov.br/appls/legislacao/secoge/govato.nsf, 2007.

[63] Governo do Estado do Rio de Janeiro. Lei nº 5.990, de 20 de Junho de 2011: Dispõe sobre a eliminação gradativa da queima da palha da cana-de-açúcar. https://alerjln1.alerj.rj.gov.br/contlei.nsf, 2011.

[64] Conselho Estadual de Política Ambiental de Minas Gerais (COPAM). Deliberação Normativa COPAM nº 133, de 15 de Abril de 2009: Regulamenta a prática da queima de cana-de-açúcar para fins de colheita. https://www.siam.mg.gov.br/sla/download.pdf?idNorma=32595, 2009. Publicada no Diário do Executivo, Minas Gerais, 16 de abril de 2009.

[65] Secretaria de Estado do Meio Ambiente e Recursos Hídricos do Paraná. Resolução SEMA nº 76, de 20 de Dezembro de 2010: Dispõe sobre eliminação gradativa da despalha da cana-de-açúcar através da queima controlada. https://www.legisweb.com.br/legislacao/?id=144572, 2010.

[66] Daniel Alves Aguiar, Bernardo Friedrich Theodor Rudorff, Wagner Fernando Silva, Marcos Adami, and Marcio Pupin Mello. Remote sensing images in support of environmental protocol: Monitoring the sugarcane harvest in São Paulo State, Brazil. *Remote Sensing*, 3(12):2682–2703, 2011.

[67] Secretaria de Estado do Meio Ambiente de São Paulo. Resolução SMA nº 088, de 19 de Dezembro de 2008: Define as diretrizes técnicas para o licenciamento de empreendimentos do setor sucroalcooleiro no Estado de São Paulo. https://smastr16.blob.core.windows.net/legislacao/sites/262/2022/07/2008resolucao_sma_088_2008-1.pdf, 2008.

[68] Tribunal de Justiça do Estado de Goiás, Corte Especial. Ação Direta de Inconstitucionalidade nº 356-8/200, processo 200702335821: Lei estadual nº 15.834/2006, 2009. Acórdão de 22 de abril de 2009, publicado no Diário da Justiça nº 338 de 21 de maio de 2009. Artigo 1º e incisos I a V declarados inconstitucionais.

[69] Secretaria da Fazenda do Estado de Goiás. Relatório da Gestão, Exercício 2008, Volume IIA, 2009.

[70] Conselho Estadual do Meio Ambiente de Goiás (CEMAm). Resolução CEMAm nº 082/2007: Dispõe sobre a eliminação da queima da palha da cana-de-açúcar para novas

usinas a serem instaladas no Estado de Goiás. https://supremoambiental.com.br/wp-content/uploads/2018/07/resolucao-n-082-cemam-2007-elim-da-queima-da-palha-da-cana-de-acucar-de-novas-usinas-sucroalcooleiras-em-goias.pdf, 2007.

[71] Governo do Estado do Paraná. Decreto nº 10.068, de 6 de Fevereiro de 2014: Estabelece critérios, prazos e procedimentos para adequação ambiental das usinas de beneficiamento de cana-de-açúcar. https://www.legisweb.com.br/legislacao/?id=265471, 2014.

[72] Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295–2326, 2014. ISSN 1468-0262. doi: 10.3982/ECTA11757. URL `http://dx.doi.org/10.3982/ECTA11757`.

[73] Sebastian Calonico, Matias D. Cattaneo, and Max H. Farrell. On the Effect of Bias Estimation on Coverage Accuracy in Nonparametric Inference. *Journal of the American Statistical Association*, Forthcoming 2016.

[74] Brantly Callaway and Pedro H. C. Sant'Anna. Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225(2):200–230, 2021. doi: 10.1016/j.jeconom.2020.12.001.

[75] Liyang Sun and Sarah Abraham. Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics*, 225(2):175–199, 2021. doi: 10.1016/j.jeconom.2020.09.006.

[76] Clément de Chaisemartin and Xavier D'Haultfœ uille. Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey. *The Econometrics Journal*, 26(3):C1–C30, 2023. doi: 10.1093/ectj/utac017.

[77] Jonathan Roth, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe. What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature. *Journal of Econometrics*, 235(2):2218–2244, 2023. doi: 10.1016/j.jeconom.2023.03.008.

[78] Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting Event-Study Designs: Robust and Efficient Estimation. *Review of Economic Studies*, 91(6):3253–3285, 2024. doi: 10.1093/restud/rdae007.

[79] Andrew Baker, Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. Difference-in-Differences Designs: A Practitioner's Guide. *Journal of Economic Literature*, 64(2):498–557, 2026. doi: 10.1257/jel.20251650.

[80] Clément de Chaisemartin and Xavier D'Haultfœ uille. Difference-in-Differences Estimators of Intertemporal Treatment Effects. *Review of Economics and Statistics*, 108(4):863–880, 2026. doi: 10.1162/rest_a_01414.

[81] Clément de Chaisemartin, Xavier D'Haultfœ uille, Félix Pasquier, Doulo Sow, and Gonzalo Vazquez-Bare. Difference-in-Differences for Continuous Treatments and Instruments with Stayers. Working Paper 2201.06898, arXiv, 2026.

[82] Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8):2586–2624, 2020.

[83] Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, 89(1):181–213, 2022.

[84] David S. Lee, Justin McCrary, Marcelo J. Moreira, and Jack Porter. Valid t-Ratio Inference for IV. *American Economic Review*, 112(10):3260–3290, 2022.

[85] Sebastian Calonico, Matias D. Cattaneo, and Rocío Titiunik. Optimal Data-Driven Regression Discontinuity Plots. *Journal of the American Statistical Association*, 110 (512):1753–1769, 2015. doi: 10.1080/01621459.2015.1017578. URL `http://dx.doi.org/10.1080/01621459.2015.1017578`.

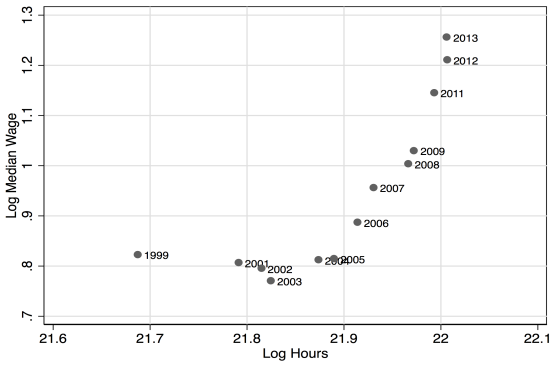
Table 1: IV Estimates of Wage Elasticity of Labor Demand

	(1) Maize	(2) Soybeans	(3) Coffee	(4) Oranges	(5) All	(6) All
$\log(A^S)$	0.852***	0.598***	0.813***	0.868***	0.840***	0.794***
	(0.0653)	(0.0900)	(0.0438)	(0.0560)	(0.0405)	(0.0446)
$\log(w^S)$	-2.938**	1.423	-2.204***	-3.201***	-2.729***	38.86*
	(0.992)	(1.477)	(0.551)	(0.726)	(0.420)	(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	

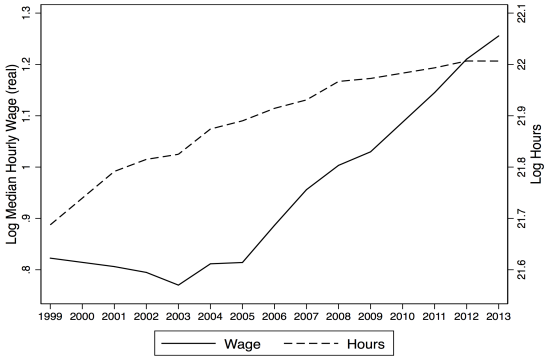
Standard errors clustered by county in parentheses.

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

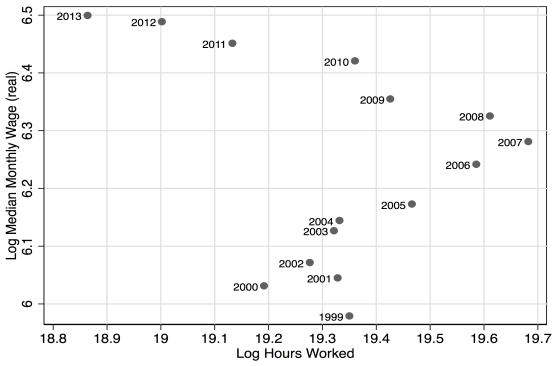
Figure 1: Employment and Real Wages in Brazil and in Sugarcane, 1999–2013



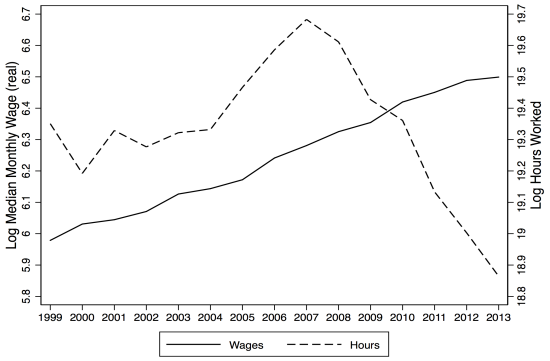
(a) All paid workers, hours on the horizontal axis



(b) All paid workers, by year



(c) Manual laborers in sugarcane, hours on the horizontal axis



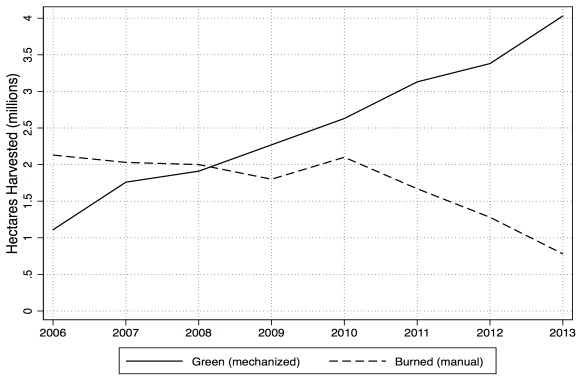
(d) Manual laborers in sugarcane, by year

Panels a and b are from PNAD micro data and include all paid workers; PNAD was not conducted in 2010 because that was a census year. Wages in these panels are real median hourly wages. Panels c and d are from RAIS administrative data and cover manual laborers in sugarcane cultivation in the six South-Central states of the study region; wages in these panels are real median monthly wages, the same series used to estimate the wage elasticity in Table 1. All wages are measured in 2003 R$. Panels a and c place hours on the horizontal axis to emphasize movements of supply and demand; panels b and d place years on the horizontal axis to emphasize the evolution of each series.

Figure 2: The Transition from Manual to Mechanical Harvesting



(a) Labor intensity, study region, 1999–2013



(b) Area harvested by method, São Paulo, 2006–2013

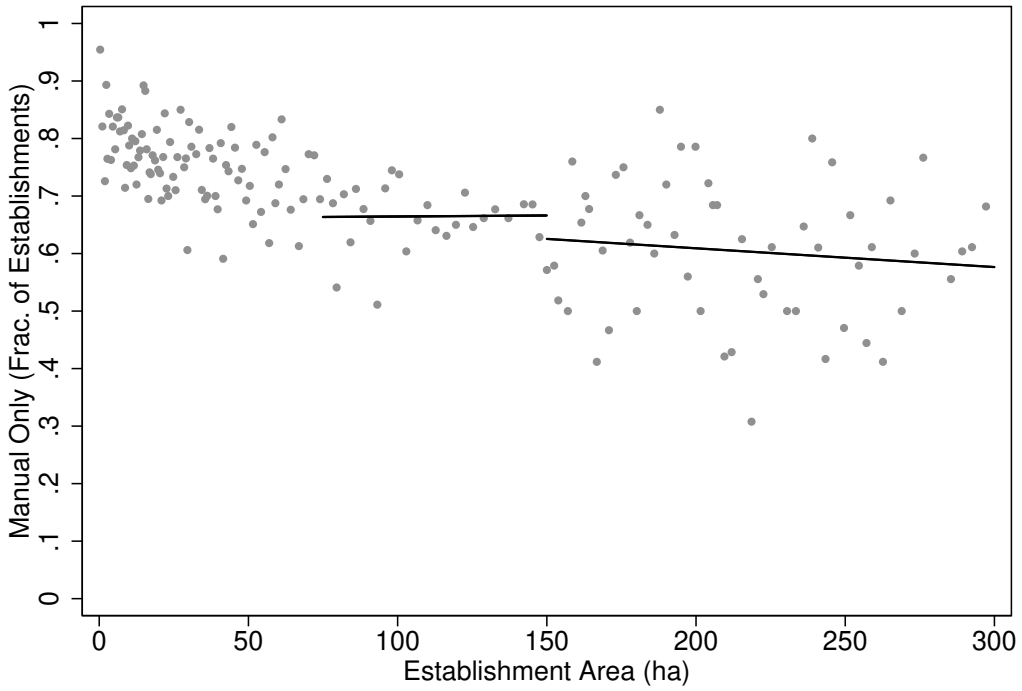
(a) Hours worked by manual laborers in sugarcane divided by area harvested, for the six South-Central states of the study region, from RAIS administrative data and the PAM survey. This is the outcome variable in the difference-in-differences design. (b) Sugarcane area harvested green, that is mechanically and without burning, against area harvested burned, from the CANASAT satellite mapping of harvest method. Panel b begins in 2006, the first year of that series, and covers São Paulo alone, the largest producing state, so its coverage is narrower in both time and space than panel a.

Table 2: The Effect of Regulation on Labor Intensity (Difference-in-Differences)

Choice	Alternative	Effect	RI p	Rank	95% set
Headline specification		-0.032	0.867	104/120	[-0.876, 1.211]
Sample	All municipalities	-0.032	0.883	106/120	
	Cane in all of 1999–2001	0.021	0.875	105/120	
	Balanced panel	-0.263	0.417	50/120	
Weighting	Harvested area	-0.108	0.733	88/120	
Outcome	Winsorized log	-0.044	0.817	98/120	
Dynamics	Lag-and-lead sample	-0.113	0.517	62/120	
	Lag and lead	-0.084	0.825	99/120	
Stringency	Binary indicator	0.141	0.492	59/120	

The headline specification regresses log labor intensity on stringency with municipality and year fixed effects, unweighted, on municipalities that harvested sugarcane in at least one year of 1999–2001 (7,710 municipality-years). Each other row changes one choice and holds the rest fixed. The effect is the change in aggregate 2013 labor intensity, in log points, implied by moving to a world without regulation (Methods); the observed change from 1999–2001 to 2013 is -1.33 log points. Inference is by randomization: the six states' stringency paths are permuted across states, giving 120 distinct assignments, and the rank is that of the observed assignment's studentized statistic (a rank of 6 or lower rejects at 5 percent). For the lag-and-lead row the test is of the sum of the three coefficients. No conventional standard errors are reported (Methods). The 95 percent confidence set inverts the same test and is reported for the headline specification only; it is a single interval.

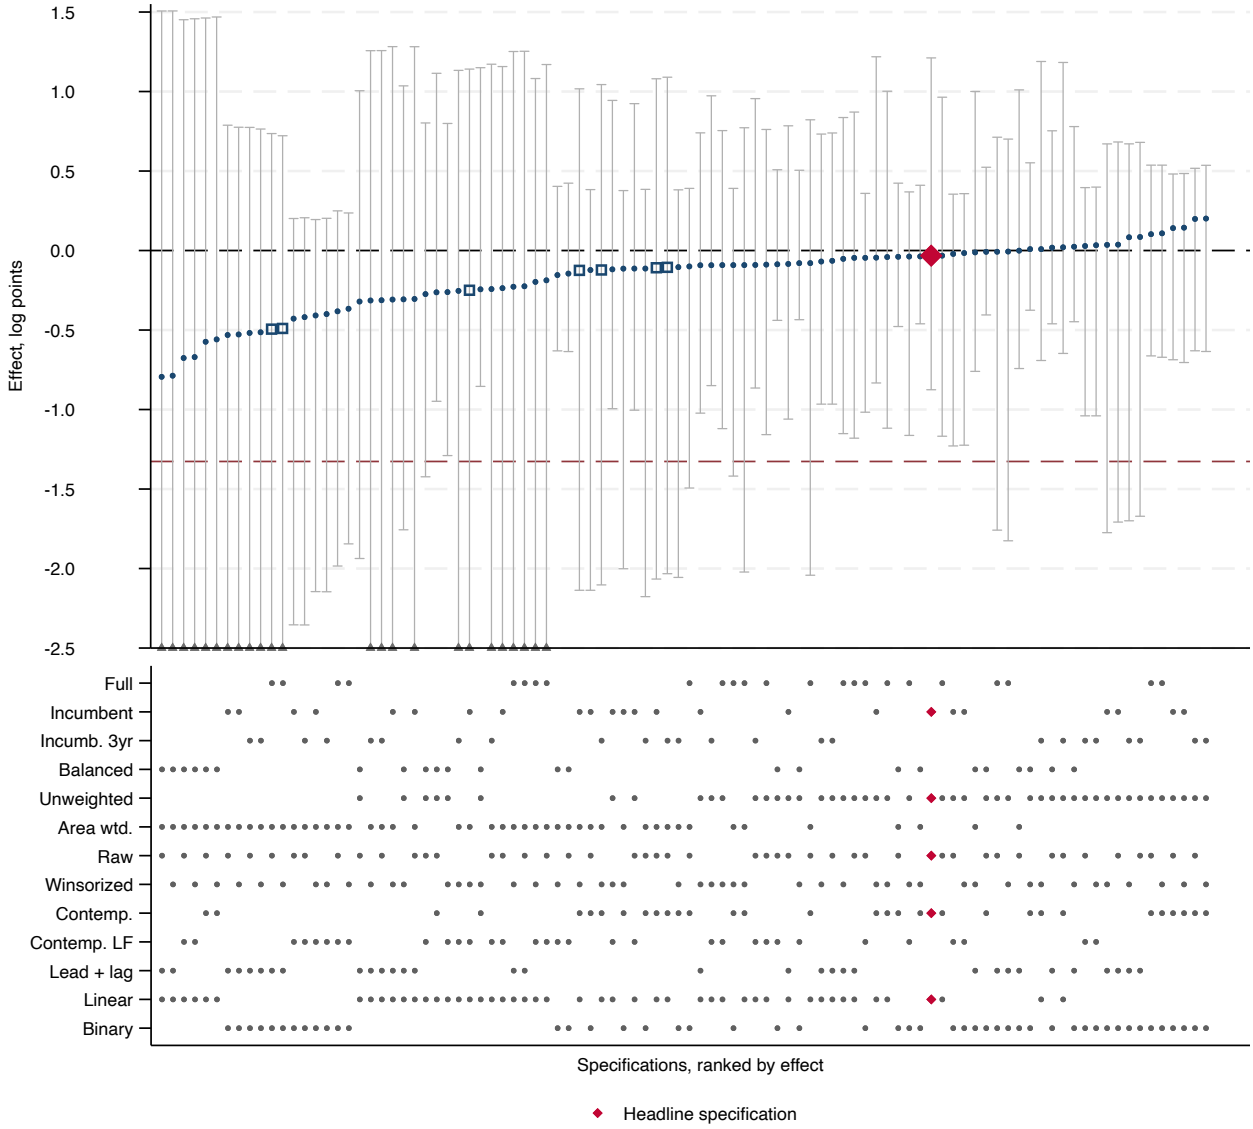
Figure 3: Regression Discontinuity at the 150-Hectare Regulatory Threshold



Source: 2006 Census of Agriculture.

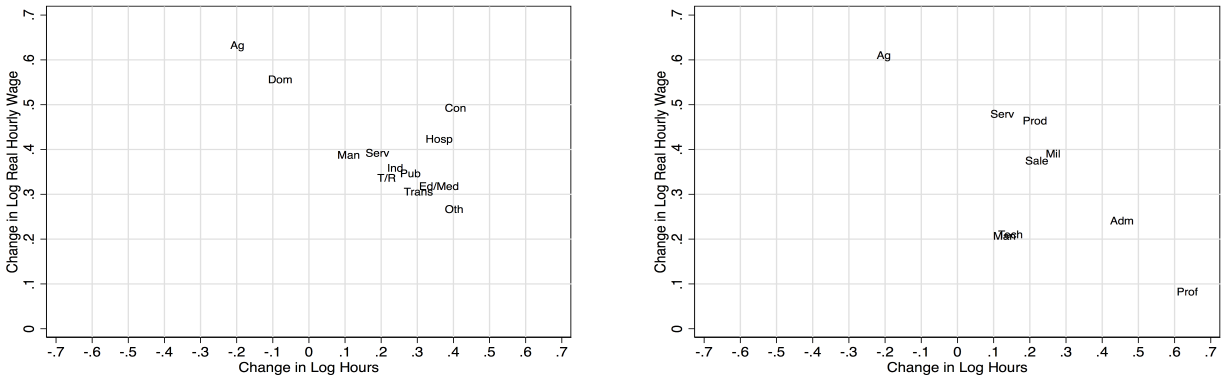
Binned scatter plot and local linear estimates of the share of establishments using manual harvesting only, against establishment area, from the 2006 Census of Agriculture. The estimated discontinuity is $\hat{\tau} = -0.041$ with a 95% confidence interval of $[-0.139, 0.047]$; $h_- = 75$, $h_+ = 755$, $N = 14{,}795$, with 1,863 observations below and 2,402 above the threshold. Local linear estimates are plotted for $x \in [150-h_-, \min\{300, 150+h_+\}]$. Bin sizes are set to mimic the variance of the underlying data[85].

Figure 4: Stringency and Labor Intensity Across 96 Difference-in-Differences Specifications



Each column is one specification of the difference-in-differences design, ranked by its estimated effect. The upper panel plots the change in aggregate labor intensity in 2013 implied by moving to a world without regulation, in log points, computed over the municipalities in that specification's estimation sample. Whiskers are 95% confidence sets obtained by inverting a randomization test that permutes the six states' stringency paths across states, which gives 120 distinct assignments; no standard errors are computed. The lower panel marks each specification's sample, weighting, treatment of the upper tail, dynamics and form of stringency, and the red diamond marks the headline specification. The black dashed line marks zero and the red dashed line marks the observed change in labor intensity between 1999–2001 and 2013; the vertical axis is not linear in the share of that change. The axis is truncated at −2.5; 24 sets, all area weighted, extend below it (triangles), the lowest to −8.20. Seven sets, all area weighted, are not single intervals (open squares), and their whiskers span the full set. All 96 sets contain zero.

Extended Data Fig. ED1: Change in Real Wages and Hours Worked by Industry and by Occupation, 2002–2013

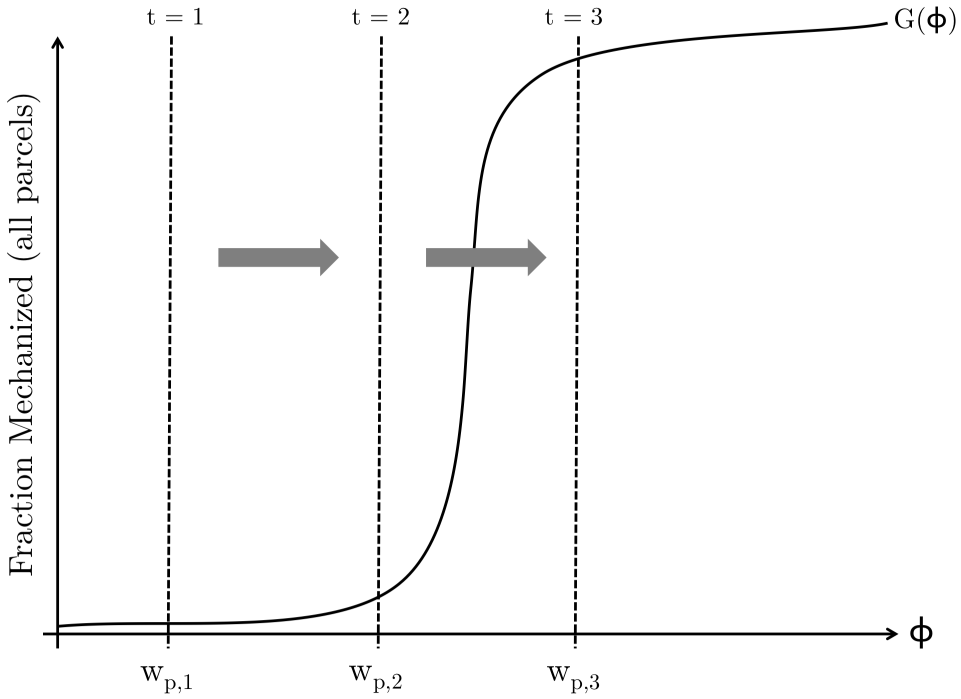


(a) By industry

(b) By occupation

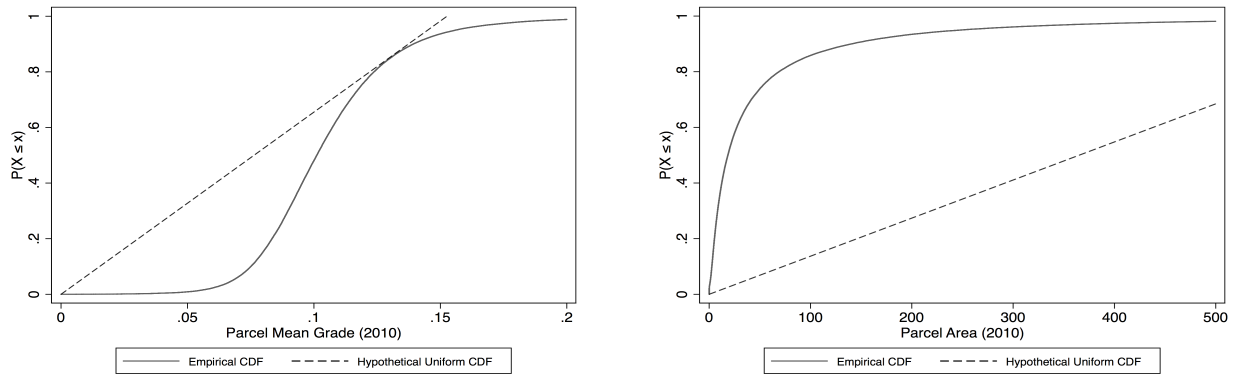
From PNAD micro data, all paid workers, hourly wages in 2003 R$. Industries are Agriculture (Ag), Industry (Ind), Manufacturing (Man), Construction (Con), Trade and repair (T/R), Hospitality (Hosp), Transport, communication and storage (Trans), Public administration (Pub), Education, health and social services (Ed/Med), Domestic services (Dom), Other services (Serv), Other (Oth). 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 and repair workers (Prod), Military (Mil). Every industry and occupation experienced meaningful wage growth. Hours worked increased in every industry except agriculture and domestic services, and in every occupation except agricultural workers. Changes to PNAD industry and occupation codes prohibit easy comparisons to earlier years.

Extended Data Fig. ED2: Harvesting Techniques and the Distribution of Threshold Wages



$G(\phi)$ is the cumulative distribution of the threshold wage ϕ across parcels. Parcels with ϕ below the prevailing manual wage w_p mechanize. Equal wage increases have very different effects depending on where they fall in the distribution.

Extended Data Fig. ED3: Empirical Distributions of Parcel Steepness and Parcel Area



(a) Cumulative distribution of parcel steepness (b) Cumulative distribution of parcel area

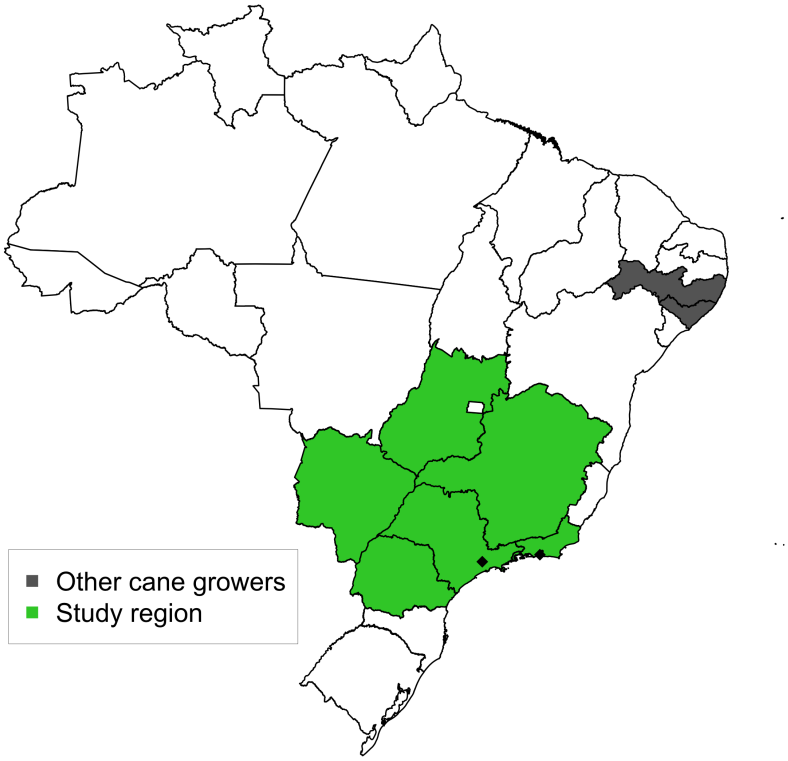
Steepness and area are the two land characteristics that most directly govern the threshold wage in the model. Both cumulative distributions have the S shape that, in the model, converts steady wage growth into abrupt aggregate mechanization. These are distributions of land characteristics, not of threshold wages, which are not directly observed.

Extended Data Fig. ED4: Quadratic and Linear IV Results



Predicted quantity of labor over the observed range of log real median wages, from the linear specification in Column (5) and the quadratic specification in Column (6) of Table 1. The quadratic specification implies a smaller elasticity at lower wages and a larger elasticity at higher wages, the curvature the model predicts.

Extended Data Fig. ED5: Major Sugarcane Producing States and the Study Region



The study region comprises the six contiguous South-Central states of Goiás, Minas Gerais, Paraná, Mato Grosso do Sul, Rio de Janeiro, and São Paulo, which together account for about 80 percent of Brazil's sugarcane output.

Extended Data Fig. ED6: Required Mechanization as a Fraction of Land Area

	São Paulo	Rio de Janeiro	Goiás	Mato Grosso do Sul	Minas Gerais	Paraná
1999						
2000						
2001						
2002	0.20					
2003	0.20					
2004	0.20					
2005	0.20					
2006	0.30					
2007	0.50					
2008	0.50		0.10		0.80	
2009	0.50		0.10		0.80	
2010	0.70		0.10		0.80	
2011	0.70		0.10	0.17	0.80	
2012	0.70	0.20	0.10	0.34	0.80	
2013	0.70	0.20	0.25	0.50	0.80	
2014	1.00	0.50	0.25	0.67	1.00	
2015	1.00	0.50	0.25	0.84	1.00	0.20
2016	1.00	0.50	0.25	1.00	1.00	0.20
2017	1.00	0.50	0.25	1.00	1.00	0.20
2018	1.00	0.80	0.50	1.00	1.00	0.20
2019	1.00	0.80	0.50	1.00	1.00	0.20
2020	1.00	1.00	0.50	1.00	1.00	0.60
2021	1.00	1.00	0.50	1.00	1.00	0.60
2022	1.00	1.00	0.50	1.00	1.00	0.60
2023	1.00	1.00	0.75	1.00	1.00	0.60
2024	1.00	1.00	0.75	1.00	1.00	0.60
2025	1.00	1.00	0.75	1.00	1.00	1.00
2026	1.00	1.00	0.75	1.00	1.00	1.00
2027	1.00	1.00	0.75	1.00	1.00	1.00
2028	1.00	1.00	1.00	1.00	1.00	1.00

A value of zero (red) means no restriction on burning; a value of one (green) means pre-harvest burning is prohibited on the entire property. These values are the treatment variable Pct in the difference-in-differences design.

Table ED1: Heterogeneity-Robust Estimates of the Effect of Stringency on Log Labor Intensity

Panel A. Static estimators			
	Estimate	Switchers	Stayers
Two-way fixed effects	-0.064	7,710 municipality-years	
Weighted average slope	-0.190	904	2,166
Average slope	0.107	904	2,166
Placebo 1, weighted average slope	-0.557	841	936
Placebo 2, weighted average slope	0.026	816	1,059
Placebo 1, average slope	-0.407	841	936
Placebo 2, average slope	0.109	816	1,059
Panel B. Event-study estimator, normalized			
	All comparisons	Never treated	Switchers
Effect 1, year of first change	-0.398	-0.508	242
Effect 2, 1 year after	-0.114	0.011	233
Effect 3, 2 years after	-0.033	0.164	224
Effect 4, 3 years after	0.134	0.302	224
Effect 5, 4 years after	0.162	0.266	215
Average total effect per unit	0.111	0.458	
Placebo 1	-0.426	-0.326	232
Placebo 2	-0.669	-0.516	219
Placebo 3	1.211	1.211	9

Point estimates of the change in log labor intensity per unit of stringency, on the headline sample (municipalities harvesting sugarcane in at least one year of 1999–2001; unweighted; unwinsorized log outcome). Panel A uses did_multiplegt_stat. Switchers are municipality-years in which stringency changes from the previous year and stayers are those in which it does not; the estimators use only these consecutive-year comparisons, so they draw on fewer observations than the two-way fixed effects estimate. The weighted average slope weights each switcher by the size of its change in stringency. Placebos repeat the comparison over earlier consecutive years. Panel B uses did_multiplegt_dyn with normalized effects: each effect is divided by the cumulative change in stringency received. The first column compares switchers with municipalities whose stringency has not yet changed, including those in states treated later; the second uses only the three never-treated states. Switcher counts are the same in both columns. Both commands report standard errors, confidence intervals and tests computed by clustering on the six states. They rest on the approximation that Methods rejects for this design and are not reported here.