EUROPEAN SUMMER SYMPOSIUM IN INTERNATIONAL MACROECONOMICS (ESSIM) 2013

Hosted by the
Central Bank of the Republic of Turkey (CBRT)

Izmir, Turkey; 21-24 May 2013

AGRICULTURAL PRODUCTIVITY AND STRUCTURAL TRANSFORMATION. EVIDENCE FROM BRAZIL

Paula Bustos, Bruno Caprettini and Jacopo Ponticelli

Agricultural Productivity and Structural Transformation. Evidence from Brazil *

Paula Bustos Bruno Caprettini Jacopo Ponticelli[†]

May 7, 2013

Preliminary and Incomplete

Abstract

We study the effects of agricultural productivity on industrial development. Classical models of structural transformation propose several channels through which productivity growth in agriculture can speed up industrial growth. However, Matsuyama (1992) notes that in open economies a comparative advantage in agriculture can slow down industrial growth. In this paper we provide direct empirical evidence on the impact of agricultural productivity on industrial development by studying the effects of adoption of new agricultural technologies in Brazil. In particular, we use the widespread adoption of genetically engineered soybean seeds as a shock to the productivity of Brazilian agriculture and measure its effects on manufacturing firms. To establish causality, we exploit exogenous differences in soil and weather characteristics across geographical areas leading to a differential impact of the new technology on yields. We find that areas more affected by technical change in soy production experienced faster manufacturing growth. Our findings imply that if technical change is strongly labour saving, increases in agricultural productivity can lead to industrialization even in an open economy

*We received valuable comments from David Atkin, Gino Gancia, Gene Grossman, Jaume Ventura, Silvana Tenreyro, Joao Pinho de Melo, Claudio Ferraz and participants at presentations held at Pontificia Universidade Catolica de Rio de Janeiro, Fundaçao Getulio Vargas, CREI, UPF, LSE and the SED annual meeting. We acknowledge financial support from the Private Enterprise Development in Low-Income Countries Project by the CEPR and UK Department for international development.

†CREI, Universitat Pompeu Fabra and Barcelona GSE (www.crei.cat).

1 Introduction

The early development literature documented that the growth path of most advanced economies was accompanied by a process of structural transformation. As economies develop, the share of agriculture in employment falls and workers migrate to cities to find employment in the industrial and service sectors (Kuznets, 1957). These findings suggest that isolating the forces that can give rise to structural transformation is key to our understanding of the development process. In particular, scholars have argued that increases in agricultural productivity are an essential condition for economic development, based on the experience of England during the industrial revolution (Nurkse, 1953 and Rostow, 1960). Classical models of structural transformation formalized their ideas by noting that productivity growth in agriculture increases income per capita, generating demand for manufacturing goods (Murphy et al., 1989, Matsuyama, 1992, Gollin et al., 2002). However, Matsuyama (1992) notes that the positive effects of agricultural productivity on industrialization occur only in closed economies, while in open economies a comparative advantage in agriculture can slow down industrial growth. This is because labor reallocates towards the agricultural sector, reducing the size of the industrial sector and its scope to benefit from external scale economies.

The sharp prediction that increases in agricultural productivity do not lead to industrial growth in open economies needs to be qualified in two dimensions. First, in the presence of transport costs and increasing returns to scale, local demand can shape specialization through home market effects even in open economies (Krugman, 1980). Second, note that classical models of the effects of agricultural productivity on industrial development such as Matsuyama (1992) assume that technical change is hicks-neutral. Thus, they predict that in an open economy an increase in agricultural productivity would induce increases in agricultural employment and a reduction in the size of the industrial sector as labor reallocates towards agriculture and wages increase. However, technical change in agriculture is often labor-saving. In this case, increases in agricultural productivity can lead to industrialization even in an open economy. This is because if the reduction in the labor intensity of agricultural production is strong enough, employment in agriculture and wages might fall, inducing an expansion of industrial employment. Then, whether technological change in agriculture induces industrial development in an open economy depends on the factor bias of technical change.

In this paper we shed light on the effects of factor biased technical change in agriculture by studying the widespread adoption of new agricultural technologies on Brazilian manufacturing firms. In particular, we study the effects of the adoption of a new agricultural technology, namely genetically engineered (GE) soybean seeds, on industrial development in Brazil. During the ten years after the technology was invented in 1996 the output of soy doubled in Brazil, becoming the most important crop in the country. GE soybeans seeds have a gene that makes them herbicide tolerant: a feature that allows farmers to adopt a new package of techniques that lowers production costs, in particular it requires less labor per unit of land to yield the same output. The adoption of GE soyben seeds can thus be characterized as labor saving technical change.

To identify the causal effects of this new technology, we use two sources of exogenous variation in the profitability of technology adoption. First, as the technology was invented in 1996, we use the introduction of the new technology as our source of time variation. Second, as the new technology had a differential impact on yields depending on geographical and weather characteristics, we use differences in soil suitability across regions as our source of cross-sectional variation.

We start by reporting that during this period the municipalities where soy expanded experienced an increase in agricultural output per worker, a reduction in labor intensity in agriculture and an expansion in industrial employment. This could respond to the adoption of labor saving agricultural technologies reducing labor demand in the agricultural sector and thus inducing a reallocation of labor towards the industrial sector. Alternatively it could be due to other shocks to local labor markets. For example: an increase in labor demand in the industrial sectors could increase wages, inducing agricultural firms to switch to less labor intensive crops, like soy. To establish the direction of causality we exploit the timing of adoption and the differential impact of the new technology on potential yields across geographical areas.

We construct a municipality-level measure of the potential profitability of technology adoption using data on potential soil yields from the FAO-GAEZ database. These yields are calculated by incorporating local soil and weather characteristics into a model that predicts yields for each crop given certain climate and soil conditions. Potential yields are a source of exogenous variation in agricultural productivity because they are a function of weather and soil characteristics, not of actual yields in Brazil. In addition, the database reports potential yields under different technologies. Thus, we exploit the predicted differential impact of the high technology on yields across geographical areas in Brazil as our source of cross-sectional variation in agricultural productivity. This research design allows us to investigate whether exogenous shocks to local agricultural productivity lead to changes in the size and efficiency of the local industrial sectors. Note that this identification strategy relies on the assumption that although Brazil is an open economy, the existence of transport costs implies that local markets are important sources of labor for industrial firms.

In a preliminary analysis of the data, we find that municipalities where the new technology is predicted to have a higher effect on potential yields of soy did experience a higher increase in the area planted with GE soy. These preliminary findings show that our instrument (the potential impact of new agricultural technologies on yields given soil and weather characteristics) is a good predictor of the profitability of GE soy adoption. In addition, these regions experienced increases in the value of agricultural output per worker and reductions in labor intensity measured as employment per hectare. Finally, these regions experienced faster employment growth and wage reductions in the industrial sector. Interestingly, the effects of technology adoption are very different for maize, a labor intensive crop that also experienced technical change during this periods as new seeds were introduced by the Brazilian government's agency of agricultural technology development (*EMBRAPA*). Regions where the FAO potential yields are predicted to increase the most when switching from the low to the high technology, did indeed experience a higher increase in the area planted with maize. In addition, they experienced increases in the value of agricultural output per worker and labor intensity. Finally, they experienced increases in wages in the industrial sector.

The differential effects of technological change in agriculture documented for GE soy and maize indicate that the factor bias of technical change is a key factor in the relationship between agricultural productivity and industrial growth in open economies. Our purpose in this paper is to investigate these effects further to isolate the channels and mechanisms through which structural transformation takes place.

The remaining of the paper is organized as follows. Section 2 gives background information on agriculture in Brazil. Section 3 describes the data used in the empirical analysis. Section 4 presents the empirical strategy and results. Section 5 concludes.

2 Agriculture in Brazil

In this section we provide background information about recent developments in the agricultural sector in Brazil. In the last decade, Brazilian labor force has been shifting away from agriculture. In 2004, industry and agriculture employed 18 millions workers each. By 2011, industry was employing 5 millions workers more than agriculture (see Figure 1). In the same years of this process of structural transformation, Brazilian farmers started introducing on a large scale the most relevant innovation in agricultural technology since the green revolution: genetically engineered (GE) crops.

Table 1 shows how the area of farms specialized in seasonal crops – i.e. crops produced from plants that need to be replanted after each harvest – increased by 11.8 millions hectares between 1996 and 2006. Out of these 11.8 millions, 8.3 millions hectares were part of farms specialized in soybean production. Figure 2 shows the time variation of the actual area planted with soy in Brazil since 1980. This area doubled in the last 10 years, with a clear breaking point in the early 2000s, in correspondence with the approval of GE soybean seeds for commercial use.

GE soy, in its Roundup Ready (RR) variety, was patented in 1996 by the multinational agricultural biotechnology corporation Monsanto. It was tested in Brazil by Monsanto since 1998, then temporarily legalized in 2003, and definitely approved for commercialization in 2005.[1] The adoption of GE soy by Brazilian farmers has been fast. According to the last Agricultural Census, in 2006 already 46.4% of Brazilian farmers were using GE soy (IBGE, 2006, p.144). The Foreign Agricultural Service of the USDA reports that GE soybean seeds were used in 85% of the area planted with soy in the 2011/2012 crop year (USDA, 2012).

Brazilian farmers opted for GE soy with the objective of reducing production costs. GE soybeans seeds have a gene that makes them herbicide tolerant: a feature that allows farmers to adopt a new "package" of techniques that lowers labor intensity for several reasons. First, since GE soybeans are resistant to herbicides, weeds control can be done more flexibly, because the chemical can be applied at any time during the season, also after the emergence of the plant (Duffy and Smith, 2001). Second, the new GE soybeans are resistant to a specific herbicide (glyphosate), which needs fewer applications: on average fields cultivated with GE soybeans require 1.55 sprayer trips against 2.45 of conventional soybeans (Duffy and Smith, 2001; Fernandez-Cornejo et al., 2002). Third, herbicide resistant soybean ease the adoption of no-tillage techniques, a system that substitute plowing with the application of chemical products that supplement the ground with the nutrients removed by the previous season's crop. The system reduces labor requirement, because application of chemicals needs fewer and shorter trips than plowing, and because no-tillage allows greater density of the crop on the field (Huggins and Reganold, 2008). Finally, farmers that adopt the new GE soybeans report also gains in the time to harvest, because combine harvesters tend to clog less when herbicides are applied later in the season, after the plant has emerged (Duffy and Smith, 2001).

These gains explain why the technology spread so fast, even though experimental evidence reports no improvements in yield with respect to conventional soybeans (Fernandez-Cornejo and Caswell, 2006)[2].

The reduction of labor intensity of soybean cultivation is apparent also in Brazilian aggregate statistics. The first four columns in table 2 report total area and total workers employed in agriculture from the Agricultural Censi of 1996 and 2006 broken down by

[1] In 1998 the CTNBio, the National Technical Bio-Safety Committee gave permission to Monsanto to conduct field tests of the RR soy (Comunicado n. 54/1998). In 2003, law 10.688 allowed the commercialization of RR soy for one season, requiring users to burn all unsold stocks after the harvest. Finally, in 2005 law 11.105 – the New Bio-Safety Law – authorized production and commercialization of RR soy (art. 35). Since Monsanto holds the exclusive rights to the commercialization of RR soy, Brazilian farmers pay royalties for its use after each harvest.

[2] In particular, in the U.S. the main advantage enjoyed by farmers that adopted GE soybeans is not in terms of higher yield, but in terms of time saved: in 2000 off-farm income of GE soybean adopters was 28% higher than that of non-adopters, while farm income was roughly the same (Fernandez-Cornejo et al., 2005).

farms' principal activity. Column 5 and 6 report the number of workers per thousand hectares across the different activities in the same two years, and the last column report the relative change. From the table it is clear that soybean production experienced the largest reduction in labor intensity, with a drop of workers per unit of land of -37%. By contrast, labor intensity in other activities either did not change (as permanent crop cultivation) or increased (as in cattle ranching and in forestry).

Table 2 also shows another important point. Soybean production is the least labor intensive agricultural activity, employing on average 22 workers per thousand hectare against 153 of other seasonal crops, 27 of cattle ranching and 39 of forestry.

Tables 3 shows correlations between the expansion of soy and the movement of other agricultural activities. It is constructed from Census data at the municipality level by dividing all Brazilian municipalities into three groups: the first (in column 1) with municipalities where the total area reaped with soybean increased between 1996 and 2006, the second (in column 2) with municipalities where area reaped with soybean decreased between 1996 and 2006 and the third (column 3) with those municipalities that did not produce soybean neither in 1996 nor in 2006. Panel A of table 3 reports the average change in the area of farms whose principal activity is the cultivation of seasonal crops (except soybean), permanent crops, cattle ranching and forestry in each of these 3 groups of municipalities. The last 2 columns of panel A report the difference between the averages in column 1 and 2 and the p-values from a test of equality of the means of these two groups. Panel B of table 3 shows the total changes in these areas.

Table 3 shows that soy expanded primarily on areas previously occupied by cattle ranching. Moreover, it also shows that in municipalities where soy expanded, other seasonal crops expanded at the same time. Table 4 looks deeper into the changes of other seasonal crops, and report the average change in the area reaped with the main crops in Brazil across the three groups of municipalities (those that expanded soybean cultivation, those that retrenched and those that do not cultivate soybean neither in 1996 nor in 2006). From table 4 it is clear that soybean cultivation expanded together with maize and (to a lesser extent) with wheat cultivation, while it does not show a strong correlation with the other main seasonal crops. Overall, the evidence shown in tables 3 and 4 is consistent with anecdotal evidence, that reports soy cultivation expanding on areas previously devoted to cattle ranching, and soybean and maize being cultivated on the same fields during different seasons of the year.

In order to study the the effects of the recent technological changes on labor intensity in agriculture it is important to understand why soybean and maize expanded together. Maize used to be cultivated as soy, during the summer season (August to December in Brazil). At the beginning of the 1980s few farmers in the South-East started producing maize during the fall (from March to July), and introduced what is now know as *milho* *safrinha* ("maize from the small season"). Adoption of a second season of maize is especially convenient in places where soy is cultivated over the summer, because maize can be planted right after the harvest of soybean. Nevertheless, a second season of maize is not feasible everywhere in Brazil, because only where the rain season ends relatively late can a second season be planted: it is for this reason Southern states like Santa Caterina and Rio Grande do Sul can not produce maize during the fall.

Cultivation of a second season of maize requires the use of very modern techniques. A second season removes much nitrogen from the soil, and this needs to be re-added with fertilizers (EMBRAPA, 2006). Herbicides are also needed to remove rapidly the residuals of the previous soy crop, and to allow the farmers to plant the second season in time (if planting is done after march 15th it is less productive). Time also forces the planting of the second season to be carried out faster than the first one (one month window against two months window), which means that it is difficult to do it without modern tractors (CONAB, 2012). Finally, *milho safrinha* puts a lot of stress on the soil and EMBRAPA advises to use no-tillage techniques, as tilling twice a year would accelerate the "compactification" of soil, reducing land productivity (EMBRAPA, 2006). As already explained no-tillage techniques require the use of modern fertilizers and tractors.

Finally, even with advanced techniques, maize is still more labor intensive than both soy and other agricultural activities like cattle ranching. In the USDA Agricultural Resources Management Survey (ARMS) labor cost of maize cultivation in 2001 and 2005 were on average 1.8 and 1.4 times higher than the labor cost for soy cultivation[3].

These considerations are relevant when trying to isolate the effect of the adoption of GE soybean on labor intensity in agriculture. Between 1996 and 2006 soybean became much less labor intensive and occupied areas that were devoted to activities that were more labor intensive. For these reasons we expect the change in soybean cultivation to be negatively correlated with labor intensity in agriculture. At the same time, in many areas soybean expanded together with the cultivation of a second season of maize. Maize cultivation is an activity that is relatively more labor intensive than cattle ranching; and so when ranches are replaced by the cultivation of soybean in the summer and maize in the fall the net effect on labor intensity is ambiguous. It is for this reason that when we correlate labor intensity with the change in area devoted to soybean, we need to control for the change in the area devoted to maize. For the same reason, we also need to find an instrument for the change of area cultivated with maize when we try to identify the causal effect of the GE technology changes on agricultural and industrial outcomes.

[3]Maize (corn) survey years are 2001 and 2005, soybean producers were surveyed by the USDA in 2002 and 2006.

3 Data

In this paper we use three main data sources: the Agricultural Census for data on agriculture, the Yearly Industrial Survey (PIA) for the data on manufacturing and the FAO Global Agro-Ecological Zones database for potential yields of soy and other crops.

The Agricultural Census is released at intervals of 10 years by the IBGE, the Brazilian National Statistical Institute. We use data from the last two rounds of the census that have been carried out in 1996 and in 2006. This allows us to observe agricultural variables both before and after the introduction of genetically engineered soybean seeds, that were legalized in Brazil in 2003). The census data is collected through direct interviews with the managers of each agricultural establishment and is made available online by the IBGE aggregated at municipality level. The main variables we use from the Census are: the value of agricultural production, the number of agricultural workers and the area devoted to agriculture in each municipality. Out of the area devoted to agriculture in each municipality we are able to disentangle the area devoted to each crop in a given Census year. This allows us to monitor how land use has changed between 1996 and 2006. The upper panel of Table 1 reports the total area reaped (in millions hectares) for the three major crops by cultivated area produced in Brazil: soybean, maize and sugar. Among these three crops, soybean registered the largest absolute increase in terms of area reaped between the last two Census years. The area devoted to soybean increased from 9.2 to 15.7 millions hectares between 1996 and 2006, more than half of the total increase in all seasonal crops. The lower panel of Table 1 reports total employment (in millions workers) in seasonal crops production and in agriculture as a whole. As reported in the last row, there are around 17 millions Brazilians - around 20% of the Brazilian active population - whose main working activity is agriculture. Notice that between 1996 and 2006, although the area cultivated with seasonal crops has increased by a third (from 36.8 to 48.2 millions hectares), the number of workers employed in seasonal crops' production decreased from 6.8 to 6.4 millions. This might be due to 2 reasons: technological change in the production of single crops (within-crop effect) and the switch from more to less labor intensive crops (across-crops effect). The introduction of genetically engineered soybean, whose production requires less workers per hectare than normal soybean, is an example of a within-crop effect. The across-crop effect derives from the fact that the production of some crops is less labor intensive to start with. Data from the 1996 Agricultural Census suggests that soy production in Brazil employs on average 42 workers per thousand hectares, while maize and sugar production employ on average respectively 106 and 138 workers per thousand hectares. Part of the reduction in the number of workers employed in seasonal crops production could be due to agricultural establishment switching from more to less labor intensive crops, e.g. from maize to soy production.

Figures 1 to 3 compare the distributions of average actual yields (tons per hectare) across Brazilian municipalities in 1996 and 2006 for, respectively, soy, maize and sugar. As for soy and maize, there was a clear shift to the right in the distribution of average yields, indicating some type of technological improvement taking place. As for sugar, on the contrary, the distribution of average yields looks very similar in 1996 and 2006.

Our second source of data is the Yearly Industrial Survey (PIA), produced by the IBGE, that monitors the performance of Brazilian firms in the extractive and manufacturing sectors. We focus on the manufacturing sector as defined by the Brazilian sector classification CNAE 1.0 (sectors 15 to 37). We use yearly data from 1996 to 2007. The population of firms eligible for the survey is composed by all firms with more than 5 employees registered in the national firm registry (CEMPRE, Cadastro Central de Empresas). The survey is constructed using two strata: the first includes a representative sample of firms having between 5 and 29 employees (*estrato amostrado*), the second includes all firms having 30 or more employees (*estrato certo*). For all firms in the survey the data is available both at firm and at plant level (when firms are composed by more than one plant). Our unit of observation is, for most outcomes, the plant. At plant level the survey includes information on: number of employees, wage bill, revenues, costs, capital investment and gross value added.

Finally, we use data on potential yields for soy and other crops from the Global Agro-Ecological Zones database produced by the FAO. Potential yields are the maximum yields attainable for a crop in a certain geographical area. They depend on the climate and soil conditions of that geographical area, and the level of technology available. The FAO-GAEZ database provides estimates of potential yields under three levels of technology: low, intermediate and high. Each of these levels is captured by the availability of certain inputs like machines and fertilizers. When the level of technology is assumed to be low, agriculture is aimed at subsistence. It is mostly labor-intensive, it uses traditional cultivars and does not use nutrients or chemicals for pest and weed control. When the level of technology is assumed to be intermediate, agriculture is partly market oriented. Production is partly mechanized, it uses improved varieties and some fertilizers and chemicals for pest and weed control. When the level of technology is high, agriculture is market oriented. Production is fully mechanized, it uses improved or high yielding varieties and "optimum" application of nutrients and chemical pest, disease and weed control. The database reports potential yields for each crop under low, medium and high technological levels available in agriculture for a worldwide grid at a resolution of 9.25 x 9.25 km. Figure 4, 5 and 6 show the potential yields for soybean in Brazil under, respectively, low and high technology. The same type of maps are also available for maize and sugar. In order to match the potential yields data with agriculture and industry variables we superimposed each of the potential yields' maps with a political map of Brazil reporting the boundaries of each municipality. Then we took the average of the potential yield, weighted by the area, within each municipality. We repeated this operation per each crop and per each of the three levels of technology. Finally, we measure technological change within each municipality by computing the difference between yields under the high and the low technology. Figure 7 illustrates the resulting measure of technological change in soy.

4 Empirics

In this section we study the effect of the adoption of a new agricultural technology, genetically engineered (GE) soybean seeds. The GE soy seeds were first commercially introduced in the U.S. in 1996 and legalized in Brazil in 2003. The advantage of this seeds relative to traditional ones is that they are herbicide resistant which implies that no-tillage planting techniques can be used.[4] The planting of traditional seeds is preceded by soil preparation in the form of tillage to kill the weeds in the seedbed that would crowd out the crop or compete with it for water and nutrients. In contrast, the planting of herbicide resistant GE soy seeds requires no tillage as the herbicide kills the weeds. Then, the GE soy seeds can be applied directly on last season's crop residue. This new technology is then expected to save on production costs, in particular requires less labor per unit of land to yield the same output. Then, the adoption of GE soy seeds can be characterized as labor saving technical change.

Note that traditional models of the effects of agricultural productivity on industrial development like Matsuyama (1992) focus on hicks neutral technical change. They predict that in an open economy an increase in agricultural productivity would induce increases in agricultural employment and a reduction in the size of the industrial sector as labor reallocates towards agriculture and wages increase. The type of technical change we study is instead labour saving. As a result, new forces emerge: if the reduction in the labor intensity of agricultural production is strong enough, employment in agriculture and wages might fall, inducing an expansion of industrial employment. Then, whether technological change in agriculture induces industrial development in an open economy depends on the factor bias of technical change. In this section we exploit the adoption of GE soy seeds to asses the relative importance of these two forces.

For this purpose, we first study the effect of the adoption of GE soy on the factor intensity of agricultural production and agricultural labor markets. Next, we asses its impact on industrial employment. We start by reporting simple correlations between

[4] Genetic engineering (GE) techniques allow a precise alteration of a plant's traits and permit targeting a single plant trait. This facilitates the development of characteristics not possible through traditional plant breeding and decrease the number of unintended characteristics. In the case of herbicide resistant GE soy seeds their genes were altered to include those of a bacteria that was herbicide resistant.

the expansion of the planted area with soy relative to other crops and agricultural and industrial labor market outcomes. Next, to establish causality, we exploit the timing of legalization and the differential impact of the new technology on potential yields across geographical areas, which depends on local weather and soil characteristics.

4.1 Basic Correlations in the Data (OLS)

We start by documenting how soy and maize expansion during the 1996-2006 period relates to changes in the agricultural production and in the industrial labor market. Here we present a set of OLS regressions in which agricultural and industrial outcomes are regressed on the percentage of farm land cultivated with soy and maize, inserted one at a time in separate regressions. Note that these results are intended to introduce to the basic correlations in the data, but do not claim to uncover any causal relation between variables.

The basic form of our equation is:

$$y_{jt} = \alpha_j + \alpha_t + \beta \left(\frac{Crop\ Area}{Agricultural\ Area} \right)_{jt} + \varepsilon_{jt} \tag{1}$$

where $\left(\frac{Crop\ Area}{Agricultural\ Area} \right)_{jt}$ is total area reaped with either soy or maize divided by total farm land[5] and y_{jt} are agricultural or industrial outcomes of interest. The units of observation are Brazilian *smallest comparable areas* (*AMC*: Área Mínima Comparável in Portuguese)[6] and we wish to control for both AMC and year fixed effects (α_j and α_t). Our source for agricultural variables is the decennal Agricultural Census, which means that for both the independent variable and the agricultural outcomes we have only two observations over the last 20 years: one in 1996 and the other in 2006. With only two periods, fixed effects and first difference estimates are identical, so we estimate (1) in first differences:

$$\Delta y_j = \Delta\alpha + \beta\ \Delta \left(\frac{Crop\ Area}{Agricultural\ Area} \right)_{j} + \Delta\varepsilon_j \tag{2}$$

where changes are defined between 1996 and 2006 throughout.

Table 5 reports correlations between changes in area reaped and changes in agricultural production: the first panel presents the results using changes in area reaped with soy while the second using changes in area reaped with maize. Together they suggest that although

[5]Total farm land includes areas devoted to crop cultivation (both permanent and seasonal crops), animal breeding and logging.

[6]Municipalities are the smallest administrative areas for which many Census variables are recorded in Brazil. Over the years however, many municipalities split and reorganized to accomodate population growth and migration, with the result that municipalities are not always comparable across years. *Smallest comparable area* is the smallest geographical breakdown for which consistency is warranted overtime.

technological progress was fast in both crops, soy production became much less labor intensive and that this in turn drove labor out of agriculture.

The first column of Table 5 reports the relationship with productivity, and shows that in places where soy and maize cultivation expanded, also the value of agricultural production per worker increased. Value per worker is defined here as the total value of crop production divided by total number of workers and refers only to seasonal crop production[7]. Column 2 confirms that labor intensity in agriculture decreased where soy cultivation expanded and increased where maize cultivation did. Agricultural labor intensity is measured here as number of workers per acre in seasonal crops cultivation. These results support the notion that changes in agricultural productivity have come with major changes in the mode of production, and in particular that in 2006 soy is being produced with relatively less labor per acre and maize with relatively more. This evidence in turn, is consistent with the qualitative description of technological change we gave above, and in particular with the fact that in 2006 much of soybean was produced with GE seeds that need much less labor. Column 3 shows that the share of workers employed in agriculture decreased in places where soy expanded and did not change significatively in places cultivated with maize. Share of workers employed in agriculture is defined as total number of workers in agriculture (from the Agricultural Census) divided by total number of workers (calculated as total number of workers in non-agricultural sectors from CEMPRE plus total number of workers in agriculture from the Agricultural Census).

Table 6 reports results from regression (2) using industrial labor market outcomes as dependent variables. We focus only on manufacturing plants (CNAE 1.0 codes 15 to 37) owned by firms that employ at least 30 employees: the sample for which the PIA Empresa survey contains the population of Brazilian firms. Industrial outcomes come from the yearly plant-level survey: we aggregate these data at AMC level and then collapse all years in 2 periods: one before 2003 and one after it, in order to maximize the number of observation used in the regressions[8]. Total employment includes both production and non-production workers; plant size is calculated as total number of workers within an AMC divided by total number of plants; wage is aggregate wage bill within an AMC divided by total number of workers there.

The first column of Table 6 shows that total employment in manufacturing significantly grew where soy expanded, but not where maize did. The second and third columns qualify this result, and show that the increase came from both the intensive and the extensive margin: column 2 shows that places where soy expanded had on average more plants, while the third column shows that these plants were on average larger. Finally, the last

[7]Both soy and maize are seasonal crops.

[8] Brazilian manufacturing activity is very concentrated, and not all AMCs host at least 1 plant from a large firm every year. Collapsing observations from more than one year allows us to use all municipalities for which we observe 1 plant at least once before and after the introduction of GE soy seeds.

column of table 3 shows that wages did not move in a statistically significant way where soy expanded while they grew where maize did.

Overall, the results on employment are consistent with a rightward shift of the labor supply schedule driven by soy expansion: although the positive sign on wages does not support this story, notice that this result might be driven by omitted variable bias, because the sign is reversed when the change in soy and maize area are inserted together in the regression.

4.2 Causality

In this section we provide direct empirical evidence on the effects of the widespread adoption of a new agricultural technology, GE soybean seeds, on industrial development in Brazil. The basic correlations in the data reported in the previous section show that areas where soy expanded experienced an increase in output per worker and a reduction in labor intensity in agriculture and an expansion in industrial employment. This could respond to the adoption of labor saving agricultural technologies reducing labor demand in the agricultural sector and thus inducing a reallocation of labor towards the industrial sector. Alternatively it could be due to other shocks to local labor markets. For example: an increase in labor demand in the industrial sectors could increase wages, inducing agricultural firms to switch to less labor intensive crops, like soy. To establish the direction of causality we exploit the timing of adoption and the differential impact of the new technology on potential yields across geographical areas. First, the new technology was commercially introduced in the U.S. in 1996, and legalized in Brazil in 2003. Thus, we use the periods before and after 2003 as our "pre and post-treatment" periods. Second, the new technology has a differential impact on potential yields depending on soil and weather characteristics. Thus, we exploit these exogenous differences on potential yields across geographical areas as our source of cross-sectional variation in the intensity of the treatment.

To implement this strategy, we need an exogenous measure of potential yields for soy and other crops, which we obtain from the FAO-GAEZ database. These potential yields are estimated using an agricultural model that predicts yields for each crop given climate and soil conditions. As potential yields are a function of weather and soil characteristics, not of actual yields in Brazil, they can be used as a source of exogenous variation in agricultural productivity across geographical areas. In addition, the database reports potential yields under different technologies or input combinations. Yields under low inputs are described as those obtained using traditional seeds and no use of chemicals, while yields under high inputs are obtained using high yielding varieties and optimum application of chemicals for weed control. Thus, the difference in yields between the high and low technology captures the effect of moving from traditional agriculture to a technology that uses optimum weed control, among other characteristics. [9] We expect this increase in yields to be a good predictor of the profitability of adopting herbicide resistant GE soy seeds. Thus, we can then exploit the predicted differential impact of the high technology on yields across geographical areas in Brazil as our source of cross-sectional variation in agricultural productivity. This research design allows us to investigate whether exogenous shocks to local agricultural productivity lead to changes in the size and efficiency of the local industrial sectors. Note that this identification strategy relies on the assumption that although Brazil is an open economy, the existence of transport costs implies that local markets are important sources of labor and demand for industrial firms.

More formally, our basic empirical strategy is to estimate an equation of the following form:

$$y_{jt} = \alpha_j + \alpha_t + \beta A^{soy}_{jt} + \varepsilon_{jt} \tag{3}$$

where y_{jt} is an outcome that varies across municipalities and time, j indexes municipalities, t indexes time, α_j are municipality fixed effects, α_t are time fixed effects and A^{soy}_{jt} = potential yield of soy under high (low) inputs if $t \geqslant$2003 ($t <$ 2003).

Note that a potential problem with this identification strategy is that the productivity of land is positively correlated across crops, thus we could be capturing the effect of overall technical change instead of the labor saving technical change associated to GE soy. For example, during this period there were also increases in yields of maize, a labor intensive crop. Thus, we need to control in the above regressions for the changes in yields of maize when switching from the low to the high technology. We then include the following variable as a control: A^{maize}_{jt} = potential yield of maize under high (low) inputs if $t \geqslant$2003 ($t <$ 2003). In addition, we want to control for changes in the prices of crops, that can also have an influence on the expansion of soy relative to other agricultural activities. Note that the overall effect of price changes would

[9] FAO-GAEZ description of each technology is as follows:

Low-level inputs/traditional management

"Under the low input, traditional management assumption, the farming system is largely subsistence based and not necessarily market oriented. Production is based on the use of traditional cultivars (if improved cultivars are used, they are treated in the same way as local cultivars), labor intensive techniques, and no application of nutrients, no use of chemicals for pest and disease control and minimum conservation measures."

High-level inputs/advanced management

"Under the high input, advanced management assumption, the farming system is mainly market oriented. Commercial production is a management objective. Production is based on improved high yielding varieties, is fully mechanized with low labor intensity and uses optimum applications of nutrients and chemical pest, disease and weed control."

$$y_{jt} = \alpha_j + \alpha_t + \beta A_{jt}^{soy} + \gamma A_{jt}^{maize} + \sum_z \theta_z p_t^z A_{j0}^z + \varepsilon_{jt} \tag{4}$$

where $z =$ soy, maize and sugar, p^z is the international price of crop z, A_{j0}^z is potential yield of crop z under low inputs for maize, medium for sugar.

First we estimate the effects of the agricultural technology shock A^{soy} on the share of agricultural land devoted to soy. As the increase in yields resulting from adoption of the high technology $(A_{high}^{soy} - A_{low}^{soy})$ is expected to be correlated with the profitability of technology adoption, we can use changes in A^{soy} as an instrument for changes share of agricultural land devoted to soy. Then, we can perform an instrumental variables estimation of equation (1) where the first stage is given by equation (3). Alternatively, we can estimate reduced form equations of the form of (3) where we study the direct effect of the change in potential yields driven by technology adoption $(A_{high}^{soy} - A_{low}^{soy})$ on the set of outcomes we are interested in, namely: value of output per worker and labor intensity of agricultural production and industrial employment. In the following subsections, we report our first stage, reduced form and instrumental variable estimates.

4.3 First Stage

We document here the relation between the change in agricultural area cultivated with soy and maize on our technological shock described above. The purpose of these regressions is twofold. On the one hand we see them as a "sanity check" on the information content of our technological shock; on the other hand these regressions represent the first stage used to address the endogeneity of the OLS regressions reported above.

The equations we estimate are:

$$\Delta \left(\frac{Crop\ Area}{Agricultural\ Area} \right)_j = \Delta\alpha + \beta \Delta A_j^{soy} + \Delta\varepsilon_j$$

and

$$\Delta \left(\frac{Crop\ Area}{Agricultural\ Area} \right)_j = \Delta\alpha + \beta \Delta A_j^{soy} + \gamma \Delta A_j^{maize} + A_j^{sugar} + \Delta\varepsilon_j \tag{5}$$

where $\Delta \left(\frac{Crop\ Area}{Agricultural\ Area} \right)$ is the change in share of farm land reaped with either soy or maize between 1996 and 2006 and it is defined as in section 3.1. A_j^{soy}, A_j^{maize} and A_j^{sugar} are potential yields in AMC j for soy, maize and sugar while the technological shocks are defined as explained in the previous section: $\Delta A^z = A_{high}^z - A_{low}^z$, with $z =$soy, maize. We control for general suitability to sugar production with intermediate inputs (A_j^{sugar}) because sugar is the other major seasonal crop in Brazil. However, we do not use its change in potential yield because sugar production does not seem to be more productive in 2006 relative to 1996 (see graph 3).

Column 1 and 3 in Table 7 show that the soy and maize shock correctly predict soy and maize expansion over the period: when inserted alone, an increase in soy (maize) potential yield is associated with significantly greater farm land reaped with soy (maize). Column 2 and 4 strengthen this results: they show that controlling for both shocks and for the potential yield of sugar, increase the effect of both shock: on soy, the effect more than doubles, while on maize the effect goes up by one third. Also, the fact that the maize shock has a significantly negative effect on the share of farm land cultivated with soy and that the potential yield of sugar has a significantly negative effect on both shares, is also consistent with optimal behavior of farmers, who choose which crop to cultivate based on the specific suitability of their plots. Note that all of these suitability measures tend to be positively correlated (this is especially true for maize and soy), so the result that the shocks correctly predict expansion or retrenchment of specific crops means that they capture the dimensions that are more relevant for the decisions of farmers.

All in all, Table 7 support the choice of our instrument. The effect of the shocks are extremely significant (the F-tests for the joint significance of the regressors range from 141.8 to 15.51) and also economically relevant. The estimated coefficient on soy implies that municipalities with a one standard deviation above the mean increase in potential soy yields increased the share of soy in planted land area by 36% of a standard deviation.

4.4 Reduced Form

Once we have established the relevance of our instruments, we turn to the study of the effect of technological change in agriculture on production and employment in agriculture and manufacturing: our reduced form. In this section we show that the change in potential soy yield is associated with a reduction in the use of labor in agriculture and with a rightward shift of the labor supply schedule in manufacturing.

We start with agricultural production: table 5 reports the results of running an equation similar to (5) on the agricultural outcomes used in table 2:

$$\Delta y_j = \Delta\alpha + \beta\Delta A_j^{soy} + \gamma\Delta A_j^{maize} + A_j^{sugar} + \Delta\varepsilon_j$$

where y_j is value produced per worker in seasonal crops, labor intensity or share of workers

employed in agriculture all defined as in section 3.1.

Table 8 reinforce the results on agriculture highlighted in section 3.1: productivity increased in places where potential soy yield increased relatively more, while in the same places labor intensity in agriculture dropped and agriculture shrunk in terms of total employment. At the same time, in places where potential maize yield increased relatively more, the number of workers per acre increased, and the share of population employed in agriculture grew. Also, the maize shock seems to be negatively associated with value produced by worker in seasonal crops, but this effect is not significantly different from 0. It is interesting to note that these results do *not* hold when we run these regressions only including AMC in the North-East of the country, the region where sugar has been cultivated historically, and neither soy nor maize are widespread (results available upon request).

Overall, Table 8 is consistent with our interpretation of the technological change brought about by GE soy seed. The new technology seems to require much less labor per acre, and this in turn seems to have driven labor out of the agricultural sector.

We now turn to the regressions on manufacturing outcomes. Here we are able to exploit the full panel structure of our data: the equation we estimate with these data have the form:

$$y_{jt} = \alpha_j + \alpha_t + \beta A_{jt}^{soy} + \gamma A_{jt}^{maize} + \sum_z \theta_z P_t^z A_{j0}^z + \varepsilon_{jt}$$

where y_{jt} are industrial outcomes of interest; A_{jt}^{soy} (A_{jt}^{maize}) are potential yield of soy (maize) under low inputs for all years before 2003 and under high input for all years from 2003 on. We observe all years from 1996 to 2007 and control for the real price of soy, maize and sugar times potential yield for these crops in 1996: A_{j0}^z is potential yield under low inputs when z = soy or maize and it is potential yield under medium inputs when z = sugar. These controls are intended to make sure that changes are truly driven by technological change rather that the evolution of commodity prices. Although international commodity prices will affect all Brazilian AMC in the same way, they might still have heterogeneous effect in places that were more suitable to the cultivation of some particular crop at the beginning. Controlling for the interaction of these prices with potential yield in 1996 makes sure that our results are not driven by the heterogeneous effects that commodity prices have on different places in Brazil. In all specification we control for both AMC and year fixed effects (α_j and α_t) and cluster standard errors at AMC level to avoid that serial correlation in our shock make the precision of our estimates artificially high (Bertrand et al. (2004)).

Table 9 shows reduced form results with industrial variables. All variables are defined as above and are built as averages at AMC level starting from plant level data. Again we focus on manufacturing plants owned by firms with at least 30 employees.

The first column of Table 9 reproduces closely the patterns highlighted in section 3.1. In particular, it confirms that industrial employment (both production and non-production workers) grew in places where soy potential yield increased more. The effect of maize potential yield is negative (as in the OLS) but not significantly different from 0. Contrary to the OLS results however, the effect on employment seems to come exclusively from the growth of existing plants: in places where the change in potential soy yield increased more, average plant size grew significantly, while the total number of plant did not change. Average wage fell in places where the potential yield of soy increased more, and grew in places where the change in potential yield of maize was greater.

Overall, the results shown in Table 9 are consistent with rightward shift of the labor supply curve in areas where the change in potential yield of soy was relatively greater, and a (less clear) leftward shift of the labor supply schedule in the places where the change in potential yield of maize was relatively greater. Together with table 8, these results support our story of a labor-saving technological change in soy production promoting a growth of industrial employment by freeing up labor in agriculture.

5 Final Remarks

The process of modern economic growth has been accompanied by a reallocation of economic activity from agriculture to industry. Isolating the forces behind structural transformation is therefore key to our understanding of economic development. Based on the experience of England during the industrial revolution, economists have argued that increases in agricultural productivity can be one of these forces. However, Matsuyama (1992) notes that the positive effects of agricultural productivity on industrialization occur only in closed economies, while in open economies a comparative advantage in agriculture can slow down industrial growth. Despite the importance of the question, there is so far scarce micro evidence on the channels through which increases in agricultural productivity can shape the reallocation of economic activity across sectors in an open economy.

In this paper we contribute to the debate by isolating the effects of the adoption of new agricultural technologies on Brazilian manufacturing firms. In particular, we study the effects of the adoption of genetically engineered (GE) soybean seeds on industrial development in Brazil. To identify the causal effects of this new technology, we exploit the timing of adoption and the differential impact of the new technology on potential yields across geographical areas.

We find that municipalities where GE soy is predicted to have expanded more experienced faster employment growth and wage reductions in the industrial sector. Interestingly, we find opposite effects when looking at municipalities where maize – a labor intensive crop that also experienced technical change during this periods – is predicted to have expanded more. The different effects documented for GE soy and maize indicate that the factor bias of technical change is a key factor in the relationship between agricultural productivity and industrial growth in open economies.

References

Bertrand, M., E. Duflo, and S. Mullainathan (2004). "How Much Should We Trust Differences-in-differences Estimates?". *The Quarterly Journal of Economics 119*(1), 249–275.

Buera, F. J. and J. P. Kaboski (2009). "Can Traditional Theories of Structural Change Fit the Data?". *Journal of the European Economic Association 7*(2-3), 469–477.

CONAB (2003-2012). "Levantamento De Avaliação Da Safra, several issues".

Duffy, M. and D. Smith (2001). "Estimated Costs of Crop Production in Iowa". *Iowa State University Extension Service FM1712*.

EMBRAPA (2006). "Cultivo do Milho". Embrapa, Sistemas de Produção.

Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011). "Income Distribution, Product Quality, and International Trade". *Journal of Political Economy 119*(4), 721–765.

Fernandez-Cornejo, J. and M. Caswell (2006). "The First Decade of Genetically Engineered Crops in the United States". *United States Department of Agriculture, Economic Information Bulletin 11*.

Fernandez-Cornejo, J., C. Hendricks, and A. Mishra (2005). "Technology Adoption and Off-Farm Household Income: The Case of Herbicide-Tolerant Soybeans". *Journal of Agricultural and Applied Economics 37*(3), 549.

Fernandez-Cornejo, J., C. Klotz-Ingram, and S. Jans (2002). "Estimating Farm-Level Effects of Adopting Herbicide-Tolerant Soybeans in the USA". *Journal of Agricultural and Applied Economics 34*, 149–163.

Foster, A. D. and M. R. Rosenzweig (2007). "Economic Development and the Decline of Agricultural Employment". *Handbook of Development Economics 4*, 3051–3083.

Gollin, D., S. Parente, and R. Rogerson (2002). "The Role of Agriculture in Development". *The American Economic Review 92*(2), 160–164.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2011). "Growth and Structural Transformation". *Handbook of Economic Growth, forthcoming*.

Hornbeck, R. and P. Keskin (2012). "Does Agriculture Generate Local Economic Spillovers? Short-run and Long-run Evidence from the Ogallala Aquifer". Technical Report Working Paper 18416, National Bureau of Economic Research.

Huggins, D. R. and J. P. Reganold (2008, Month). "No-Till: the Quiet Revolution". *Scientific American*.

IBGE (2006). "Censo Agropecuário 2006".

Krugman, P. (1980). "Scale Economies, Product Differentiation, and the Pattern of Trade". *The American Economic Review 70*(5), 950–959.

Kuznets, S. (1957). "Quantitative Aspects of the Economic Growth of Nations: II. Industrial Distribution of National Product and Labor Force". *Economic Development and Cultural Change*, 1–111.

Laitner, J. (2000). "Structural Change and Economic Growth". *The Review of Economic Studies 67*(3), 545–561.

Matsuyama, K. (1992). "Agricultural Productivity, Comparative Advantage, and Economic Growth". *Journal of Economic Theory 58*(2), 317–334.

Murphy, K. M., A. Shleifer, and R. W. Vishny (1989). "Industrialization and the Big Push". *The Journal of Political Economy 97*(5), 1003–1026.

Nurkse, R. (1953). *"Problems of Capital Formation in Underdeveloped Countries"*. Oxford University Press.

Rostow, W. (1960). *"The Stages of Economic Growth: A Non Communist Manifesto"*. Cambridge University Press, London.

USDA (2012). "Agricultural Biotechnology Annual". BR0813.

Vogel, S. J. (1994). "Structural Changes in Agriculture: Production Linkages and Agricultural Demand-Led Industrialization". *Oxford Economic Papers*, 136–156.

Figures and Tables

Figure 1: Labor Force in Agriculture, Industry and Services, 2002-2011

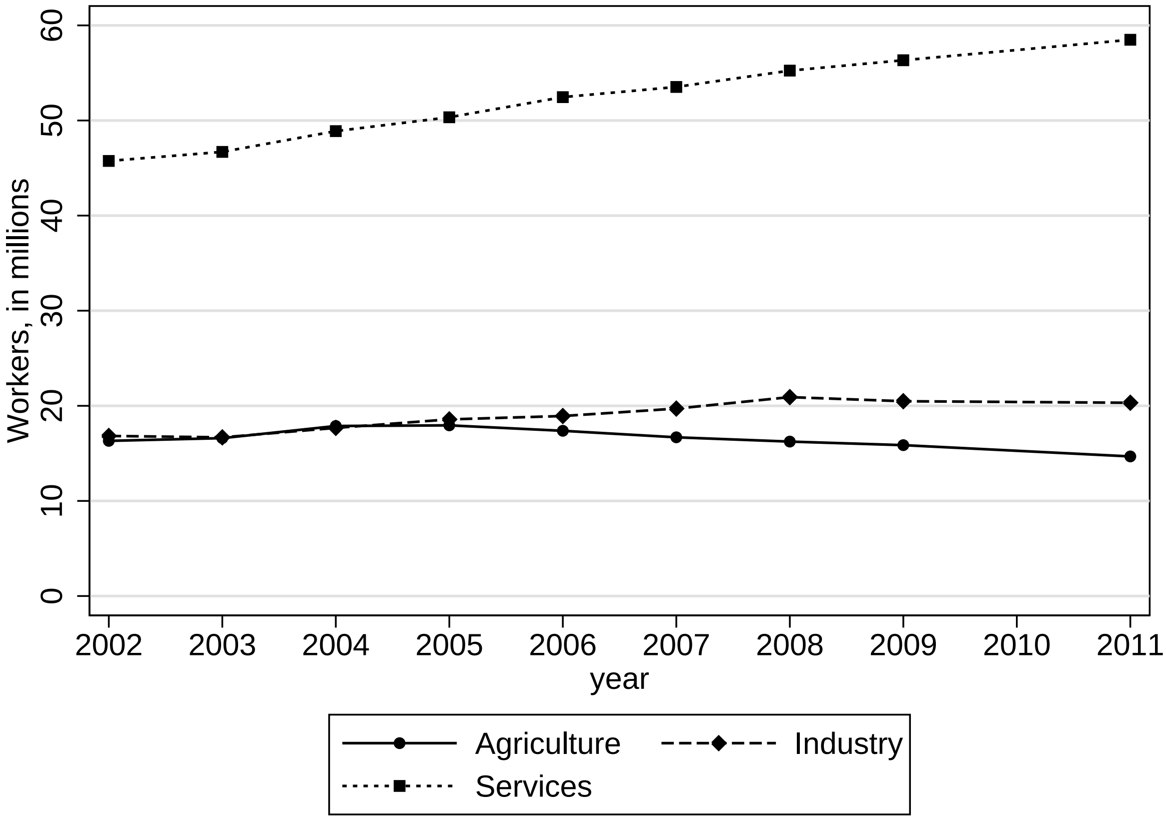

Notes: The Figure depicts the evolution between 2002 and 2011 of the total number of workers (expressed in millions) employed by sector in Brazil. The sectors are: Agriculture (including pasture, fishery and forestry), Industry (including manufacturing, construction and extractive industries) and Services. Data come from PNAD, a national household survey representative at country level and carried out yearly (with the exception of the population census years) by the Brazilian National Statistical Institute. The sectoral classification used is the CNAE-Domiciliar.

Figure 2: Area Planted with Soy, 1980-2010

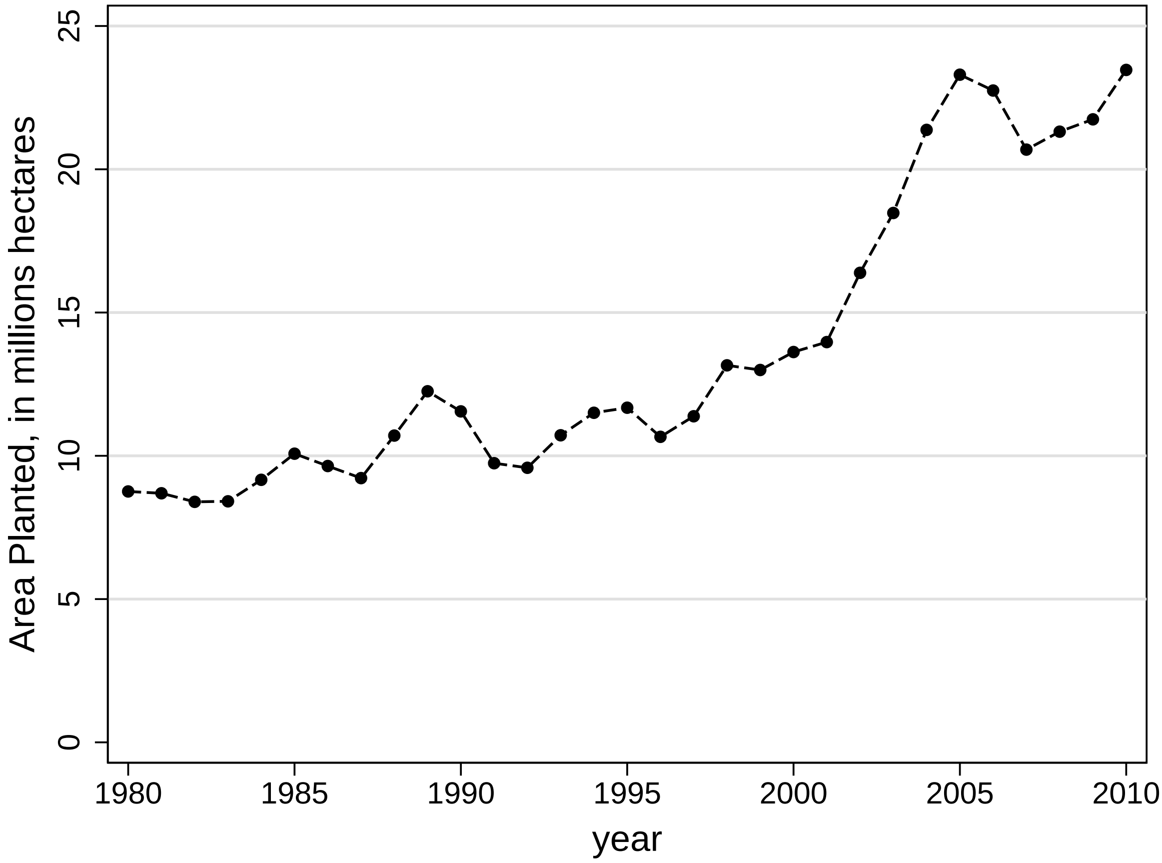

Notes: The Figure depicts the evolution between 1980 and 2010 of the total area planted with soy in Brazil (expressed in millions hectares). Data come from monthly surveys carried out by CONAB, Companhia Nacional de Abastecimento, an agency created by the Brazilian Ministry of Agriculture. Data is constructed by interviewing on the ground farmers, agronomists and financial agents in the main cities of the country.

Figure 3: Labor Force in Soy Production, 2002-2011

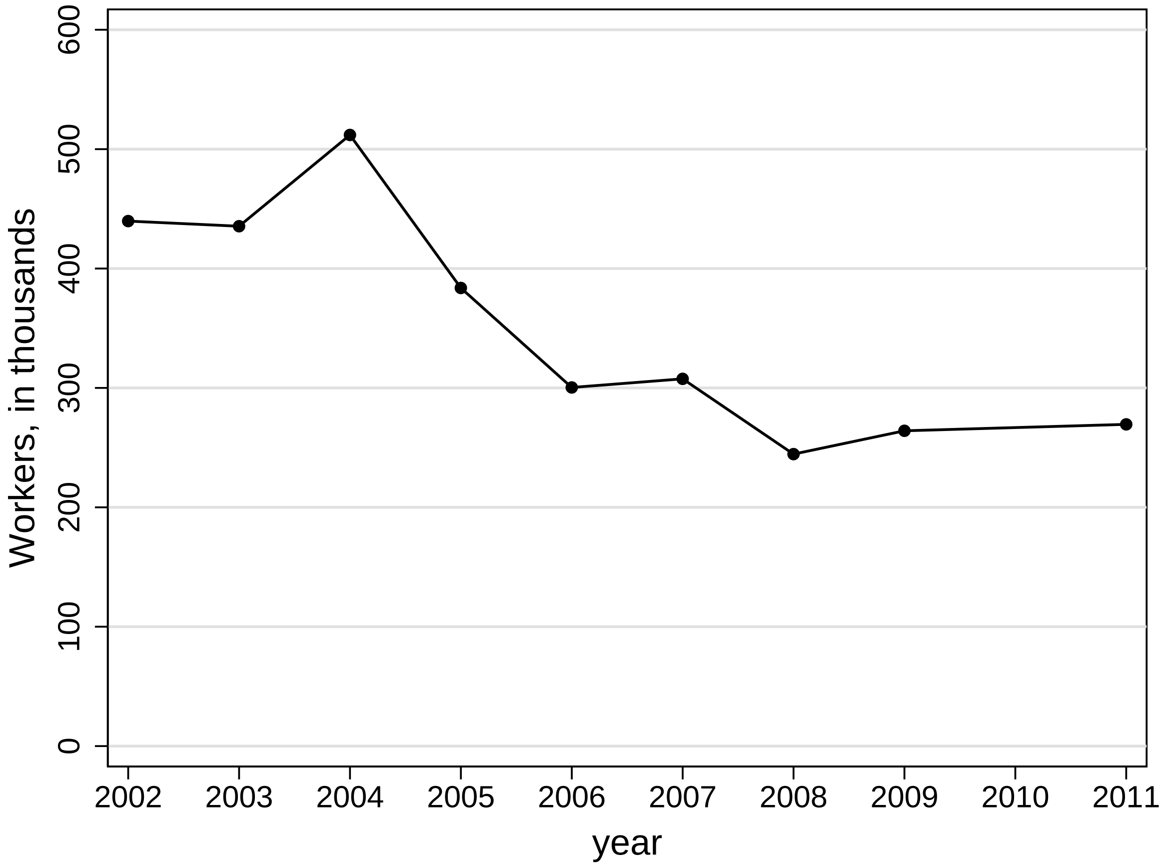

Notes: The Figure depicts the evolution between 2002 and 2011 of the total number of workers employed in soy production (expressed in thousands) in Brazil. Data come from PNAD, a national household survey representative at country level and carried out yearly (with the exception of the population census years) by the Brazilian National Statistical Institute. The sectoral classification used is the CNAE-Domiciliar (soy production has the code 01107).

Figure 4: Distribution of Actual Soy Yields across Brazilian Municipalities in 1996 and 2006

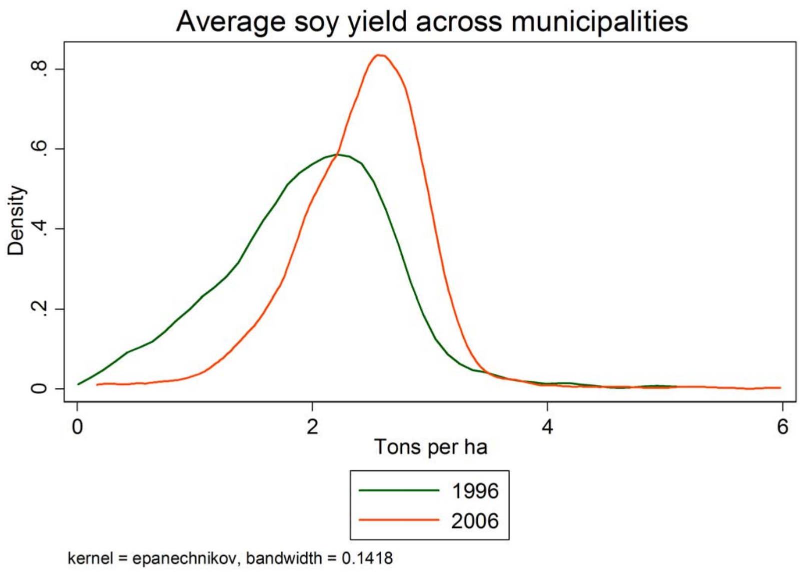

Notes: All data from Brazilian Agricultural Censi of 1996 and 2006, IBGE.

Figure 5: Distribution of Actual Maize Yields across Brazilian Municipalities in 1996 and 2006

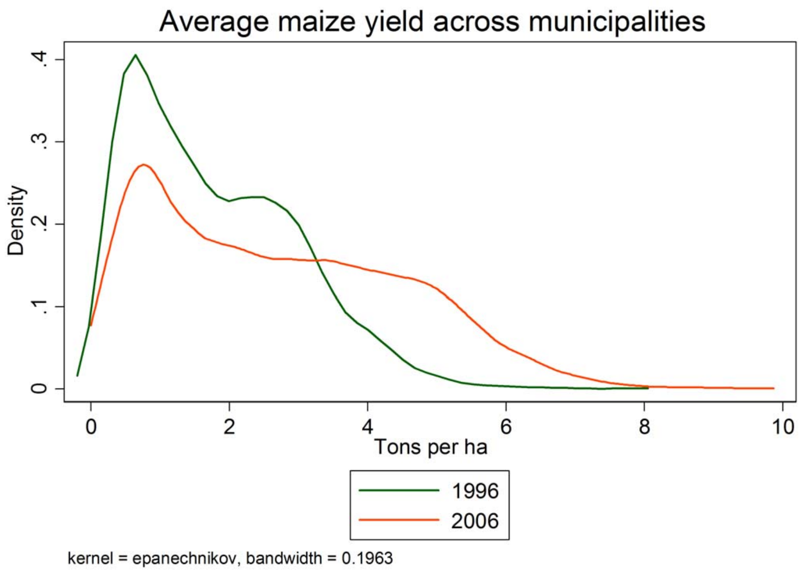

Notes: All data from Brazilian Agricultural Censi of 1996 and 2006, IBGE.

Figure 6: Distribution of Actual Sugar Yields across Brazilian Municipalities in 1996 and 2006

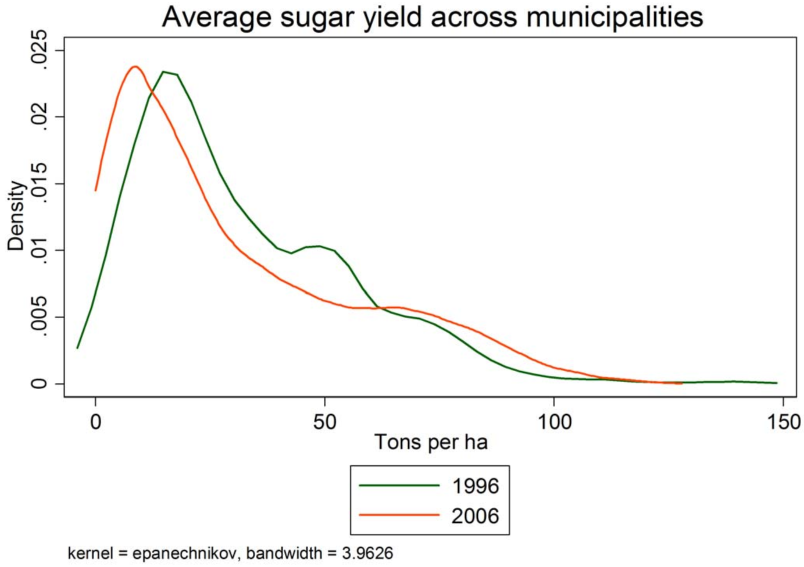

Notes: All data from Brazilian Agricultural Censi of 1996 and 2006, IBGE.

Figure 7: Potential Soy Yield Under Low Agricultural Technology

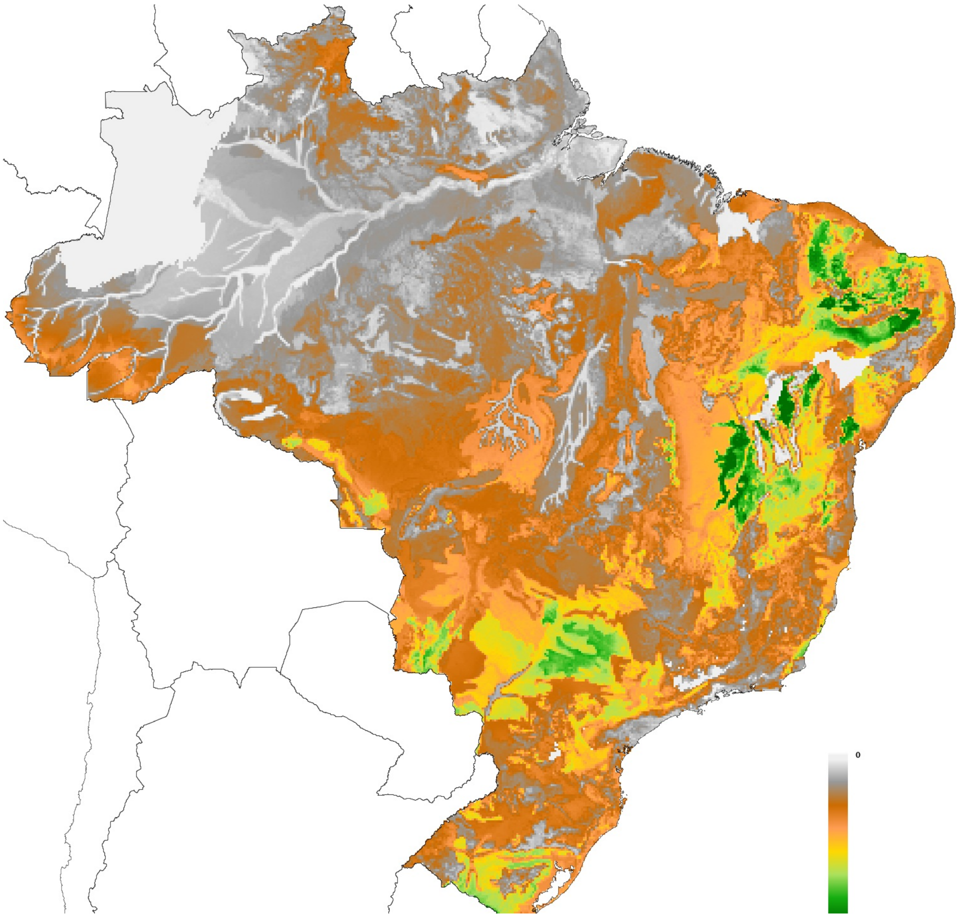

Notes: All data from FAO GAEZ.

Figure 8: Potential Soy Yield Under Intermediate Agricultural Technology

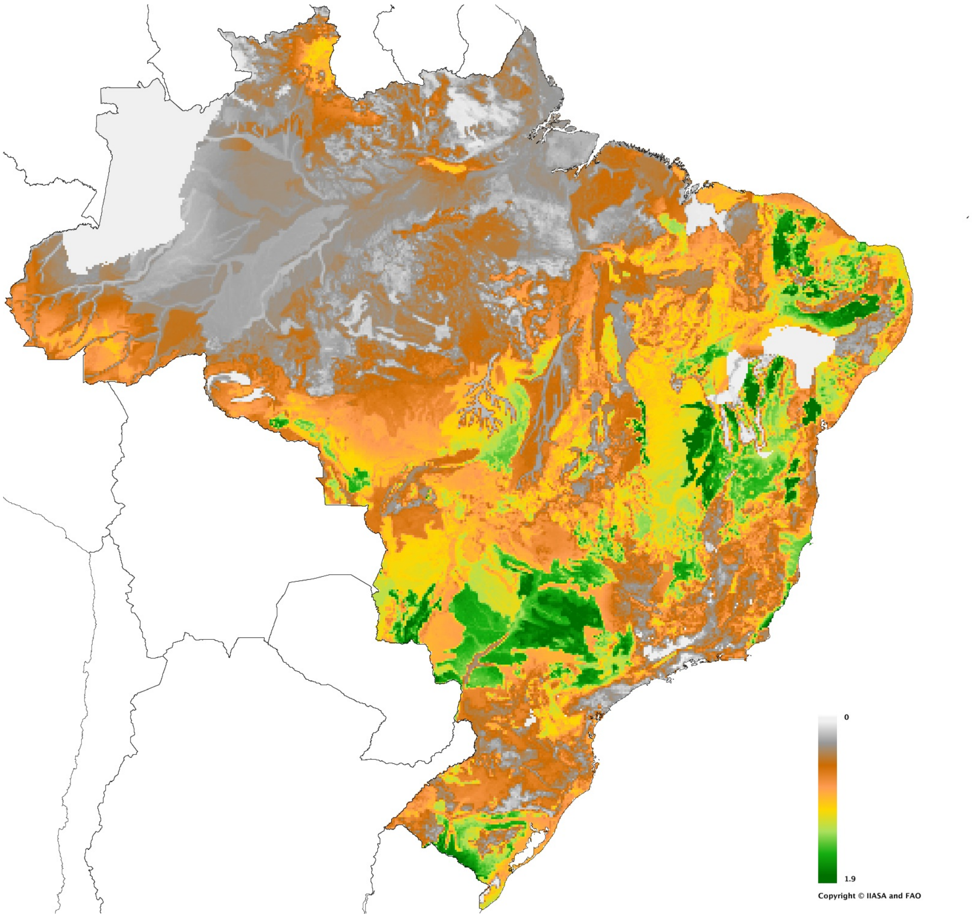

Notes: All data from FAO GAEZ.

Figure 9: Potential Soy Yield Under High Agricultural Technology

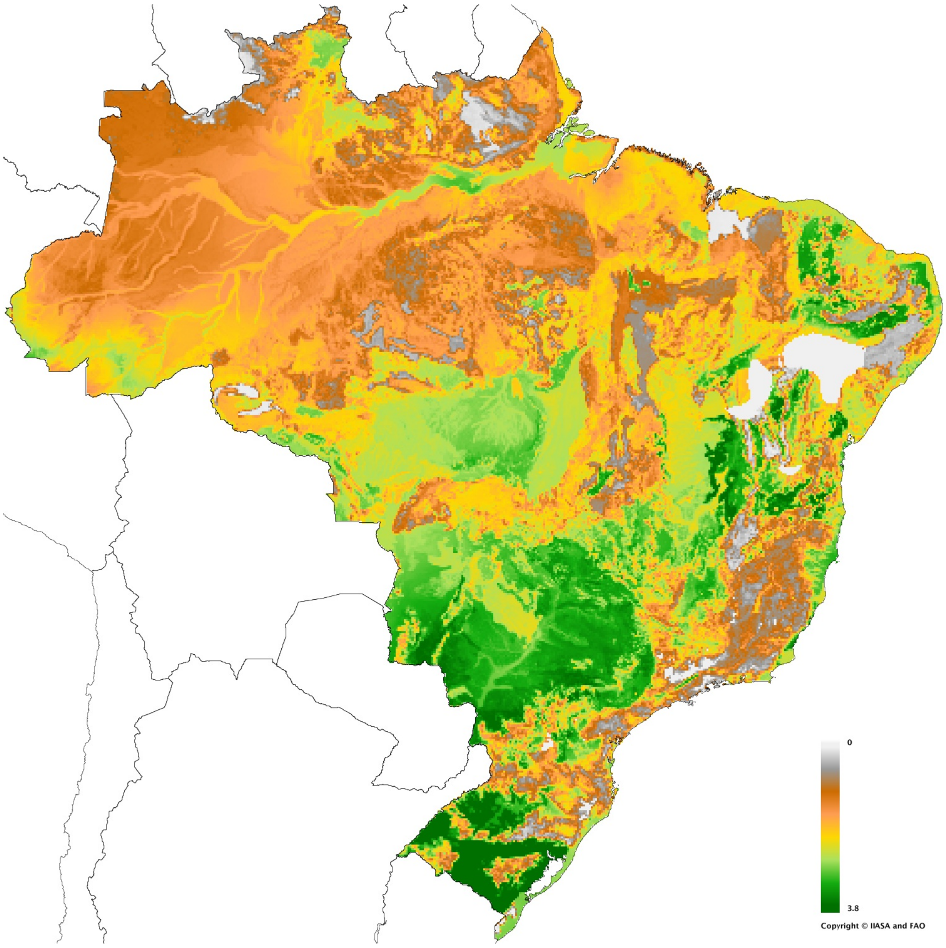

Notes: All data from FAO GAEZ.

Figure 10: Technological Change in Soy: Potential Yield under High Technology minus Potential Yield under Low Technology

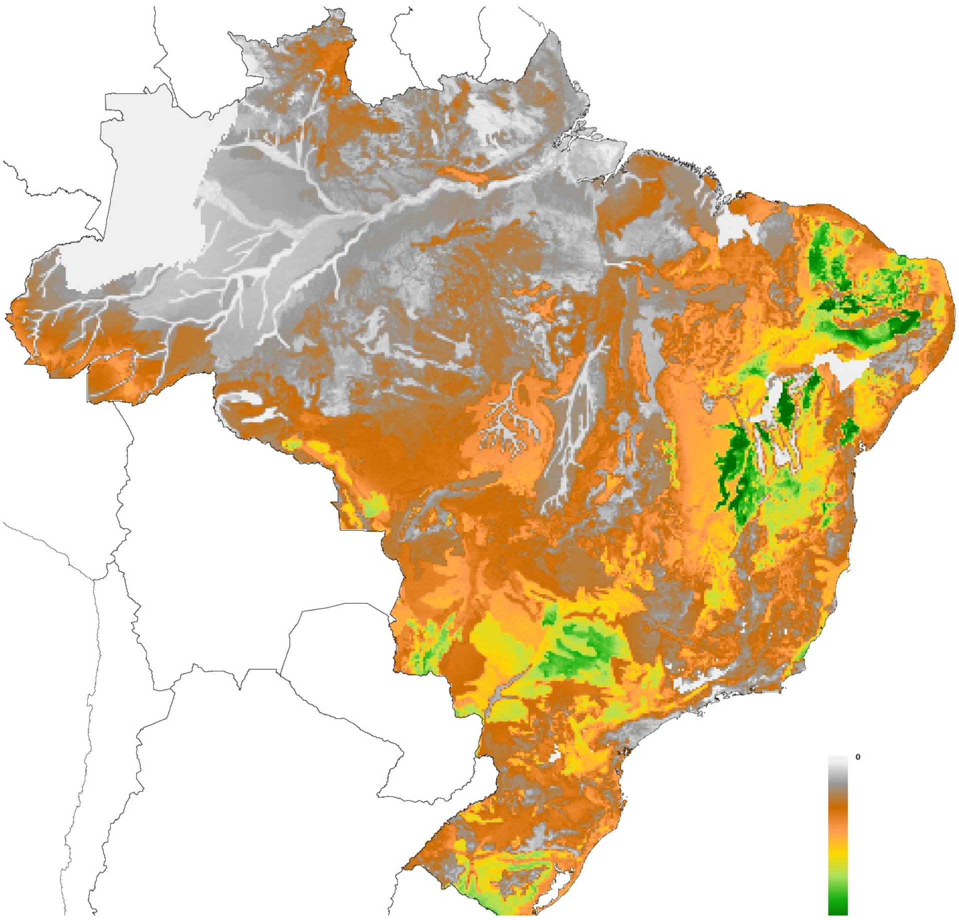

Notes: Authors' calculations starting from FAO GAEZ data.

Table 1: Land use (millions ha)

	1996	2006	Change	% change
Permanent crops	7.5	11.7	4.1	55%
Seasonal crops	34.3	44.6	10.4	30%
Cattle ranching	177.7	168.3	-9.4	-5%
Forest	110.7	91.4	-19.2	-17%
Not usable	15.2	8.2	-6.9	-46%
Other	8.3	9.0	0.7	8%
Total	353.6	333.2	-20.4	-6%

Notes: The Table reports the total land use in Brazil (expressed in millions hectares). Data is available for 1996 and 2006 and come from the last two Brazilian Agricultural Censi carried out by the Brazilian National Statistical Institute and it is sourced from the the IBGE Sidra repository (table 317 for 1996 and table 1011 for 2006). Seasonal crops include (among others) cereals (e.g. maize, wheat and rice), soybean, cotton, sugar cane and tobacco. Permanent crops include (among others) coffee and cocoa. Not usable land includes lakes and areas that are not suitable for neither crop cultivation nor cattle ranching. Other uses is not exactly comparable across years: in 1996 it includes resting area for seasonal crops; in 2006 it includes area devoted to pasture, flowers and buildings.

Table 2: Factor Intensity in Brazilian Agriculture: 1996-2006

	Area		Workers		Labour intensity			Change in labour intensity	
Principal activity:	1996	2006	1996	2006	1996	2006	Average 1996-2006	Absolute	Relative
Seasonal crops	63.0	74.8	6.8	6.3	107.6	83.7	94.6	-23.9	-22%
soy	*15.1*	*23.4*	*0.4*	*0.4*	*28.6*	*17.9*	*22.1*	*-10.7*	*-37%*
all cereals	*19.1*	*23.8*	*1.8*	*1.8*	*92.4*	*76.8*	*83.7*	*-15.6*	*-17%*
other	*28.8*	*27.6*	*4.6*	*4.0*	*159.2*	*145.4*	*152.5*	*-13.8*	*-9%*
Permanent crops	17.2	17.3	2.2	2.2	126.8	127.4	127.1	0.6	0%
Cattle ranching	213.9	221.8	4.8	6.8	22.6	30.6	26.7	8.1	36%
Forest	16.0	13.0	0.5	0.6	33.9	46.1	39.4	12.2	36%

Note: The table reports total land in farms and total number of workers employed by principal activity of the farm. Data are sourced from the IBGE Sidra repository. Land in farm by principal activity in 1996 comes from table 491 and for 2006 from table 797. Total number of workers in 1996 is reported in table 321 and in 2006 in table 956. Cereals are rice, wheat, maize and other cereals. Labor intensity is computed as number of workers per 1000 hectares. The definition of "principal activity" of the farm changed somehow between 1996 and 2006. In 1996 higher specialization was required for farms to be classified under one of the categories reported, and those that did not produce at least 2/3 of the value within a single category were classified under the "mixed activity" category. In 2006 farms were classified according to the activity that accounted for the simple majority of production and no "mixed activity" category existed. Source: Agricultural Census 1996 and 2006.

Table 3: Correlations between changes in land by activity: 1996-2006

Panel A. Average change in land used (ha)

	Land reaped with soy				
	increased (1)	decreased (2)	no soy (3)	Difference (1) - (2)	p-value
Soy	6’912	-1’295		8’206	0.00***
Seasonal crops, except soy (area reaped)	2’188	335	703	1’853	0.00***
Permanent crops	1’054	919	629	135	0.48
Cattle ranching	-8’326	-3’882	760	-4’443	0.02**
Forest	-157	-8’911	-3’165	8’754	0.00***
Unusable land	-2’516	-1’882	-727	-633.9367	0.03**
Observations	1’069	848	3’647		

Panel B. Total change in land use (millions ha)

	Land reaped with soy		
	increased (1)	decreased (2)	no soy (3)
Soy	7.39	-1.10	
Seasonal crops, except soy (area reaped)	2.34	0.28	2.56
Permanent crops	1.11	0.77	2.25
Cattle ranching	-8.90	-3.29	2.77
Forest	-0.17	-7.56	-11.50
Unusable land	-2.69	-1.59	-2.64

Note: The table reports the change in land of farms by their principal activity. Panel A reports the average change across all municipalities that experienced an increase in the land reaped with soybean (column 1) a decrease (column 2) or were not producing soy neither in 1996 nor in 2006 (column 3). The last 2 columns report the difference in the change in land across the municipalities in the first 2 groups and the *p*-value from a test of equality of means. Panel B reports the total change for the three groups of municipalities. The first row in both panels reports the change in land reaped with soybean, the second line reports the area reaped with seasonal crops except soybean. All data are sourced from the IBGE Sidra repository. Total area reaped comes from table 501 in 1996 and table 1823 for 2006. Land in farm by principal activity in 1996 comes from table 491 and for 2006 from table 797. The definition of “principal activity” of the farm changed somehow between 1996 and 2006. In 1996 higher specialization was required for farms to be classified under one of the categories reported, and those that did not produce at least 2/3 of the value within a single category were classified under the “mixed activity” category. In 2006 farms were classified according to the activity that accounted for the simple majority of production and no “mixed activity” category existed. Source: Agricultural census 1996 and 2006.

Table 4: Correlation between changes in area reaped with seasonal crops: 1996-2006

Panel A. Average change in land reaped (ha)

	Land reaped with soy				
	increased (1)	decreased (2)	no soy (3)	Difference (1) - (2)	*p-value*
Soy	6’912	-1’295		8’206	0.00***
Maize	1’272	-483	97	1’754	0.00***
Wheat	517	-30	0	547	0.00***
Tobacco	194	34	15	160	0.00***
Rice	-265	-283	-17	18	0.87
Cotton	180	182	-44	-2	0.99
Beans	-120	27	91	-148	0.20
Cassava	102	283	350	-180	0.06*
Sugar	356	602	127	-246	0.19
Other seasonal crops	592	916	576	-325	0.24
Observations	1’069	848	3’647		

Panel B. Total change in land reaped (millions ha)

	Land reaped with soy			
	increased (1)	decreased (2)	no soy (3)	Total
Soy	7.39	-1.10		6.29
Maize	1.35	-0.40	0.34	1.29
Wheat	0.48	-0.02	0.00	0.46
Tobacco	0.19	0.03	0.05	0.27
Rice	-0.24	-0.20	-0.05	-0.49
Cotton	0.17	0.14	-0.14	0.17
Beans	-0.13	0.02	0.33	0.23
Cassava	0.10	0.22	1.17	1.49
Sugar	0.32	0.44	0.38	1.15
Other seasonal crops	0.63	0.78	2.10	3.51

Note: The table reports the change in area reaped with the main seasonal crops in Brazil. Panel A reports the average change across all municipalities that experienced an increase in the land reaped with soybean (column 1) a decrease (column 2) or were not producing soy neither in 1996 nor in 2006 (column 3). The last 2 columns report the difference in the change in area reaped across the municipalities in the first 2 groups and the p-value from a test of equality of means. Panel B reports the total change for the three groups of municipalities. Data are sourced from the IBGE Sidra repository. Total area reaped comes from table 501 in 1996 and table 1823 for 2006. Source: Agricultural census 1996 and 2006.

Table 5: OLS Regressions: Changes in Agricultural Production on Changes of Area Reaped with Soy and Maize.

	Δ Value per Worker	Δ Labor Intensity	Δ % Agri Workers
Panel A			
Δ % Soy Area	3.303***	-0.630***	-0.0734**
	(0.281)	(0.210)	(0.0358)
N	3,841	3,838	3,921
Panel B			
Δ % Maize Area	2.907***	0.679***	0.0204
	(0.209)	(0.160)	(0.0252)
N	4,062	4,053	4,112

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 6: OLS Regressions: Changes in Industrial Labor Market on Changes of Area Reaped with Soy and Maize.

	Δ Total Employment	Δ Number of Plants	Δ Plant Size	Δ Wage
Panel A				
Δ % Soy Area	1.087***	0.547**	0.662**	0.066
	(0.379)	(0.268)	(0.334)	(0.155)
N	2,048	2,063	2,048	2,048
Panel B				
Δ % Maize Area	0.117	0.132	0.046	0.250**
	(0.260)	(0.175)	(0.188)	(0.122)
N	2,172	2,187	2,172	2,172

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 7: First stage: changes in area reaped with soy and maize on potential yield shocks.

	Δ % Soy Area		Δ % Maize Area	
ΔA^{soy}	0.012***	0.025***		0.002
	(0.001)	(0.002)		(0.003)
ΔA^{maize}		-0.003***	0.003***	0.004***
		(0.001)	(0.001)	(0.001)
A^{sugar}		-0.007***		-0.006***
		(0.001)		(0.001)
N	3,921	3,921	4,112	4,112
R-squared	0.054	0.074	0.006	0.013
F-test for joint significance	141.28	52.77	23.95	15.51

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 8: Reduced form: changes in agricultural production on potential yield shocks.

	Δ Value per Worker	Δ Labor Intensity	Δ % Agri Workers
ΔA^{soy}	0.143***	-0.088**	-0.027***
	(0.044)	(0.035)	(0.005)
ΔA^{maize}	-0.025	0.049***	0.010***
	(0.016)	(0.013)	(0.002)
A^{sugar}	-0.036*	-0.027	0.002
	(0.021)	(0.017)	(0.002)
N	4,150	4,146	4,254
R-squared	0.003	0.007	0.013

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 9: Reduced form results: changes in industrial labor market on potential yield shocks.

	Total Employment	Number of Plants	Plant Size	Wage
A^{soy}	0.087***	-0.006	0.094***	-0.046***
	(0.029)	(0.016)	(0.025)	(0.013)
A^{maize}	-0.021	-0.001	-0.020	0.019***
	(0.015)	(0.008)	(0.013)	(0.006)
$P^z A^z$ controls	Yes	Yes	Yes	Yes
AMC & year FE	Yes	Yes	Yes	Yes
N	25,258	25,517	25,258	25,235
R-squared	0.922	0.948	0.809	0.777

Standard errors clustered at AMC level in parentheses
*** p<0.01, ** p<0.05, * p<0.1