Agricultural Productivity and Structural Transformation. Evidence from Brazil*

Paula Bustos Bruno Caprettini Jacopo Ponticelli†

First Draft: July 2012

This Draft: November 2014

Abstract

We study the effects of the adoption of new agricultural technologies on structural transformation. To guide empirical work, we present a simple model where the effect of agricultural productivity on industrial development depends on the factor bias of technical change. We test the predictions of the model by studying the introduction of genetically engineered soybean seeds in Brazil, which had heterogeneous effects on agricultural productivity across areas with different soil and weather characteristics. We find that technical change in soy production was strongly labor saving and lead to industrial growth, as predicted by the model.

Keywords: Agricultural Productivity, Structural Transformation, Industrial Development, Labor Saving Technical Change, Genetically Engineered Soy.

JEL Classification: F16, F43, 014, Q16.

*We received valuable comments from David Atkin, Vasco Carvalho, Gino Gancia, Gene Grossman, Juan Carlos Hallak, Chang-Tai Hsieh, Nina Pavcnik, Joao Pinho de Melo, Silvana Tenreyro, Jaume Ventura, and participants at presentations held at Pontificia Universidade Catolica de Rio de Janeiro, Fundaçao Getulio Vargas, CREI, UPF, LSE, Universidad Torcuato di Tella, SED annual meeting, CEPR ESSIM, CEPR ERWIT, Barcelona GSE Summer Forum, Princeton IES Summer Workshop, AEA Annual Meeting, Harvard University, Brown University, Copenhagen Business School, University of Southern California, NBER SI International Trade and Investment, NBER SI Development and Entrepreneurship, NBER SI Development Economics, MIT, Chicago Booth, and Federal Reserve Bank of New York. We acknowledge financial support from the Private Enterprise Development in Low-Income Countries Project by the CEPR and UK Department for International Development.

†Bustos: CREI, Universitat Pompeu Fabra and Barcelona GSE, pbustos@crei.cat. Caprettini: Universitat Pompeu Fabra, bruno.caprettini@upf.edu. Ponticelli: University of Chicago Booth School of Business, jacopo.ponticelli@chicagobooth.edu.

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 [Clark (1940), 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.[1] Classical models of structural transformation formalize their ideas by showing how productivity growth in agriculture can release labor or generate demand for manufacturing goods.[2] 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. Despite the richness of the theoretical literature, there is scarce direct empirical evidence testing the mechanisms proposed by these models.[3]

In this paper we provide direct empirical evidence on the effects of technical change in agriculture on the industrial sector by studying the recent widespread adoption of new agricultural technologies in Brazil. First, we analyze the effects of the adoption of genetically engineered soybean seeds (GE soy). This new technology requires less labor per unit of land to yield the same output, thus can be characterized as land-biased technical change. In addition, we study the effects of the adoption of second-harvest maize (*milho safrinha*). This type of maize permits to grow two crops a year, effectively increasing the land endowment. Thus, it can be characterized as labor-biased technical change.[4] The simultaneous expansion of these two crops permits to assess the effect of agricultural productivity on structural transformation in open economies.

To guide empirical work, we build a simple model describing a two-sector small open economy where technical change in agriculture can be factor-biased. The model predicts that a Hicks-neutral increase in agricultural productivity induces a reduction in the size of the industrial sector

[1] See, for example, Rosenstein-Rodan (1943), Nurkse (1953), Lewis (1954), Rostow (1960).

[2] See Baumol (1967), Murphy, Shleifer, Vishny (1989), Kongsamut, Rebelo and Xie (2001), Gollin, Parente and Rogerson (2002), Ngai and Pissarides (2007).

[3] Empirical studies of structural transformation include Foster and Rosenszweig (2004, 2008), Nunn and Qian (2011), Michaels, Rauch and Redding (2012), Hornbeck and Keskin (2012). We discuss this literature in more detail below.

[4] Land augmenting technical change is labor-biased when the production displays an elasticity of substitution between land and labor smaller than one.

as labor reallocates towards agriculture, as in Matsuyama (1992). Similar results are obtained when technical change is labor-biased. However, if technical change is strongly labor-saving, labor demand in agriculture falls and workers reallocate towards manufacturing. In sum, the model predicts that the effects of agricultural productivity on structural transformation in open economies depend on the factor-bias of technical change.

In a first analysis of the data we find that regions where the area cultivated with soy expanded experienced an increase in agricultural output per worker, a reduction in labor intensity in agriculture and an expansion in industrial employment. These correlations are consistent with the theoretical prediction that the adoption of strongly labor saving agricultural technologies reduces labor demand in the agricultural sector and induces the reallocation of workers towards the industrial sector. However, causality could run in the opposite direction. For example: an increase in labor demand in the industrial sector could increase wages, inducing agricultural firms to switch to less labor intensive crops, like soy.

We propose to establish the direction of causality by using two sources of exogenous variation in the profitability of technology adoption. First, in the case of GE soy, as the technology was invented in the U.S. in 1996, and legalized in Brazil in 2003, we use this last date as our source of variation across time. 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. Similarly, in the case of maize, we exploit the timing of expansion of second-harvest maize and cross-regional differences in soil suitability.

We obtain an exogenous measure of technological change in agriculture by using estimates of potential soil yields across geographical areas of Brazil from the FAO-GAEZ database. These yields are calculated by incorporating local soil and weather characteristics into a model that predicts the maximum attainable yields for each crop in a given area. 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 traditional and new agricultural technologies. Thus, we exploit the predicted differential impact of the new technology on yields across geographical areas in Brazil as our source of cross-sectional variation in agricultural productivity. Note that this empirical strategy relies on the assumption that although goods can move across geographical areas of Brazil, labor markets are local due to limited labor mobility. This research design allows us to investigate whether exogenous shocks to local agricultural productivity lead to changes in the size of the local industrial sector. We use municipalities as our geographical unit of observation, that are assumed to behave as the small open economy described in the model.[5]

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. In addition, these regions experienced increases in the value of agricultural output per worker and reductions in labor intensity measured as employment per hectare. These regions experienced faster employment growth and wage reductions in the industrial sector. Interestingly, the effects of technology adoption are different for maize. Regions where the FAO potential maize yields are predicted to increase the most when switching from the traditional to the new technology did indeed experience a higher increase in the area planted with maize and in the value of agricultural output. However, they also experienced increases in labor intensity, reductions in industrial employment and increases in wages.

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 determinant of the relationship between agricultural productivity and structural transformation in open economies. If technical change is labor-biased, as in the case of maize, agricultural productivity growth leads to a reduction in industrial employment, as predicted by Matsuyama (1992). However, if technical change is strongly labor saving, as in the case of GE soy, agricultural productivity growth leads to employment growth in the industrial sector.

Our estimates can be used to quantify the effect of factor-biased agricultural technical change on structural transformation. First, we estimate the elasticity of the agricultural employment share to changes in agricultural labor productivity induced by soy technical change: 1 log point increase in agricultural labor productivity corresponds to 0.136 percentage points decrease in the agricultural employment share. Next, we illustrate the magnitude of our estimates by calculating how much of the differences in the speed of structural transformation across Brazilian regions can soy technical change explain. A municipality experiencing a one standard deviation larger increase in potential soy yield experienced an 11 log points larger increase in agricultural labor productivity, and a corresponding 1.52 percentage points larger decrease in agricultural employment share. This estimate corresponds to 21% of a standard deviation in the change of agricultural employment share between 2000 and 2010 (7.3 percentage points). A parallel calculation reveals that a municipality shocked with a one standard deviation increase in potential soy yield would experience an increase in the manufacturing employment share of a similar magnitude, which corresponds to 31% of a standard deviation in the change of manufacturing employment share between 2000 and 2010 (5.4

[5] Because the size of municipalities is small in coastal areas of Brazil, we show that our results are robust to using a larger unit of observation, Micro-regions.

percentage points).

The estimates of the effects of technical change on the agricultural and manufacturing employment shares discussed above have a similar magnitude. At the same time, the estimates of the effects on the service and public sectors are very small and not statistically different from zero. This implies that labor reallocated mostly from agriculture to manufacturing. The absence of an effect of agricultural technical change on services is, at first sight, puzzling. A simple extension of our model to include a non-tradable sector would imply that as higher productivity increases income, the demand for all goods increases. Because non traded goods need to be produced locally, some labor must reallocate from agriculture towards the service sector. Our ongoing empirical work on this matter indicates that the absence of an average effect on services is related to the asymmetric effects of soy technical change on the income of land owners and workers and the fact that in some areas of Brazil land owners do not reside locally.

Finally, let us note that we perform the following robustness checks. First, we show that our estimates are stable when we allow municipalities with different initial levels of development to be on differential structural transformation trends. Second, we relax the assumption of constant factor endowments within municipalities. We show that we obtain similar estimates in the subsample of Brazilian municipalities where the land endowment did not increase. Similarly, we show that contemporaneous migration patterns are consistent with the predictions of the model: there is out (in) migration in areas more affected by strongly land biased (labor biased) technical change. Third, we show that our estimates are not driven by pre-existing trends in manufacturing employment nor migration flows. Fourth, we show that our results are robust to using a larger unit of observation, micro-regions. Fifth, we show that at least 2/3 of our estimated effect of agricultural technical change on the manufacturing employment share is not driven by the processing of soy and maize in downstream industries nor larger agricultural sector demand for manufacturing inputs. Sixth, we show that our estimates are not driven by contemporaneous changes in commodity prices. Finally, we show that our main results remain statistically significant when we correct standard errors to account for spatial correlation.

Related Literature

There is a long tradition in economics of studying the links between agricultural productivity and industrial development. Nurkse (1953) and Rostow (1960) argued that agricultural productivity growth was an essential precondition for the industrial revolution. Schultz (1953) held the view that an agricultural surplus is a necessary condition for a country to start the development process. Classical models of structural transformation formalized their ideas by proposing two main mechanisms through which agricultural productivity can speed up industrial growth in closed economies.

First, the demand channel: agricultural productivity growth rises income per capita, which generates demand for manufacturing goods if preferences are non-homothetic [Murphy, Shleifer, Vishny (1989), Kongsamut, Rebelo and Xie (2001), Gollin, Parente and Rogerson (2002)]. The higher relative demand for manufactures generates a reallocation of labor away from agriculture. Second, the supply channel: if productivity growth in agriculture is faster than in manufacturing and these goods are complements in consumption, then the relative demand of agriculture does not grow as fast as productivity and labor reallocates towards manufacturing [Baumol (1967), Ngai and Pissarides (2007)].[6,7]

The view that agricultural productivity can generate manufacturing growth was challenged by scholars studying industrialization experiences in open economies. These scholars argued that high agricultural productivity can retard industrial growth as labor reallocates towards the comparative advantage sector [Mokyr (1976), Field (1978) and Wright (1979)] . Matsuyama (1992) formalized these ideas by showing how the demand and supply channels are not operative in a small open economy that faces a perfectly elastic demand for both goods at world prices. The open economy model we present in this paper differs from Matsuyama's in one key dimension. In his model, there is only one type of labor thus technical change is, by definition, Hicks-neutral. In our model agricultural production uses both land and labor, and technical change can be factor-biased. Thus, a new prediction emerges: when technical change is strongly labor saving an increase in agricultural productivity leads to industrial growth even in open economies.

Our work also builds on the empirical literature studying the links between agricultural productivity and economic development.[8] The closest precedent to our work is Foster and Rosenzweig (2004, 2008) who study the effects of the adoption of high-yielding-varieties (HYV) of corn, rice, sorghum, and wheat during the Green Revolution in India. To guide empirical work, they present a model where agricultural and manufacturing goods are tradable and technical change is Hicks-neutral. Consistent with the model, they find that villages with higher improvements in crop yields experienced lower manufacturing growth. Our findings are in line with theirs in the case of Maize, where technical change is labor-biased. However, we find the opposite effects in the case of soy, where technical change in strongly labor saving. Thus, relative to theirs, our work highlights the importance of the factor bias of technical change in shaping the relationship between agricultural

[6] The agricultural and manufacturing goods are complements in consumption of the elasticity of substitution between the two goods is less than one.

[7] Another mechanism generating a reallocation of labor from agriculture to manufacturing is faster growth in the relative supply of one production factor when there are differences in factor intensity across sectors [See Caselli and Coleman (2001), and Acemoglu and Guerrieri (2008)]. For a recent survey of the structural transformation literature see Herrendorf, Valentinyi and Rogerson (2013).

[8] This literature is surveyed by Syrquin (1988) and Foster and Rosensezweig (2008).

productivity and industrial development in open economies.

Finally, our work is also related to recent empirical papers studying the effects of agricultural productivity on urbanization [Nunn and Qian (2011)], the links between structural transformation and urbanization [Michaels, Rauch and Redding (2012)], and the effects of agriculture on local economic activity [Hornbeck and Keskin (2012)].

The remaining of the paper is organized as follows. Section 2 gives background information on agriculture in Brazil. Section 3 presents the theoretical model. Section 4 describes the data. Section 5 presents the empirical strategy and results. Section 7 concludes.

2 Agriculture in Brazil

In this section we provide background information about recent developments in the Brazilian agricultural sector. As Figure 1 shows, in the last decade, Brazilian labor force has been shifting away from agriculture and increasing in manufacturing and services. At the end of the 1990s, agriculture employed around 16 million workers, while manufacturing less than 8 million. By 2011, this gap was almost closed with agriculture and manufacturing employing, respectively, 12 and 10.5 million workers.

During the same period, agricultural productivity increased significantly. Figures 2 and 3 compare the distributions of, respectively, average soy yields and average maize yields (expressed in tons per hectare) across Brazilian municipalities in 1996 and 2006. The figures show a clear shift to the right in the distribution of average yields for both soy and maize, the two major crops produced in Brazil. Productivity growth went hand-in-hand with an expansion in the area planted. Table 1 shows that the land cultivated with seasonal crops – i.e. crops produced from plants that need to be replanted after each harvest, such as soy and maize – increased by 10.4 million hectares between 1996 and 2006. Out of these 10.4 million, 6.2 million hectares were converted to soy cultivation.

During this period new agricultural technologies were adopted in the cultivation of both soy and maize. In the case of soy, Brazilian farmers started introducing on a large scale genetically engineered (GE) seeds. In the case of maize, Brazilian farmers started introducing a second harvesting season, which requires the use of advanced cultivation techniques.

2.1 Technical Change in Soy: Genetically Engineered Seeds

The first generation of GE soy seeds, the Roundup Ready (RR) variety, was commercially released in the U.S. in 1996 by the agricultural biotechnology firm Monsanto. In 1998 the Brazilian National

Technical Commission on Biosecurity (CTNBio) authorized Monsanto to field-test GE soy in Brazil for 5-years as a first step before commercialization. However, reports from the Foreign Agricultural Service of the United States Department of Agriculture (USDA) document that smuggling of GE soy seeds from Argentina – where they were approved for cultivation since 1996 – was already taking place from 2001 (USDA, 2001, p. 63). Eventually, pressure from soy farmers led the Brazilian government to legalize cultivation of GE soy seeds in 2003.[9]

The main advantage of GE soy seeds relative to traditional seeds is that they are herbicide resistant. This allows the use of no-tillage planting techniques.[10] The planting of traditional seeds is preceded by soil preparation in the form of "tillage", the operation of removing the weeds in the seedbed that would otherwise crowd out the crop or compete with it for water and nutrients. In contrast, planting GE soy seeds requires no tillage, as the application of herbicide will selectively eliminate all unwanted weeds without harming the crop. As a result, GE soy seeds can be applied directly on last season's crop residue, allowing farmers to save on production costs since less labor is required per unit of land to obtain the same output.[11]

The new technology spread quickly: in 2006 GE seeds were planted in 46.4% of the area cultivated with soy in Brazil, according to the last Agricultural Census (IBGE, 2006, p.144). In the following years the technology continued spreading to the point that it covered 85% of the area planted with soy in Brazil in the 2011-2012 harvesting season, according to the Foreign Agricultural Service of the USDA (USDA, 2012). The timing of adoption of GE soy coincides with a fast expansion in the area planted with soy in Brazil. Figure 4 documents the evolution of the area planted with soy since 1980. The figure shows that this area grew slightly between 1980 and 1996, but experienced a fast expansion afterwards. In particular, note that growth in the soy area accelerated after 2001 when the USDA documents that GE soy seeds started to be smuggled from

[9] In 2003, law 10.688 allowed the commercialization of GE soy for one harvesting season, requiring farmers to burn all unsold stocks after the harvest. This temporary measure was renewed in 2004. Finally, in 2005, law 11.105 – the New Bio-Safety Law – authorized production and commercialization of GE soy in its Roundup Ready variety (art. 35).

[10] Genetic engineering (GE) techniques allow a precise alteration of a plant's traits. This allows to target a single plant's trait, facilitating the development of plant characteristics with a precision not attainable through traditional plant breeding. In the case of herbicide resistant GE soy seeds, soy genes were altered to include those of a bacteria that was herbicide resistant.

[11] GE soybeans seeds allow farmers to adopt a new "package" of techniques that lowers labor intensity for several reasons. First, since GE soybeans are resistant to herbicides, weed control can be done more flexibly. Herbicides can be applied at any time during the season, even after the emergence of the plant (Duffy and Smith, 2001). Second, GE soybeans are resistant to a specific herbicide (glyphosate), which needs fewer applications: fields cultivated with GE soybeans require an average of 1.55 sprayer trips against 2.45 of conventional soybeans (Duffy and Smith, 2001; Fernandez-Cornejo et al., 2002). Third, no-tillage production techniques require less labor. This is because the application of chemicals needs fewer and shorter trips than tillage. In addition, no-tillage allows greater density of the crop on the field (Huggins and Reganold, 2008). Finally, farmers that adopt GE soybeans report gains in the time to harvest (Duffy and Smith, 2001). These cost savings might explain why the technology spread fast, even though experimental evidence in the U.S. reports no improvements in yield with respect to conventional soybeans (Fernandez-Cornejo and Caswell, 2006)

Argentina.

The expansion of the area planted with soy can affect labor demand in the agricultural sector through two channels. First, soybean production is one of the least labor-intensive agricultural activities, as documented in Table 2.[12] As a result, the expansion of soy cultivation over areas previously devoted to other agricultural activities tends to reduce the labor intensity of agricultural production (*across-crops effect*). Second, during the period under study the expansion of the area cultivated with soy was accompanied by a steady reduction in the number of people employed in soy cultivation, as documented in Figure 6. The expansion of the area cultivated, together with the constant reduction in employment determined the drop in labor intensity of soy cultivation shown in Figure 5. This effect also tends to reduce the labor intensity of agricultural production (*within-crop effect*), as documented in Table 2.

2.2 Technical Change in Maize: Second Harvesting Season

During the last two decades Brazilian agriculture experienced also important changes in maize cultivation. Maize used to be cultivated as soy, during the summer season that takes place between August and December. At the beginning of the 1980s a few farmers in the South-East started producing maize after the summer harvest, between March and July. This second season of maize cultivation spread across Brazil, where it is now known as *milho safrinha* (small-harvest maize).

Cultivation of a second season of maize requires the use of modern cultivation techniques for several reasons. First, more intensive land-use removes nitrogen from the soil, which needs to be replaced by fertilizers (EMBRAPA, 2006). Second, the planting of a second crop requires careful timing, as yields drop considerably due to late planting. Then, herbicides are used to remove residuals from the first harvest on time to plant the second crop. In addition, the second season crop needs to be planted one month faster than the first, which usually requires higher mechanization (CONAB, 2012). Finally, because a second-harvest implies a more intensive use of the soil, farmers have to rely mostly on no-tillage techniques (EMBRAPA, 2006).

Note that, even with advanced cultivation techniques, maize is still more labor intensive than both soy and other agricultural activities like cattle ranching (see Table 2). 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.[13,14]

[12] In 2006 it required less than 20 workers per 1000 hectares against the 84 of the average seasonal crop and the 127 of the average permanent crop.

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

[14] In Table 2 we do not report productivity for Brazilian farms whose main activity was maize cultivation because publicly available data on area in farms and number of workers by principal activity is available only for farms whose principal activity is either soy or cereals, a category that includes rice, wheat, maize and other cereals. The breakdown

Figure 7 documents the evolution of the area cultivated with maize since 1980. The figure shows that, although the total area devoted to maize has increased only slightly, the area devoted to second season maize has expanded steadily since the beginning of the 1990s.[15]

3 Model

In this section we present a simple model to illustrate the effects of factor-biased technical change on structural transformation in open economies. We consider a small open economy where there are two sectors, agriculture and manufacturing, and two production factors, land and labor.

3.1 Setup

This small open economy has a mass one of residents, each endowed with L units of labor. There are two goods, *manufactures* and *agriculture*, both of which are tradable. Production of the manufactured good requires only labor and labor productivity in manufacturing is A_m, so that

$$Q_m = A_m L_m \tag{1}$$

where Q_m denotes production of the manufactured good and L_m denotes labor allocated to the manufacturing sector. Production of the agricultural good requires both labor and land, and takes the CES form:

$$Q_a = A_a \left[\gamma \left(A_L L_a \right)^{\frac{\sigma-1}{\sigma}} + (1-\gamma) \left(A_T T_a \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{2}$$

where Q_a denotes production of the agricultural good, the two production factors are labor (L_a) and land (T_a), A_a is hicks-neutral technical change, A_L is labor-augmenting technical change and A_T is land-augmenting technical change. The parameter $\gamma \in (0,1)$, and the parameter $\sigma > 0$ captures the elasticity of substitution between land and labor. The production function described by equation (2) implies the following ratio of marginal product of land to marginal product of labor:

$$\frac{MPT_a}{MPL_a} = \frac{1-\gamma}{\gamma} \left(\frac{A_T}{A_L} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{T_a}{L_a} \right)^{-\frac{1}{\sigma}}$$

Thus, if land and labor are complements in production ($\sigma < 1$), labor-augmenting technical change is land-biased. That is, increases in A_L rise the marginal product of land relative to labor for

of area and number of worker by type of cereal cultivated is available in 1996, and in this year labor intensity in maize is 100.4 workers per 1000 ha of area cultivated, slightly above the labor intensity of all the farms that cultivate cereals (92.4, as reported in Table 2).

[15] Data on area cultivated with maize broken down by the season of harvest of maize are available only at the aggregate level. For this reason in section 5, when we study municipality-level data, we will not be able to distinguish between the two maize cultivation seasons.

a given amount of land per worker. Similarly, land-augmenting technical change is labor-biased. Finally, technical change is strongly labor-saving if improvements in technology reduce the marginal product of labor. In the case of labour-augmenting technical change, this requires $\frac{\partial MPL_a}{\partial A_L} < 0$, which imposes a stronger condition on the elasticity of substituiton:[16]

$$\sigma < \frac{(1-\gamma)(A_T T)^{\frac{\sigma-1}{\sigma}}}{\gamma (A_L L_a)^{\frac{\sigma-1}{\sigma}} + (1-\gamma)(A_T T)^{\frac{\sigma-1}{\sigma}}} < 1. \tag{3}$$

Note that this condition is more likely to be satisfied the more complementary are land and labor in production and the more important is land relative to labor in production.[17]

Consumers have homotetic preferences over the agricultural and manufacturng good: $U(C_a, C_m)$ where $\frac{\partial U}{\partial C_i} > 0$ and $\frac{\partial^2 U}{\partial C_i^2} < 0$ for $i = a, m$.

3.2 Equilibrium

We consider a small open economy that trades with a world economy where the relative price of the agricultural good is $\frac{P_a}{P_m} = \left(\frac{P_a}{P_m}\right)^*$.Profit maximization implies that the value of the marginal product of labor must equal the wage in both sectors, thus:

$$P_a MPL_a = w = P_m MPL_m. \tag{4}$$

This implies that, in equilibrium, the marginal product of labor is determined by international prices and manufacturing productivity:

$$MPL_a = \left(\frac{P_m}{P_a}\right)^* A_m. \tag{5}$$

The equilibrium allocation of labor can be determined by substituting the land market clearing condition, $T_a = T$, in equation 5:

$$A_a \left[\gamma (A_L L_a)^{\frac{\sigma-1}{\sigma}} + (1-\gamma)(A_T T)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}-1} \gamma (A_L L_a)^{\frac{\sigma-1}{\sigma}-1} A_L = \left(\frac{P_m}{P_a}\right)^* A_m. \tag{6}$$

The above equation 6 implicitely defines the equilibrium level of employment in agriculture, L_a^{eq}. In turn, the equilibrium level of employment in manufacturing, L_m^{eq}, can be determined using the labor market clearing condition, $L_m + L_a = L$. Once L_m^{eq} and L_a^{eq} are determined output in each sector can be found using the production functions described in equations 2 and 1. Equilibrium consumption is finally determined by: $\frac{\partial U/\partial C_a}{\partial U/\partial C_m} = \left(\frac{P_a}{P_m}\right)^*$ and the zero trade balance condition

[16] See Acemoglu (2010) for a discussion and more general definition of strongly labor-saving technical change.

[17] See Appendix A for a formal proof.

$(Q_a - C_a) = \left(\frac{P_m}{P_a}\right)^* (Q_m - C_m).$

3.3 Technological Change and Structural Transformation

In this section we assess the response of the employment share of agriculture to three types of technological change: Hicks-neutral, labor-augmenting and land-augmenting. We assume that land and labor are complements in production, thus labor-augmenting technical change is land-biased and land-augmenting technical change is labor-biased.

Hicks-neutral technical change

An increase in A_a generates a reallocation of labor from manufacturing to agriculture, that is $\frac{\partial L_a^{eq}}{\partial A_a} > 0$ and $\frac{\partial L_m^{eq}}{\partial A_a} < 0$. To see why this is the case, note that, in equlibrium, the marginal product of labor in agriculture is given by international prices and manufacturing productivity, thus it must stay constant when A_a increases. However, the increase in agricultural productivity rises the marginal product of labor in agriculture because $\frac{\partial MPL_a}{\partial A_a} > 0$. Thus, employment in agriculture must increase to reduce the marginal product of labor to its equilibrium level, because $\frac{\partial MPL_a}{\partial L_a} < 0$ (see Appendix A for a proof).

Land-augmenting technical change (labor-biased)

An increase in A_T generates a reallocation of labor from manufacturing to agriculture, that is $\frac{\partial L_a^{eq}}{\partial A_T} > 0$ and $\frac{\partial L_m^{eq}}{\partial A_T} < 0$. To see why this is the case, note that the land-augmenting technical change rises the marginal product of labor in agriculture because $\frac{\partial MPL_a}{\partial A_T} > 0$ as long as $\sigma < 1$ (see Appendix A for a proof). Thus, employment in agriculture must increase to bring the marginal product of labor back to its equilibrium level, because $\frac{\partial MPL_a}{\partial L_a} < 0$.

Labor-augmenting technical change (land-biased)

A. Strongly labor saving

If land and labor are strong complements in production, that is, the elasticity of substitution, σ, satisfies the condition stated in equation (3), labor-augmenting technical change is not only land-biased but also strongly labor-saving. In this case, an increase in A_L generates a reallocation of labor from agriculture to manufacturing, that is $\frac{\partial L_a^{eq}}{\partial A_L} < 0$ and $\frac{\partial L_m^{eq}}{\partial A_L} > 0$. This is because technical change induces a reduction in the marginal product of labor in agriculture, that is $\frac{\partial MPL_a}{\partial A_L} < 0$. However, in equlibrium, the marginal product of labor in agriculture is given by international prices and manufacturing productivity, thus it must stay constant when A_L changes. Thus, as $\frac{\partial MPL_a}{\partial L_a} < 0$, employment in agriculture must fall to bring the marginal product of labor back to its equilibrium level.

B. *Weakly labor saving*

If the elasticity of substitution, σ, is smaller than one but does not satisfy the condition stated in equation (3), labor-augmenting technical change is land-biased but not strongly labor-saving. Thus, an increase in A_L generates a reallocation of labor from manufacturing to agricluture, that is $\frac{\partial L_a^{eq}}{\partial A_L} > 0$ and $\frac{\partial L_m^{eq}}{\partial A_L} < 0$. This is because technical change induces an increase in the marginal product of labor in agriculture, that is $\frac{\partial MPL_a}{\partial A_L} > 0$. Thus, agricultural employment must increase to bring the marginal product of labor back to its equilibrium level.

3.4 Empirical Predictions

The model predicts that, in a small open economy, a Hicks-neutral increase in agricultural productivity induces a reduction in the size of the industrial sector as labor reallocates towards agriculture, as in Matsuyama (1992). Similar results are obtained when technical change is labor-biased. However, if technical change is strongly labor-saving, labor demand in agriculture falls and workers reallocate towards manufacturing. In sum, the model predicts that the effects of agricultural productivity on structural transformation in open economies depend on the factor-bias of technical change.

In the following section, we test the predictions of the model by studying the simultaneous expansion of two new agricultural technologies: GE soy and second-harvest maize. In the case of soy, the advantage of GE seeds relative to traditional ones is that they are herbicide resistant, which reduces the need to plow the land. As a result, this new technology requires less labor per unit of land to yield the same output and can be characterized as labor-augmenting technical change. As discussed above, the effect of labor-augmenting technical change on structural transformation depends on the elasticity of substitution between land and labor in the agricultural production function. In the case where land and labor are strong complements, then technical change is expected to reduce the labor intensity of agricultural production and employment in agriculture as labor reallocates towards manufacturing. Thus, in this case, we expect that the adoption of GE soy reduces the labor intensity of agricultural production and reallocates labor from agriculture towards manufacturing. Note, however, that if the complementarity between land and labor is not strong enough, we obtain the opposite prediction: the labor intensity of agricultural production increases and labor reallocates towards agriculture. In the case of maize, farmers started introducing a second harvesting season, which requires the use of advanced cultivation techniques and inputs. Second-harvest maize (*milho safrinha*) permits to grow two crops a year, effectively increasing the land endowment. In the case where land and labor are complements in production, land-augmenting technical change can be characterized as labor-biased. Thus, we expect that the adoption of second-harvest maize increases the labor intensity of agricultural production and reallocates labor from manufacturing towards agriculture.

4 Data

In this paper we use three main data sources: the Agricultural Census for data on agriculture, the Population Census for data on the sectoral composition of employment and wages, and the FAO Global Agro-Ecological Zones database for potential yields of soy and maize. To perform robustness checks we also use manufacturing plant-level data from the Brazilian Yearly Industrial Survey (PIA).[18]

The Agricultural Census is released at intervals of 10 years by the *Instituto Brasileiro de Geografia e Estatística* (IBGE), the Brazilian National Statistical Office. 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, which were commercially released in the U.S. in 1996 and 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 distinguish 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.

Data on the sectoral composition of the economy and average wages is constructed using the Brazilian Population Census. The Census is carried out every 10 years and it covers the entire Brazilian population. We use data from the last two rounds of the census (2000 and 2010) so that we can observe the variables of interest both before and after the legalization of the new technology.[19] Data on the sector of employment is collected both in 2000 and 2010 through a special survey that is administered to a sample of around 11% of the Brazilian population (*questionário da amostra*). The sample is selected to be representative of the Brazilian population within narrow cells defined by geographical district, sex, age and urban or rural situation. The variables we focus on are the sector in which the person was working during the previous week and its wage. [20] For each municipality, we compute the employment share in manufacturing as the number of people

[18] In this section we briefly discuss the main data sources and variables of interest. For detailed variable definition and data sources please refer to Appendix B.

[19] To perform some of the robustness checks we also use the 1991 Population Census.

[20] The sector classification is comparable across the census of 2000 and 2010 and it is the CNAE Domiciliar 1.0. The broader categories of CNAE Domiciliar 1.0 follow the structure of the ISIC classification version 3.1.

working in CNAE sectors from 15 to 37 divided by the total number of people employed in that municipality.

Our third source of data is the Global Agro-Ecological Zones database produced by the FAO, which provides data on potential yields for soy and maize. 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 different theoretical levels of technology. We focus on the two extreme levels of technology used in production: *low* and *high*. When the level of technology is assumed to be *low*, agriculture is not mechanized, it uses traditional cultivars and does not use nutrients or chemicals for pest and weed control. When the level of technology is *high* instead, production is fully mechanized, it uses improved and high yielding varieties and "optimum" application of nutrients and chemical pest, disease and weed control.[21,22] The database reports potential yields for each crop under low and high technological levels for a worldwide grid at a resolution of 9.25 × 9.25 km. Figures 8 and 9 show the potential yields for soybean in Brazil under, respectively, low and high technology. Figure 10 and 11 show the correspondent maps for maize.

In order to match the potential yields data with agriculture and industry variables we superimposed each of the potential yields' maps with political maps of Brazil reporting the boundaries of either municipality or micro-regions (a larger administrative unit of observation that encompass several municipalities). Next, we compute the average potential yield of all cells falling within the boundaries of every geographical unit. We repeated this operation for both soy and maize and for each of the two levels of technology. Our measure of technical change in soy or maize production within each municipality is obtained as the potential yield under high technology minus the potential yield under low technology. Figure 12 illustrates the resulting measure of technical change in soy at the municipality level, while Figure 13 shows the same measure at the micro-region level.

Finally, in order to perform some robustness checks, we use data from the *Pesquisa Industrial*

[21] The description of each technology in the FAO-GAEZ dataset documentation 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."

[22] Note that the GAEZ documentation provides rather general definitions of low and high technologies. In direct consultations with the authors, they indicate that these definitions are stylized to represent a range of management conditions. In particular, the term "improved crop varieties" means crop varieties that are developed/ discovered to be well adapted to the conditions where they are grown. As such, model parameters are adjusted to approximate yields under low input or high input production but does not produce estimates of yields for each individual seed.

Anual (PIA), the Yearly Industrial Survey carried out by the IBGE. This survey monitors the performance of Brazilian firms in the extractive and manufacturing sectors. We focus on the manufacturing sector as defined by CNAE 1.0 (sectors 15 to 37).[23] We use yearly data from 1996 to 2006. 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 sample of firms having between 5 and 29 employees (*estrato amostrado*) and it is representative at the sector and state level. The second includes all firms having 30 or more employees (*estrato certo*). We focus on the sample of firms with 30 or more employees which is representative at municipality level. The variables we focus on are: total employment and average wages.

5 Empirics

In this section we study the effects of the adoption of new agricultural technologies on structural transformation in Brazil. For this purpose, we first study the effect of the adoption of GE soy and second season maize on agricultural productivity and the factor intensity of agricultural production. This first step permits to characterize the factor-bias of technical change. Next, we assess the impact of factor-biased technical change on the allocation of labor across sectors.

In the following section we report simple correlations between the expansion of the area planted with soy and maize and agricultural and industrial labor market outcomes in each municipality. Note that these correlations are not informative about the causal relation between these variables. In section 5.2, we present an empirical strategy that attempts to establish the direction of causality by exploiting the timing of adoption and the differential impact of the new technology on potential yields across geographical areas.

5.1 Basic Correlations in the Data

We start by documenting how the expansion of soy and maize cultivation during the 1996-2006 period relates to changes in agricultural production and industrial employment. These basic correlations in the data attempt to answer the following question: did areas where soy (maize) expanded experience faster (slower) structural transformation? In section 5.1.1 we present a set of OLS estimates of equations relating agricultural outcomes to the percentage of farm land cultivated with soy and maize. In section 5.1.2 we present a second set of OLS estimates of equations relating

[23] The broad category of CNAE 1.0 are identical to the broad categories of CNAE Domiciliar version 1.0 and of the ISIC classification version 1.0.

manufacturing outcomes to the percentage of farm land cultivated with soy and maize.

The basic form of the equations to be estimated in this section is:

$$y_{jt} = \alpha_j + \alpha_t + \beta \left(\frac{\text{Soy Area}}{\text{Agricultural Area}} \right)_{jt} + \gamma \left(\frac{\text{Maize Area}}{\text{Agricultural Area}} \right)_{jt} + \varepsilon_{jt} \tag{7}$$

where j indexes municipalities, t indexes time, α_j are municipality fixed effects and α_t are time fixed effects. y_{jt} is an outcome that varies across municipalities and time and $\frac{\text{Soy (Maize) Area}}{\text{Agricultural Area}}$ is the total area reaped with soy (maize) divided by total farm land.[24,25] Our source for agricultural variables is the Agricultural Census, thus we observe them for the years 1996 and 2006. Because fixed effects and first difference estimates are identical when considering only two periods, we estimate (7) in first differences:

$$\Delta y_j = \Delta\alpha + \beta\ \Delta \left(\frac{\text{Soy Area}}{\text{Agricultural Area}} \right)_j + \gamma\ \Delta \left(\frac{\text{Maize Area}}{\text{Agricultural Area}} \right)_j + \Delta\varepsilon_j \tag{8}$$

5.1.1 Agricultural Outcomes: Productivity, Labor Intensity and Employment Share

Table 4 reports OLS estimates of equation 8 for three agricultural outcomes. The first is labor productivity, measured as the value of output per worker in agriculture.[26] The second is labor intensity, measured as the number of workers per unit of land in agriculture. The third outcome is the employment share of agriculture, which attempts to capture the extent of structural transformation.[27]

The first two columns of Table 4 show that in areas where soy cultivation expanded, the value of agricultural production per worker increased and labor intensity in agriculture decreased. In contrast, in areas where maize cultivation expanded the labor intensity decreased. This evidence is consistent with our characterization of technical change in soy as land-biased and technical change in maize as labor-biased. The estimated coefficients imply that a 1 percentage point increase in soy area share corresponds to a 0.57 log points increase in labor productivity, and a 0.48 log points reduction in labor intensity. In the case of maize, the estimated coefficients imply that a 1 percentage point increase in maize area share corresponds to a 1.59 log points increase in labor

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

[25] Borders of municipalities often change, thus, to make them comparable across time, IBGE has defined *Área Mínima Comparável* (AMC), smallest comparable areas, which we use as our unit of observation.

[26] This is the broadest measure of labor productivity in agriculture that can be obtained using the publicly available municipality-level data. An advantage of this broad productivity measure is that it captures both the within and across crop effects of technical change.

[27] The share of workers employed in agriculture is defined as total number of workers in agriculture divided by total number of workers in all sectors. This variable is obtained from the Population Census and its first differences are computed between the years 2000 and 2010.

productivity, and a 0.74 log points increase in labor intensity.

Next, we analyze the relationship between the expansion in soy and maize area and sectoral employment shares. Note that we source information on sectoral employment shares from the Population Census which reports information for the years 2000 and 2010. Thus, our estimation of equation (8) relates changes in employment shares between 2000 and 2010 to changes in the area planted with soy and maize between 1996 and 2006. In both cases the initial year precedes the timing of legalization of soybean seeds in Brazil (2003), as well as the first date in which smuggling of GE soy seeds was documented (2001).

Column 3 of Table 4 shows that the employment share of agriculture decreased in places where soy expanded while estimates for maize are not statsitically significant. The estimated coefficient implies that a 1 percentage point increase in soy area share corresponds to a 0,05 percentage point reduction in agricultural employment share.

5.1.2 Manufacturing Outcomes: Employment Share, Total Employment and Wages

We now turn to the question of whether manufacturing employment expanded (contracted) in areas where soy (maize) expanded. Table 5 reports OLS estimates of equation 8 for three manufacturing outcomes: manufacturing employment share, the level of employment in manufacturing, and the average wage in the manufacturing sector.

The first column of Table 5 shows that municipalities where soy expanded experienced a faster increase in the employment share in manufacturing. In contrast, this share remained unchanged in municipalities where maize expanded. Interestingly, in areas where soy expanded, not only the share but also the level of manufacturing employment increased, as shown in column 2. The estimated coefficient on the effect of the expansion of soy cultivation in manufacturing employment share indicates that municipalities experiencing a 1 percentage point increase in soy area share had a 0.085 percentage point increase in manufacturing employment share and a 0.99 log point increase in manufacturing employment.

The finding that manufacturing employment increased in areas where soy expanded suggests that soy technical change is not only land-biased but also strongly labor-saving. In this case, our model predicts that technology adoption reduces labor demand in agriculture inducing a reallocation of labor towards manufacturing.

5.2 The Effect of Agricultural Technological Change on Structural Transformation

In this section we provide empirical evidence on the causal effects of the adoption of new agricultural technologies 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 while industrial employment expanded. These findings are consistent with the sequence of events predicted by the model, namely that the adoption of strongly labour saving agricultural technologies reduces labor demand in the agricultural sector and induces a reallocation of labor towards the industrial sector. However, these correlations are not informative about the direction of causality. For example, these correlations are consistent with the following alternative sequence of events: productivity growth in the industrial sector increases labor demand and wages, inducing agricultural firms to switch to less labor-intensive crops, like soy. In this section we attempt to establish the direction of causality.

Our empirical strategy relies on the assumption that goods can be traded across geographical areas of Brazil but labor markets are local. We investigate whether exogenous shocks to local agricultural productivity lead to changes in the size of the local industrial sector. Thus, our ideal unit of observation would be a region containing a city and its hinterland with limited migration across regions. We attempt to approximate this ideal using municipalities as our main level of geographical aggregation. This approach is adequate for municipalities in the interior of the country, which typically include both rural and urban areas. However, municipalities tend to be mostly urban in more densely populated coastal areas. To address this concern, we control for the share of rural population in each municipality. In addition, we show that our estimates are robust to using a larger unit of observation: micro-regions.[28] Figures 12 and 13 contain maps of Brazil displaying both levels of aggregation.

We propose to identify the causal effect the new technologies on structural transformation by exploiting the timing of adoption and the differential impact of the new technology on potential yields across geographical areas. Let us first consider whether the timing of adoption is likely to be exogenous with respect to developments in the Brazilian economy. GE soy seeds were commercially released in the U.S. in 1996, and legalized in Brazil in 2003. Given that the seeds were developed in the U.S., their date of approval for commercialization in the U.S., 1996, is arguably exogenous with respect to developments in the Brazilian economy. In contrast, the date of legalization, 2003, responded partly to pressure from Brazilian farmers. In addition, smuggling of GE soy

[28] These are groups of several municipalities created by the 1988 Brazilian Constitution and used for statistical purposes by IBGE.

seeds across the border with Argentina is reported since 2001. Thus, in our empirical analysis we would ideally compare outcomes before and after 1996. This is possible when variables are sourced from the Agricultural Census. For variables sourced from the Population Census we compare outcomes before and after 2000. Because this year predates both legalization and the first reports of smuggling, the timing can still be considered exogenous. In contrast, the cultivation techniques necessary to introduce a second harvesting season for maize were developed within Brazil. Thus, the timing of its expansion can not be considered exogenous to other developments in the Brazilian economy. Nevertheless, to the extent that the diffusion of this new technology across space depends on exogenous local soil and weather characteristics, we think it is reasonable argue that its diffusion is exogenous to developments in the local industrial sector.

Second, these new technologies have a differential impact on potential yields depending on soil and weather characteristics. Thus, we exploit these exogenous differences in 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 maize, which we obtain from the FAO-GAEZ database. These potential yields are estimated by FAO 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. Crucially for our analysis, the database reports potential yields under different technologies or input combinations. Yields under the low technology are described as those obtained using traditional seeds and no use of chemicals, while yields under the high technology are obtained using improved seeds, optimum application of fertilizers and herbicides and mechanization. Thus, the difference in yields between the high and low technology captures the effect of moving from traditional agriculture to a technology that uses improved seeds and optimum weed control, among other characteristics. We thus expect this increase in yields to be a good predictor of the profitability of adopting herbicide resistant GE soy seeds.

More formally, our basic empirical strategy consists in estimating the following equation:

$$y_{jt} = \alpha_j + \alpha_t + \beta \ A_{jt}^{soy} + \varepsilon_{jt} \tag{9}$$

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_{jt}^{soy} is equal to the potential soy yield under high inputs from 2003 onwards and to the potential soy yield under low inputs in the years before 2003. A_{jt}^{soy} can be thought of as the empirical counterpart of the labor augmenting technical change A_L presented in our model.

In the case of agricultural outcomes, our period of interest spans the ten years between the last two censuses which took place in 1996 and 2006.[29] We thus estimate a first-difference version of equation 9:

$$\Delta y_j = \Delta\alpha + \beta\Delta A_j^{soy} + \delta \; Rural_{j,1991} + \Delta\varepsilon_{jt} \tag{10}$$

where the outcome of interest, Δy_j is the change in outcome variables between 1996 and 2006; ΔA_j^{soy} is the potential yield of soy under the high technology minus the potential yield of soy under the low technology. Figure 12 contains a map of Brazilian municipalities displaying this measure of technical change. Green municipalities are the ones where potential yields increase the most when switching to the high technology. As mentioned above, in municipalities with high urbanization rates we expect agricultural outcomes to be less affected by agricultural technical change. To address this potential source of attenuation bias, we control for the share of rural population in 1991.

In the case of maize, we follow a similar strategy. As noted in Section 2, the cultivation of second harvest maize requires the use of modern techniques that are intensive in the use of fertilizers, herbicides and tractors. Thus, we expect that the the difference in FAO-GAEZ potential yields between the high and low technology captures the profitability of introducing a second harvesting season for maize. Thus, we augment the equation described above to include the following variable: A_{jt}^{maize} which is equal to the potential maize yield under high inputs from 2003 onwards and to the potential maize yield under low inputs in the years before 2003. A_{jt}^{maize} can be thought of as the empirical counterpart of the land augmenting technical change A_T presented in our model.

$$\Delta y_j = \Delta\alpha + \beta\Delta A_j^{soy} + \gamma\Delta A_j^{maize} + \delta \; Rural_{j,1991} + \Delta\varepsilon_j \tag{11}$$

where ΔA_j^{maize} is the potential yield of maize under high inputs minus the potential yield of maize under low inputs.

A potential concern with our identification strategy is that, although the soil and weather characteristics that drive the variation in $\Delta A_j^{soy}$and ΔA_j^{maize} across geographical areas are exogenous, they might be correlated with initial levels of development across Brazilian municipalities. For example, to the extent that municipalities with heterogeneous initial levels of development experience different growth paths, our estimates could be capturing differential structural transformation trends across municipalities. To assess the extent of this potential concern we first compare observ-

[29] Recall that in the case of sectoral employment shares and manufacturing outcomes, our period of analysis spans the ten years between the last two population censuses which took place in 2000 and 2010.

able characteristics of municipalities with high and low levels of our exogenous measure of technical change in agriculture. Whenever significant differences emerge, we show that our estimates are stable when we introduce controls for differential trends across municipalities with heterogeneous initial characteristics.

Table 6 compares municipalities above and below the median change in potential soy yields (ΔA_j^{soy}) in terms of observable characteristics in 1991, before the introduction of GE soy.[30] Municipalities above the median potential increase in soy yields are characterized by smaller shares of rural population and agricultural employment. In addition, they display a larger manufacturing employment share, literacy rate, and income per capita than municipalities below the median. Thus, in what follows, we always show that our estimates are stable when we introduce controls for differential trends across municipalities with heterogeneous initial characteristics in our baseline specification 11, as follows:

$$\Delta y_j = \Delta\alpha + \beta\Delta A_j^{soy} + \gamma\Delta A_j^{maize} + \delta \; Rural_{j,1991} + \theta X_{j,1991} + \Delta\varepsilon_j \tag{12}$$

where $X_{j,1991}$ are the set of municipality characteristics discussed above.

In the following subsections we report the results obtained using our measure of technical change to explain changes in agricultural production and in the sectoral composition of the economy. Section 5.2.1 reports the relationship between our measure of technical change and the expansion of soy and maize cultivation. Section 5.2.2 shows the relationship between this measure and agricultural outcomes. Finally, section 5.2.3 presents results using manufacturing outcomes.

5.2.1 Agricultural Outcomes: Soy and Maize Expansion

In this section we document the relationship between technical change measured by the increase in the FAO-GAEZ potential yields of soy and maize, and the actual change in agricultural area cultivated with each crop. The objective of this exercise is to check whether the change in potential yields is a good proxy of the profitability of the adoption of the new agricultural technologies. If this is the case, we expect the increase in the potential yield of a given crop to predict the actual expansion in the area cultivated with that crop between 1996 and 2006.

First, we expect that areas with a higher increase in potential soy yields when switching to the high technology are those adopting genetically engineered soy on a larger scale. Thus, we start by estimating equation 10 where the outcome of interest, Δy_j is the change in the share of agricultural

[30] Municipalities below the median level of ΔA_{jt}^{soy} experience, on average, a 1.06 tons per hectare increase in potential soy yield, while those with above the median experience a 2.5 tons per hectare increase.

land devoted to GE soy between 1996 and 2006. Note that because this share was zero everywhere in 1996, the change in the area share corresponds to its level in 2006. Estimates are shown in column 1 of Table 7: the increase in potential soy yield predicts the expansion in GE soy area as a share of agricultural area between 1996 and 2006. The point estimate remains stable when controlling for initial municipality characteristics, as shown in column 2.

In columns 3 and 4 of Table 7 we perform a falsification test by looking at whether our measure of technical change in soy explains the expansion in the area planted with non-GE soy. In this case, the coefficients are negative and significant. This finding supports our claim that the change in potential soy yield captures the benefits of adopting GE soy vis-à-vis traditional soy seeds.

Next, we jointly analyze the effects of technical change in soy and maize on the area planted with each crop. For this purpose, we use the broader measure of planted area with soy instead of GE soy.[31] This permits to control for municipality fixed effects by focusing on changes in area planted rather than levels. We start by estimating equation 12 where the outcome of interest, Δy_j is the change in share of agricultural land devoted to either soy or maize between 1996 and 2006. Estimates are reported in Table 8. First, note that while soy technical change has a positive effect on the are planted with soy (column 1), it does not have a significant effect on the area planted with maize (column 4). Similarly, maize technical change only has a positive effect on the area planted with maize (columns 2 and 3). These findings suggest the change in potential yields when switching to the high technology are good measures of crop-specific technical change in soy and maize during this period. In addition, both estimates are stable when we add controls for municipality characteristics. This finding suggests that the differential expansion of these crops across municipalities is not driven by differential trends across municipalities with different initial levels of development.

The size of the estimated coefficient on ΔA_j^{soy} implies that a one standard deviation increase in potential soy yield corresponds to a 1.1 percentage points increase in the share of soy in agricultural land (25% of a standard deviation). To understand the magnitude of our estimate, this is an increase of agricultural land devoted to soy by 816 hectares in response to a 0.85 tons per hectare increase in potential soy yield.

The size of the estimated coefficient on ΔA_j^{mze} implies that a one standard deviation increase in potential maize yield corresponds to a 0.5 percentage points increase in the share of maize in agricultural land (8% of a standard deviation). This means that, in response to a 1.8 tons per hectare increase in potential maize yield, agricultural land devoted to maize increases by 396

[31] In the case of maize, we can only focus on the broader measure of area planted with maize as the publicly available Agricultural Census data does not contain information on the season of planting of maize at the municipality level.

hectares.

5.2.2 Agricultural Outcomes: Productivity, Labor Intensity and Employment Share

In this section we study the effects of agricultural technical change on productivity and employment in agriculture. Table 9 reports the results of estimating equation (12) when the dependent variables are three agricultural outcomes: the value of agricultural production per worker, labor intensity, and the share of workers employed in agriculture, all defined as in section 5.1.1.

The estimated coefficients on ΔA_j^{soy} indicate that areas with higher increase in potential soy yield experienced a larger increase in the value of agricultural production per worker and a larger reduction in labor intensity between 1996 and 2006 (columns 1 and 3). This findings confirm our characterization of technical change in soy as land-biased. Next, we study the effect of agricultural technical change in soy production on sectoral employment shares, starting from agriculture. Employment shares are calculated using the population Census. Therefore, our estimation of equation (12) relates changes in soy and maize potential yields to changes in employment shares between 2000 and 2010. The estimated coefficient on ΔA_j^{soy} indicates that areas with higher increase in potential soy yield experienced a reduction in agricultural employment share (column 5). This finding suggests that technical change in soy is not only land-biased but also strongly labor saving. Taken together, the results confirm the conclusions drawn from the simple correlations in the data reported in Table 4. Finally, let us note that the estimated coefficients on ΔA_j^{soy} are larger in absolute value in the specifications that control for lagged municipality characteristics (columns 2, 4 and 6). This finding confirms that our estimates are not capturing differential growth trends across municipalities.

These estimates can be used to compute the elasticity of the agricultural employment share to changes in agricultural labor productivity due to GE soy adoption. We compute this elasticity as the ratio of the estimated coefficient on ΔA_j^{soy} when the outcome is agricultural employment share, and the estimated coefficient on ΔA_j^{soy} when the outcome is agricultural labor productivity. Using our more conservative estimates – i.e. those that include all municipality controls in columns 2 and 6 – this ratio is equal to: $-0.018/0.132 = -0.136$.[32] The size of this elasticity implies that a 1 log point increase in agricultural labor productivity corresponds to a 0.136 percentage points decrease in agricultural employment share. To illustrate the magnitude of these estimates, we compute how much of the differences in the speed of structural transformation across Brazilian regions can soy technical change explain, as follows. Note that a municipality shocked with a one standard

[32] We compute this elasticity in the same way we would compute a Wald estimator in an instrumental variable setting, where the estimated coefficient on ΔA_j^{soy} in column 2 is the first stage coefficient, and the estimated coefficient on ΔA_j^{soy} in column 6 is the reduced form coefficient.

deviation increase in potential soy yield experienced an increase in agricultural labor productivity of 11 log points ,[33] and a corresponding 1.52 percentage points decrease in agricultural employment share.[34] This estimate corresponds to 21% of a standard deviation in the change of agricultural employment share between 2000 and 2010 (7.3 percentage points, see Table 3).

In the case of maize, the estimated coefficients on ΔA_j^{mze} indicate that areas with higher increase in potential maize yield experienced a smaller increase in the value of agricultural production per worker and a larger increase in labor intensity between 1996 and 2006 (columns 1 and 3). The first result is not consistent with the simple correlations in the data reported in Table 4: municipalities where the area planted with maize expanded experienced an increase in the value of agricultural output per worker. This could respond to the fact that in many municipalities affected by technical change in maize the area planted did not expand but the same area was harvested twice (see Figure 7). Thus the coefficients on the expansion of the area planted in maize reported in Tables 4 and the coefficient on the increase in potential maize yields reported in Table 9 capture different variation. In contrast, the finding that labor intensity increased is consistent with the simple correlations in the data reported in Table 4 and confirms our characterization of technical change in Maize as labor-biased. Finally, we find that areas more affected by technical change in maize production had a faster increase in the employment share of agriculture.

The results presented in Table 9 suggest that the introduction of new agricultural technologies in Brazil had a sizable impact on agricultural labor markets. Areas where the potential impact of GE soy adoption was higher experienced an increase in the value of agricultural production per worker, a reduction in the number of workers per unit of land, and a reduction in the employment share of agriculture. These findings are consistent with our characterization of the adoption of GE soy as a land-biased technical change. In the case of maize, areas where the potential impact of the introduction of a second harvesting season was higher experienced an increase in labor intensity and in the employment share of agriculture. This results are consistent with our characterization of the introduction of a second harvesting season as labor-biased technical change.

5.2.3 Manufacturing Outcomes: Employment Share, Employment and Wages

In this section we study the effect of agricultural technical change on manufacturing employment and wages. Table 10 reports the results of estimating equation (12) where the dependent variables

[33] This number is computed multiplying one standard deviation in ΔA_j^{soy} by the estimated coefficient on ΔA_j^{soy} in our specification with municipality controls when the outcome is agricultural labor productivity (column 2 of Table 9): $0.851 \times 0.132 = 0.112$

[34] This number is computed multiplying the predicted increase in agricultural labor productivity for one standard deviation in ΔA_j^{soy} by the elasticity of agricultural employment share to agricultural labor productivity: $0.112 \times -0.136 = -0.0152$.

are three manufacturing outcomes: the employment share of manufacturing, the level of manufacturing employment, and the average wage in manufacturing, all defined as in section 5.1.2. Manufacturing outcomes are calculated using the Population Census. Therefore, our estimation of equation (12) relates changes in soy and maize potential yields with changes in manufacturing outcomes between 2000 and 2010.

The estimates indicate that areas where potential soy yields increased relatively more, experienced a larger increase in the manufacturing employment share. A comparison of point estimates reported in the first row of columns 1 and 2 shows that estimates are stable when introducing controls for lagged municipality characteristics. In addition, columns 3 and 4 show that not only the share of manufacturing employment increased but also its absolute level. Finally, columns 5 and 6 show that manufacturing wages fall. These estimates are consistent with the empirical predictions of our model: technical change in soy is strongly labor saving thus reduces labor demand in agriculture, and induces an expansion of the manufacturing sector through an increase in labor supply and lower wages.

These estimates can be used to compute the elasticity of manufacturing employment share to changes in agricultural labor productivity due to GE soy adoption. We compute this elasticity as in section 5.2.2: we divide the estimated coefficient on ΔA_j^{soy} when the outcome is manufacturing employment share by the estimated coefficient on ΔA_j^{soy} when the outcome is agricultural labor productivity. When we estimate the specification including controls for lagged municipality characteristics, this ratio is equal to: $0.019/0.132 = 0.146$. This elasticity implies that a 1% increase in agricultural labor productivity corresponds to a 0.146 percentage points increase in manufacturing employment share. As in the previous section, we illustrate the magnitude of these estimates by computing how much of the differences in the speed of structural transformation across Brazilian regions can be explained by technical change in soy. Recall that a municipality shocked with a one standard deviation increase in potential soy yield experienced an increase in agricultural labor productivity of 11 log points,[35] and a corresponding 1.64 percentage points increase in manufacturing employment share.[36] This estimate corresponds to 31% of a standard deviation in the change of manufacturing employment share between 2000 and 2010 (5.4 percentage points, see Table 3).

In the case of maize, the estimates reported in columns 1 and 2 of Table 10 indicate that areas where potential maize yields increased relatively more experienced a smaller increases in the

[35] This number is computed as before, by multiplying one standard deviation in ΔA_j^{soy} by the estimated coefficient on ΔA_j^{soy} in our specification with municipality controls when the outcome is agricultural labor productivity (column 2 of Table 9): $0.851 \times 0.132 = 0.112$

[36] This number is computed multiplying the predicted increase in agricultural labor productivity for one standard deviation in ΔA_j^{soy} by the elasticity of manufacturing employment share to agricultural labor productivity: $0.112 \times 0.146 = -0.0164$.

manufacturing employment share. In addition, columns 3 and 4 show that not only the share of manufacturing employment fell but also its absolute level. Finally, columns 5 and 6 show that manufacturing wages increased. These estimates are consistent with the empirical predictions of our model: technical change in maize is labor-biased thus increases labor demand in agriculture, generating an increase in wages and a reallocation of labor away from the manufacturing sector.

Taken together, the estimates reported in this section are consistent with the empirical predictions of our model. They show that the effects of agricultural productivity on the industrial sector depend on the factor bias of technical change. In the case of soy, our estimates indicate that strongly labor saving technologies (like GE soy seeds), by reducing the demand for labor in agriculture, promote the growth of the manufacturing sector through an increase in labor supply and lower wages. On the other hand, in the case of maize, our estimates show that land-augmenting technical change (like the introduction of a second harvesting season), by increasing the labor intensity of agriculture, result in a decrease of manufacturing employment and increasing wages.

Finally, let us note that the estimates of the effects of technical change on the agricultural and manufacturing employment shares discussed above have a similar magnitude. To make this point clearer, we reproduce them in Table 11, where we also include estimates for the service and public sectors. The point estimates of the effect of soy technical change on the agriculture and manufacturing employment shares have similar size: they are -0.018 and -0.019, respectively, with a standard error of 0.002. At the same time, the estimates of the effects on the service and public sectors are very small and not statistically different from zero. This implies that labor reallocated only from agriculture to manufacturing.

The absence of an effect of agricultural technical change on services is, at first sight, puzzling. A simple extension of our model to include non traded sectors would imply that as higher productivity increases income, the demand for all goods increases. Because non traded goods need to be produced locally, some labor must reallocate from agriculture towards the service sector. Our ongoing empirical work on this matter indicates that the absence of an average effect on services is related to the asymmetric effects of soy technical change on the income of land owners and workers and the fact that in some areas of Brazil land owners do not reside locally.

6 Robustness Checks

6.1 Variable Factor Endowments

The model presented in Section 3 describes a small open economy where goods can be freely traded but factor endowments are fixed. Our empirical strategy thus relies on the assumption that each unit of observation behaves as a small open economy: goods can be traded across municipalities but labor markets are local and there is a fixed supply of land. However, Brazil was characterized by significant internal migration flows during the last decade.[37] In addition, Brazil has vast areas of underutilized land, which were in part converted to agricultural activities during the period under study.[38] Thus, in this section, we investigate the role of migration and the expansion in the agricultural frontier.

We first investigate the impact of agricultural technical change on migration flows. For this purpose, we estimate net migration rates for every municipality between 2000 and 2010 using data from the population Census.[39] Next, we estimate the baseline specification described by equation 12 using the 2000-2010 net migration rate in each municipality as dependent variable. Estimation results are presented in the first column of Table 12. The estimated coefficient on the change in soy potential yields is negative and significant, indicating that municipalities with higher increases in potential soy yields experienced a net outflow of migrants between 2000 and 2010. This evidence is consistent with the model's predictions: municipalities more affected by labor saving technical change (GE soy) experienced a larger contraction in labor demand and as a consequence some workers migrated to other municipalities. In addition, the estimated coefficient on the change in maize potential yields is positive and significant, indicating that municipalities with higher increase in potential maize yield experienced a net inflow of migrants in the same period. This last finding is also consistent with the model. Taken together, this findings suggest that the presence of migration flows across municipalities dampen the effects of technical change on sectoral employment shares, as part of the adjustment occurs through migration flows.

Next, we study the role of the expansion in the agricultural frontier. During this period the frontier expanded not only over the Amazon rainforest but also in the Cerrado. This is a tropical savanna eco-region in central Brazil where soils used to be too acidic and nutrient poor. In the late 1990s these soils were treated by the Brazilian Agricultural Research Corporation, EMBRAPA, which enabled agricultural activities to expand over these areas. To assess the extent to which our estimated effects of agricultural technical change on structural transformation are affected by expansions in the land endowment we proceed as follows. First, we define frontier municipalities as those which experienced an increase in land use for agricultural activities between 1996 and 2006 and split the sample of municipalities in two groups: frontier and non-frontier (see map in Figure 16). Next, we estimate our baseline specification described by equation 12 separately for each

[37] In 2010 15% of the population aged between 15 and 70 years old had moved to their current municipality during the previous 10 years.

[38] Between 1996 and 2006 the land used for cultivation or cattle ranching increased by 7% to 154 million hectares in the regions of the North, North-East and Center-West.

[39] A detailed explanation of how net migration rates are constructed is contained in Appendix B.

subsample. Our estimates of the effect of soy technical change on the agricultural and manufacturing employment shares in the subsample of non-frontier (frontier) municipalities are slightly larger in absolute value (smaller) than estimates using the full sample, as shown in Columns 4 to 7 of Table 12. This finding is consistent with the predictions of the model: the expansion of the agricultural frontier increases the land endowment, partially mitigating the reduction in agricultural labor demand due to strongly labor saving technical change. In the case of maize, estimates of the effect of technical change on the agricultural and manufacturing employment shares in the subsample of non-frontier municipalities are slightly larger in absolute value than estimates using the full sample but estimates are smaller and not statistically significant in the frontier. These findings are also consistent with the model's predictions for the effects of labor biased technical change.

Finally, we study whether migration patters differ in frontier and non frontier municipalities. Columns 2 and 3 of Table 12 show that the effect of soy technical change on migration is similar for both samples. In contrast, the positive effect of maize technical change on migration is concentrated in non-frontier municipalities.

6.2 Pre-Existing Trends

In this section we show that our results are robust to controlling for pre-existing trends. One concern is that municipalities that are better suited for adopting GE soy were already experiencing faster structural transformation before the legalization of this technology in Brazil. If that is the case, our exogenous measure of technical change would capture a long term trend instead of the effect of GE soy adoption.

In order to test for the existence of pre-existing trends, we use data from the Population Censuses of 1980, 1991, 2000 and 2010. We thus estimate a model similar to the one presented in our baseline equation 12, but with an additional time period, as follows.

$$\begin{aligned}\Delta y_{jt} \quad = \quad & \alpha_t + \beta_1 \Delta A_j^{\text{soy}} + \beta_2 \Delta A_j^{\text{soy}} \times After_t + \gamma_1 \Delta A_j^{\text{maize}} + \\ & \gamma_2 \Delta A_j^{\text{maize}} \times After_t + \theta X_{jt-1} + \Delta\varepsilon_{jt} \end{aligned} \tag{13}$$

where the outcome of interest, Δy_{jt} is the decadal change in outcome variables between the start of a period (year t) and the end. Each period spans a decade: 1991 to 2000 and 2000 to 2010 and we and include separate time dummies for each decade (α_t). $After_t$ is a time dummy equal to 1 if $t = 2000$. Thus, β_1 captures the effect of soy technical change that is common in the period before (1991-2000) and after (2000-2010) the adoption of GE soy seeds. I contrast, β_2 captures the differential effect of soy technical change after the introduction of GE soy seeds. Similarly, the coefficient γ_2 captures the differential effect of maize technical change in the period 2000-2010. Finally, X_{jt-1} are a set of ten-year-lagged municipality characteristics including the share of rural population, average income per capita, population density and literacy rate. [40]

We perform this test for two manufacturing outcomes: manufacturing employment level and manufacturing wages. We do not perform this test for manufacturing and agriculture employment shares since, due to important changes in the definition of employment introduced by the IBGE after the 1991 Census, we are unable to measure employment shares in a consistent way across the 1991 and 2000 Censuses.[41] Results for manufacturing employment are reported in column 1 of Table 13. The coefficient on ΔA^{soy} estimates the effect of soy technical change that is common in the period before (1991-2000) and after (2000-2010) the adoption of GE soy seeds (β_1). That this coefficient is very small and not statistically different from zero indicates that there are no pre-existing trends in manufacturing employment. In contrast, the coefficient on the interaction term between ΔA^{soy} and $After_t$, which estimates the differential effect of soy technical change on manufacturing employment after the introduction of GE soy seeds (β_2) is positive and precisely estimated. Similarly, in the case of maize, we do not find pre-existing trends in manufacturing employment.

Column 2 of Table 13 shows the results of estimating equation 13 when the outcome variable is the average wage in manufacturing. In this case, ΔA^{soy} had an opposite effect on manufacturing wages between 1991 and 2000 with respect to the 2000-2010 period. Therefore, the existence of these pre-existing trends in manufacturing wages attenuates our baseline estimated effects of soy and maize technical change on wages in the period 2000-2010, presented in Table 10.

Finally, we check for pre-existing trends in migration. A potential concern is that areas that

[40] The municipality characteristics correspond to the year 1991 when the outcome variables are observed in changes between 2000 and 2010, and to year 1980 when the outcome variables are observed in changes between 1991 and 2000.

[41] Between the 1991 and 2000 Censuses the Brazilian Statistical Institute (IBGE) changed its definition of employment in two important ways. First, it started to count zero-income workers as employed. In order to homogenize the Brazilian Census with international practices, the IBGE started to consider employed anyone who helped another household member with no formal compensation, as well as agricultural workers that produced only for their own consumption (IBGE, 2003; p. 218). Zero-income workers are more common in agriculture than in other sectors, and in 1991 were only partially included in the labor force. In the 1991 Census 15% of agricultural workers reported zero income, against 34% in 2000 and 35% in 2010. Second, the IBGE changed the reference period for considering a person employed: while in 1991 such period included the last 12 months, in 2000 it only included the reference week of the Census. This new rule implied that workers performing temporary and seasonal activities that were not employed during the reference week were counted in the 1991 census but not the in the 2000 census. This second change is likely to be especially problematic for the agricultural sector, considering that the reference week in the 2000 Census was in the middle of the Brazilian winter. This is why, to test for pre-existing trends, we focus on the absolute number of workers employed in manufacturing as an outcome (instead of its share in total employment). This measure is less likely to be affected by the changes introduced between the two censuses because: (1) there are very few zero-income workers in manufacturing (0.5%, 1.9% and 1% of manufacturing workers declare zero income in 1991, 2000 and 2010, respectively) and (2) manufacturing is less seasonal than other sectors.

are better suited for adopting GE soy experienced a pattern of migration prior to the legalization of GE soy that affected farmers' incentive to adopt this new technology. For example, if these areas experienced large out-migration in the decade before GE soy was legalized, farmers would have had a higher incentive to adopt a labor-saving technology to cope with labor scarcity. Column 3 of Table 13 shows the results of estimating equation 13 when the outcome variable is net migration rate. The coefficient on ΔA^{soy} shows that there are no differential pre-existing trends in migration for areas that have a higher increase in potential soy yields.

These tests validate our interpretation that the effect of our estimated effects of agricultural technical change on structural transformation are due to the introduction of new agricultural technologies rather than to pre-existing trends in areas that were more affected by these new technologies.

6.3 Larger Unit of Observation: Micro-Regions

In the empirical analysis performed so far we assumed that municipalities are a good approximation of the relevant labor market faced by Brazilian agricultural workers. A potential issue is that local labor market boundaries do not overlap with a municipality's administrative boundaries. In particular, some municipalities might be too small to properly capture labor flows between urban and rural areas, provided that manufacturing activities mostly take place in the former, and agricultural activities in the latter. In order to take into account this concern we aggregate our data at a larger unit of observation: micro-regions. These regions are groups of territorially contiguous municipalities created, mostly for statistical purposes, by the Brazilian Statistical Institute (IBGE). Table 14 reports the results of estimating equation 12 using micro-regions as a unit of observation. The outcome variables are the same as in Table 10: change in manufacturing employment share, change in manufacturing employment (in logs) and change in average manufacturing wage (in logs). The estimates are broadly consistent and similar in magnitude to those reported in Table 10, both for soy and maize.

6.4 Input-Output Linkages

Our theoretical model predicts that agricultural technical change can have an effect on manufacturing employment through labor market forces. In the case of soy, for example, the adoption of new agricultural technologies releases agricultural workers that find employment in the manufacturing sector. In this section we investigate to which extent our findings reflect the strength of another channel through which agricultural technical change can affect manufacturing employment:

input-output linkages.

Soy and maize farming require inputs produced by other sectors, including manufacturing. Therefore, for example, an expansion of area farmed with soy in a given municipality might mechanically drive an increase in manufacturing employment in industries that produce inputs used in soy production, such as chemicals or fertilizers. To the extent that manufacturing firms producing chemicals and fertilizers used in agriculture are located in the same municipality in which soy expanded, the effect of agricultural technical change on manufacturing that we show in Table 10 could be explained by an increase in the agricultural demand for manufacturing inputs. A similar argument applies for manufacturing industries that use soy and maize as intermediate inputs, such as the food processing industry.

In order to assess the contribution of these direct linkages on our estimates, we construct a measure of manufacturing employment that excludes the sectors directly linked to soy and maize production through input-output chains. We identify the relevant sectors using the Brazilian input-output matrix (IBGE, 2008). On the output side, soy and maize are used directly as intermediate inputs in only one manufacturing sector: the food and beverage sector, that in 2005 purchased roughly half of the total Brazilian production of both crops. On the input side, the input-output data is less detailed. Thus, we use information on the inputs used by agricultural and breeding farms. Almost half of the inputs purchased by these farms are supplied by manufacturing sectors (49%) and four industries account for 84% of the total value of inputs: inorganic chemicals, fertilizers, diesel oil and maize oil. We use this information to construct measures of employment and wages in manufacturing that exclude those industries that are providing inputs, or receiving outputs, from the soy and maize sectors.

Table 15 reports estimates of our baseline specification described by equation (12) using as outcome variables this narrower definition of manufacturing employment and wages. When we exclude workers employed in sectors directly linked to soy and maize, the estimated coefficients on ΔA^{soy} tend to decrease in size with respect to our baseline estimates presented in Table 10. However, the estimated coefficients on ΔA^{soy} still account for roughly two-thirds of the effect on manufacturing employment share and for almost 90% of the effect on manufacturing employment shown in Table 10.[42] The effect of technical change in soy on manufacturing wages, instead, decreases substantially, and it is not precisely estimated. In the case of maize, the estimated coefficients on ΔA^{maize} are essentially unaffected by excluding workers in downstream and upstream

[42] In our specification with all initial municipality controls, the point estimate on ΔA^{soy} when the outcome is manufacturing employment share goes from 0.019 to 0.012. We can reject the null hypothesis that these two coefficients are equal. When the outcome is manufacturing employment instead, the point estimate on ΔA^{soy} goes from 0.174 to 0.154. In this case, the two coefficients are not statistically different.

manufacturing sectors when the outcomes are manufacturing employment share and level. As in the case of soy, the effect on manufacturing wages decreases in size and it is not precisely estimated.

Taken together, the results presented in this section imply that at least 2/3 of our estimated effect of agricultural technical change on the manufacturing employment share is not driven by the processing of soy and maize in downstream industries nor larger agricultural sector demand for manufacturing inputs. A more detailed analysis is needed to separate the role of labor market and input-output forces in the remaining 1/3 of the total estimated effect, which we think is an interesting avenue for further work.

6.5 Commodity Prices

In this section we show that our results are robust to controlling for international commodity prices. To the extent that variation in international prices of soy and maize affect agricultural outcomes in all Brazilian municipalities proportionally, their effects are captured by the constant term in equation (12). However, price changes might have heterogeneous effects across municipalities with different suitability to the cultivation of soy and maize. For example, an increase in the international price of soy could induce farmers to expand the area devoted to soy relatively more in municipalities that are initially more suitable for its cultivation.

Figures 14 and 15 display the evolution of international prices of soy and maize, expressed in 2000 US$. These Figures show how the international prices of both commodities have been in an upward trend starting from year 2007. This pattern most likely does not affect our estimates when we use variation across the last two Agricultural Censuses, i.e. 1996 and 2006. If anything, for both soy and maize, their international price level in 2006 was lower than in 1996. However, when we use variation across the last two Population Censuses, i.e. between 2000 and 2010, the end of period year is characterized by high international soy and maize prices with respect to the initial year.

The data from the Population Censuses does not allow us to control for yearly variation in soy and maize prices. We therefore rely on an alternative source of data for manufacturing outcomes: the Annual Manufacturing Survey (PIA). The Annual Manufacturing Survey is carried out yearly, allowing us to both exclude years of high international commodity prices and fully control for price variation. It covers the universe of manufacturing firms with at least 30 employees in Brazil, and it is therefore representative at municipality level for this class of firms. We focus on two variables from this survey: manufacturing employment and average wages.[43]

[43] The average wage is defined as the aggregate wage bill (in real terms) divided by the total number of workers employed in a municipality.

We estimate an equation of the following form:

$$y_{jt} = \alpha_j + \alpha_t + \beta A^{soy}_{jt} + \gamma A^{maize}_{jt} + \sum_z \gamma_z P^z_t A^z_{j0} + \delta Rural_{j1991} \times t + \theta X_{jt1991} \times t + \varepsilon_{jt} \qquad (14)$$

where y_{jt} is total employment or average wage in a given municipality; A^{soy}_{jt} is equal to the potential soy yield under low inputs for all years before 2003 and to the potential soy yield under high inputs starting from 2003 (same criteria is used to define A^{maize}_{jt}). We control for the prices of soy and maize by multiplying the potential yield under low inputs of each crop by the time varying international price of each crop. Finally, we add as controls the share of rural population and the same set of initial municipality characteristics used in our main specification, all interacted with a time trend. In all specifications we control for both municipality and year fixed effects (α_j and α_t) and cluster standard errors at the municipality level to address potential serial correlation in the error term.

The results obtained using data from the Annual Manufacturing Survey are consistent with those obtained using the Population Census (see Table 16): areas with higher increase in potential soy yield experienced a larger increase in manufacturing employment and a larger decrease in average manufacturing wages. The effect on wages is less precisely estimated than in Table 10, and it loses statistical significance when we add all controls. Importantly, when we control for differential effects of international prices in columns 2 and 5, our point estimates do not change. In terms of magnitude, the point estimates we obtain with this specification for the coefficients on both ΔA^{soy} and ΔA^{maize} are similar to those obtained with the same outcomes using the Population Census data.

6.6 Spatial Correlation

The maps we present in Figures 8 through 13 suggest that the potential yield of both soy and maize are correlated across space. Therefore, in this section we show that our estimates remain significant when we allow the residuals to be correlated within geographical areas larger than a single municipality.

For the regressions in Tables 8, 9 and 10, we compute standard errors clustered at two additional levels of aggregation: *micro-regions* and *meso-regions*.[44] Tables 17, 18 and 19 report the coefficients of A^{soy} and A^{maize} showed in Tables 8, 9 and 10, respectively. The first row below the coefficients

[44] Both micro-regions and meso-regions are statistical divisions of Brazil proposed by the *IBGE* to facilitate the collection of data. There are 558 micro-regions and 137 meso-regions. We do not present results with standard errors clustered at *Federal Units* level (the highest level of aggregation) because there are only 27 Federal Units and clustered standard errors are inconsistent when the number of cluster is lower than 50.

reports the baseline robust standard errors for comparison. The second and third row below these coefficients report the standard errors clustered at micro and meso-region level, along with their significance levels. Table 17 show how the effect of the change in potential soy yield on the expansion of soy cultivation becomes less precise after clustering at micro and meso-region level, but remains by and large significant at least at the 5% level. Moreover, Tables 18 and 19 show that, in the case of soy, all our results for agriculture and manufacturing (with the sole exception of manufacturing wages when clustering at meso-region level) remain significant after clustering standard errors to allow for spatial correlation within larger units of observation.

7 Final Remarks

The process of modern economic growth is accompanied by structural transformation, i.e. the reallocation of economic activity from agriculture to industry. Identifying 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 is one important force behind structural transformation. However, as underlined by Matsuyama (1992), in open economies a comparative advantage in agriculture could instead slow down industrial growth. Despite the importance of the question, there is so far scarce evidence on the channels through which agricultural productivity can shape the reallocation of economic activity across sectors in an open economy.

In this paper we argue that the effect of agricultural productivity on industrial development depends crucially on the factor-bias of technical change. In particular, predictions of models of structural transformation in open economies hold when technical change is hicks neutral or labor-biased, but are reversed when land-biased technical change is strongly labor saving.

We provide direct empirical evidence on these mechanisms by isolating the effects of adoption of two new agricultural technologies in Brazil: genetically engineered soybean seeds and a second harvesting season for maize. We argue that the first technical change is land-biased, while the second is labor-biased, and exploit this setup to study the effect of the diffusion of these agricultural technologies on the manufacturing sector. To identify the causal effects of this new technology, we exploit the timing of adoption and the differential impact of the new technologies on potential yields across geographical areas.

We find that in municipalities where the new technology had a larger potential impact on soy yields, there was faster GE soy adoption, a reduction of labor intensity in agriculture and an expansion of manufacturing employment. In contrast, in municipalities where the new technology had a larger potential impact on maize yields, there was an increase of labor intensity in agriculture and a contraction of manufacturing employment. These different effects documented for 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

Acemoglu, D. (2010). "When Does Labor Scarcity Encourage Innovation?". *Journal of Political Economy 118*, 1037–1078.

Acemoglu, D. and V. Guerrieri (2008). "Capital Deepening and Non-Balanced Economic Growth". *Journal of Political Economy 116*, 467–498.

Baumol, W. J. (1967). "Macroeconomics of Unbalanced Growth: the Anatomy of Urban Crisis". *American Economic Review 57*(3), 415–426.

Bertrand, M., E. Duflo, and S. Mullainathan (2004). "How Much Should We Trust Differences-in-differences Estimates?". *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.

Caselli, F. and W. J. Coleman II (2001). "The US Structural Transformation and Regional Convergence: A Reinterpretation". *Journal of Political Economy 109*(3), 584–616.

Clark, C. (1940). *"The Conditions of Economic Progress"*. London: MacMillan & Co.

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

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.

Field (1978). "Sectoral Shifts in Antebellum Massachusetts: A Reconsideration". *Exploration in Economic History 15*, 146–171.

Foster, A. D. and M. R. Rosenzweig (2004). "Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India, 1970–2000". *Economic Development and Cultural Change 52*(3), 509–542.

Foster, A. D. and M. R. Rosenzweig (2008). "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". *American Economic Review 92*(2), 160–164.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2013). "Growth and Structural Transformation". *NBER Working Paper*.

Hornbeck, R. and P. Keskin (2012). "Does Agriculture Generate Local Economic Spillovers? Short-run and Long-run Evidence from the Ogallala Aquifer". *NBER Working Paper*.

Huggins, D. R. and J. P. Reganold (2008). "No-Till: the Quiet Revolution". *Scientific American 299*, 70–77.

IBGE (1996). *"Censo Agropecuário 1995-1996"*. Rio de Janeiro, Brazil: Instituto Brasileiro de Geografia e Estatística (IBGE).

IBGE (2003). *"Metodologia do Censo Demográfico"*, Volume 25. Rio de Janeiro, Brazil: Instituto Brasileiro de Geografia e Estatística (IBGE).

IBGE (2006). *"Censo Agropecuário 2006"*. Rio de Janeiro, Brazil: Instituto Brasileiro de Geografia e Estatística (IBGE).

Kongsamut, P., S. Rebelo, and D. Xie (2001). "Beyond Balanced Growth". *Review of Economic Studies 68*(4), 869–882.

Krugman, P. (1980). "Scale Economies, Product Differentiation, and the Pattern of Trade". *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 5*, 1–111.

Lewis, W. A. (1954). "Economic Development with Unlimited Supplies of Labour". *The Manchester School 22*(2), 139–191.

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

Michaels, G., F. Rauch, and S. J. Redding (2012). "Urbanization and Structural Transformation". *Quarterly Journal of Economics 127*(2), 535–586.

Mokyr, J. (1976). *"Industrialization in the Low Countries, 1795-1850"*. Yale University Press New Haven.

Murphy, K. M., A. Shleifer, and R. Vishny (1989). "Income Distribution, Market Size, and Industrialization". *Quarterly Journal of Economics 104*(3), 537–564.

Ngai, L. R. and C. A. Pissarides (2007). "Structural Change in a Multisector Model of Growth". *American Economic Review 97*(1), 429–443.

Nunn, N. and N. Qian (2011). "The Potato's Contribution to Population and Urbanization: Evidence From A Historical Experiment". *Quarterly Journal of Economics 126*(2), 593–650.

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

Rosenstein-Rodan, P. N. (1943). "Problems of Industrialisation of Eastern and South-Eastern Europe". *Economic Journal 53*, 202–211.

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

Schultz, T. (1953). *"The Economic Organization of Agriculture"*. New York: MacGraw-Hill.

Shryock, H. S., J. S. Siegel, and E. A. Larmon (1980). *The methods and materials of demography*, Volume 2. Department of Commerce, Bureau of the Census.

Syrquin, M. (1988). "Patterns of Structural Change". In H. Chenery and T. Srinivasan (Eds.), *Handbook of Development Economics*, Volume 1, Chapter 7, pp. 203–273. Amsterdam and New York: North Holland.

USDA (1996–2006). "Agricultural Resource Management Survey (ARMS)". United States Department of Agriculture, Economic Research Service.

USDA (2001). "Agriculture in Brazil and Argentina: Developments and Prospects for Major Field Crops". United States Department of Agriculture, Economic Research Service.

USDA (2012). "Agricultural Biotechnology Annual". United States Department of Agriculture, Economic Research Service.

Vogel, S. J. (1994). "Structural Changes in Agriculture: Production Linkages and Agricultural Demand-Led Industrialization". *Oxford Economic Papers 46*, 136–156.

Wright (1979). "Cheap Labor and Southern Textiles before 1880". *Journal of Economic History 39*, 655–680.

Figures and Tables

Figure 1
Employment in agriculture, industry, services and construction (1992-2009)

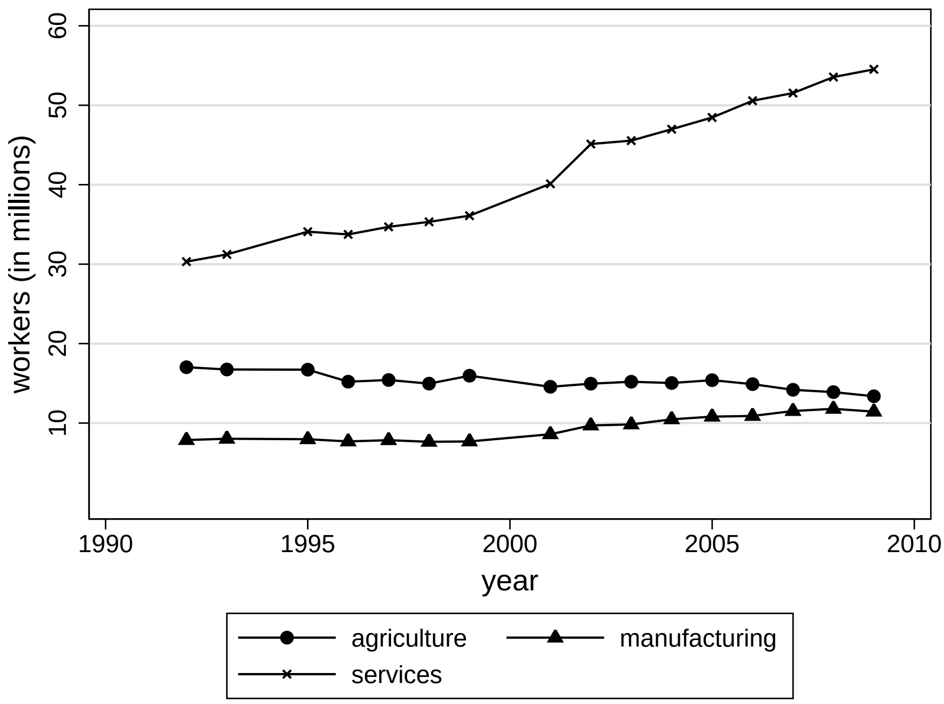

Notes: The Figure shows the evolution between 1992 and 2011 of the total number of workers (expressed in million) employed by sector in Brazil. The sectors are: Agriculture (codes A and B in the CNAE-Domiciliar classification), Industry (code D), Services (codes: F, G, H, I, J, K, L, M, N, O, P, Q, E). Workers in the Extractive Industry (code C) are not included in any of the categories above. Data source is PNAD, a national household survey representative at state level carried out yearly by the IBGE (the survey was not carried out in 1994 and in the census years: 1991, 2000 and 2010). Since the PNAD coverage changed over time, to harmonize the sample across years we exclude: (i) workers located in the states of: Rondonia, Acre, Amazonas, Roraima, Pará and Amapá (North macro-region) because only urban areas (and not rural areas) of these states were covered until 2004; (ii) workers located in the states of: Tocantins, Mato Grosso do Sul, Goias and the Distrito Federal because the sample of households in these states is not complete in the years from 1992 to 1997.

Figure 2
Distribution of soy yields across municipalities (1996-2006)

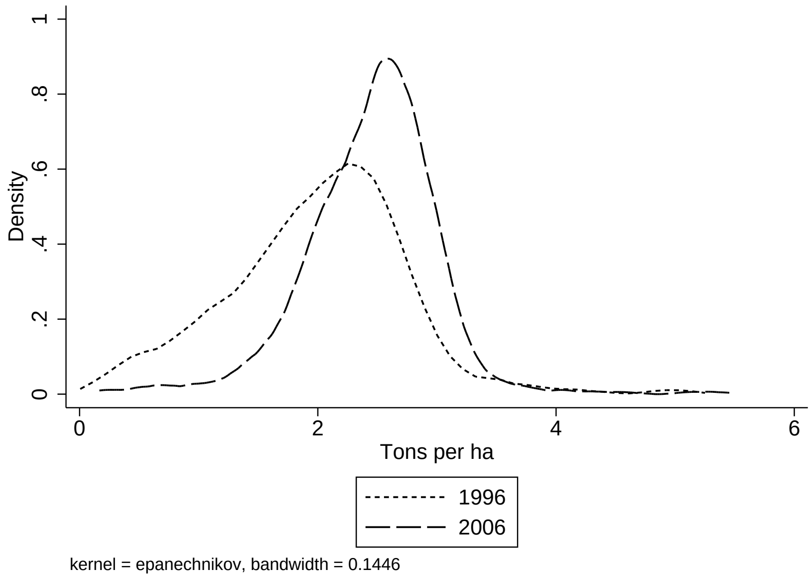

Notes: Soy yield is defined as tons per hectare farmed with soybean. Data sources are the Brazilian Agricultural Censi of 1996 and 2006, IBGE. See data appendix for details.

Figure 3
Distribution of maize yields across municipalities (1996-2006)

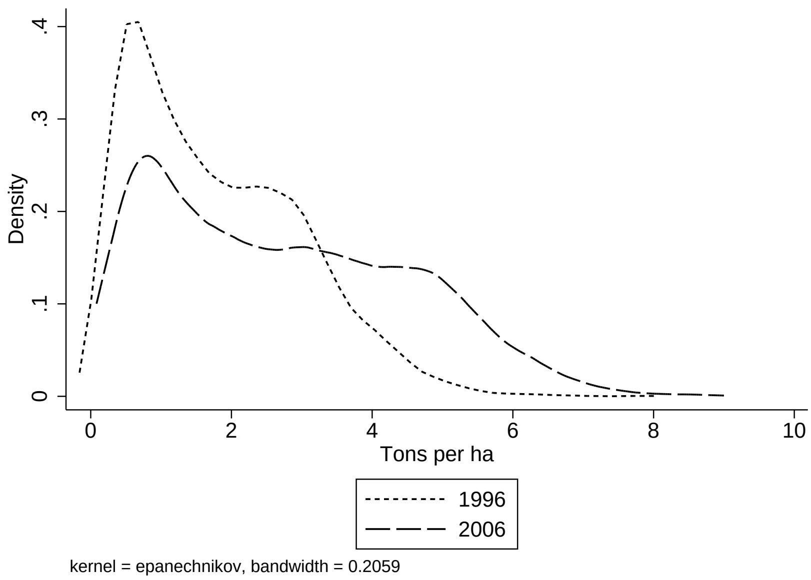

Notes: Maize yield is defined as tons per hectare farmed with maize. Data sources are the Brazilian Agricultural Censi of 1996 and 2006, IBGE. See data appendix for details.

Figure 4
Area planted with soy (1980-2011)

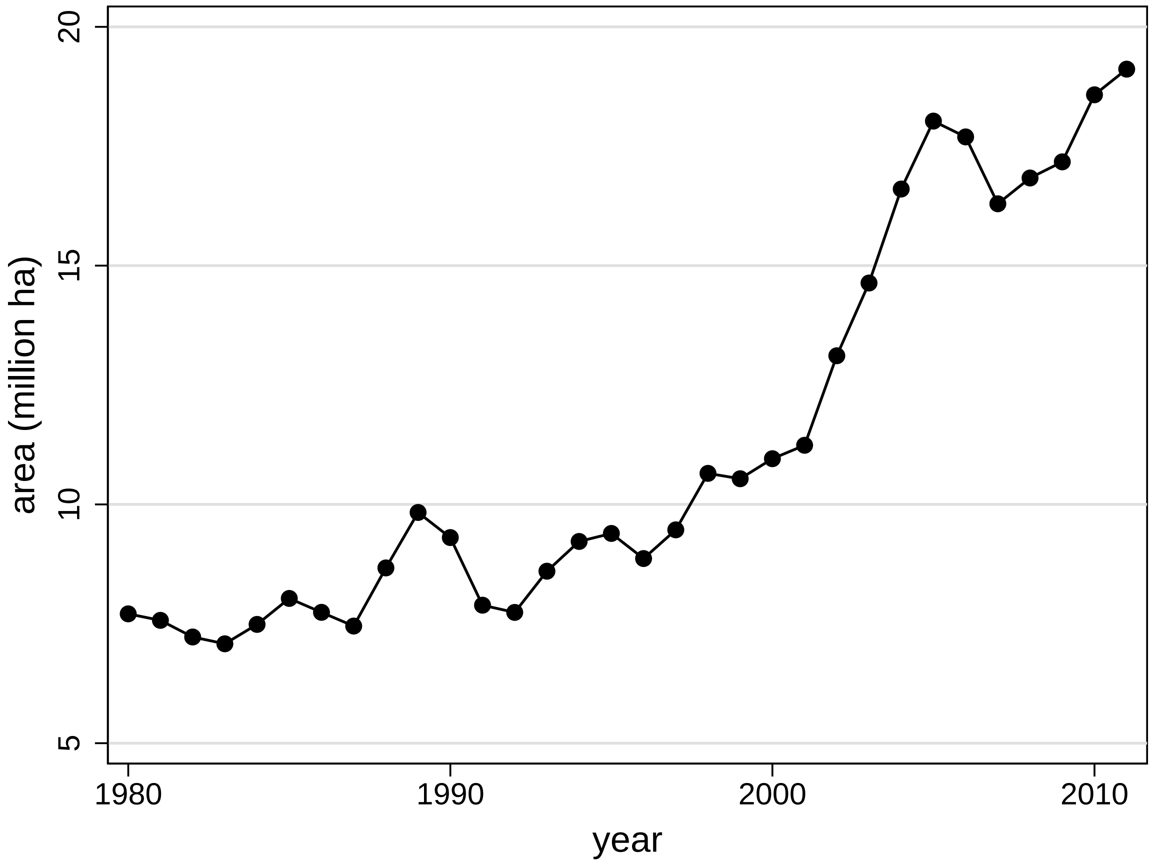

Notes: The Figure depicts the evolution between 1980 and 2011 of the total area planted with soy in Brazil (expressed in million hectares). Data source is CONAB, Companhia Nacional de Abastecimento, which is an agency within the Brazilian Ministry of Agriculture. CONAB carries out monthly surveys to monitor the evolution of the harvest of all major crops in Brazil: the surveys are representative at state level and are constructed by interviewing on the ground farmers, agronomists and financial agents in the main cities of the country. All data can be downloaded at: http://www.conab.gov.br/conteudos.php?a=1252&t=.

Figure 5
Area planted per worker in soy production (1980-2011)

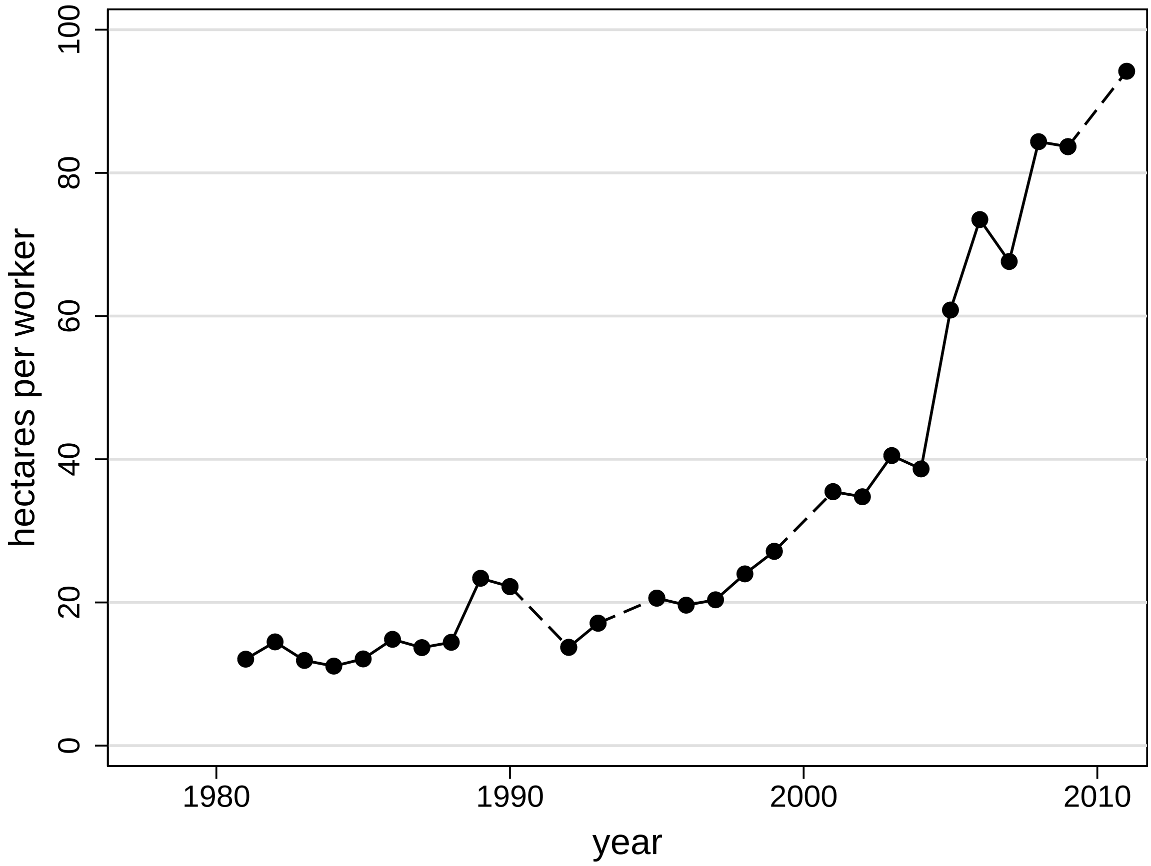

Notes: The Figure depicts the evolution between 1980 and 2011 of the area planted with soy divided by the total number of workers employed in soy production in Brazil. Data sources are CONAB for area planted with soy and PNAD for the total number of workers in soy production. CONAB, Companhia Nacional de Abastecimento, is an agency within the Brazilian Ministry of Agriculture. CONAB carries out monthly surveys to monitor the evolution of the harvest of all major crops in Brazil: the surveys are representative at state level and are constructed by interviewing on the ground farmers, agronomists and financial agents in the main cities of the country. All data can be downloaded at: http://www.conab.gov.br/conteudos.php?a=1252&t=. PNAD is a national household survey representative at state level carried out yearly by the IBGE (the survey was not carried out in 1994 and in the census years: 1991, 2000 and 2010). Since the PNAD coverage changed over time, to harmonize the sample across years we exclude: (i) workers located in the states of: Rondonia, Acre, Amazonas, Roraima, Pará and Amapá (North macro-region) because only urban areas (and not rural areas) of these states were covered until 2004; (ii) workers located in the states of: Tocantins, Mato Grosso do Sul, Goias and the Distrito Federal because the sample of households in these states is not complete in the years from 1992 to 1997. We harmonized data from CONAB with the PNAD coverage such that numerator and denominator are constructed using the same subset of states.

Figure 6
Employment in soy production (1980-2011)

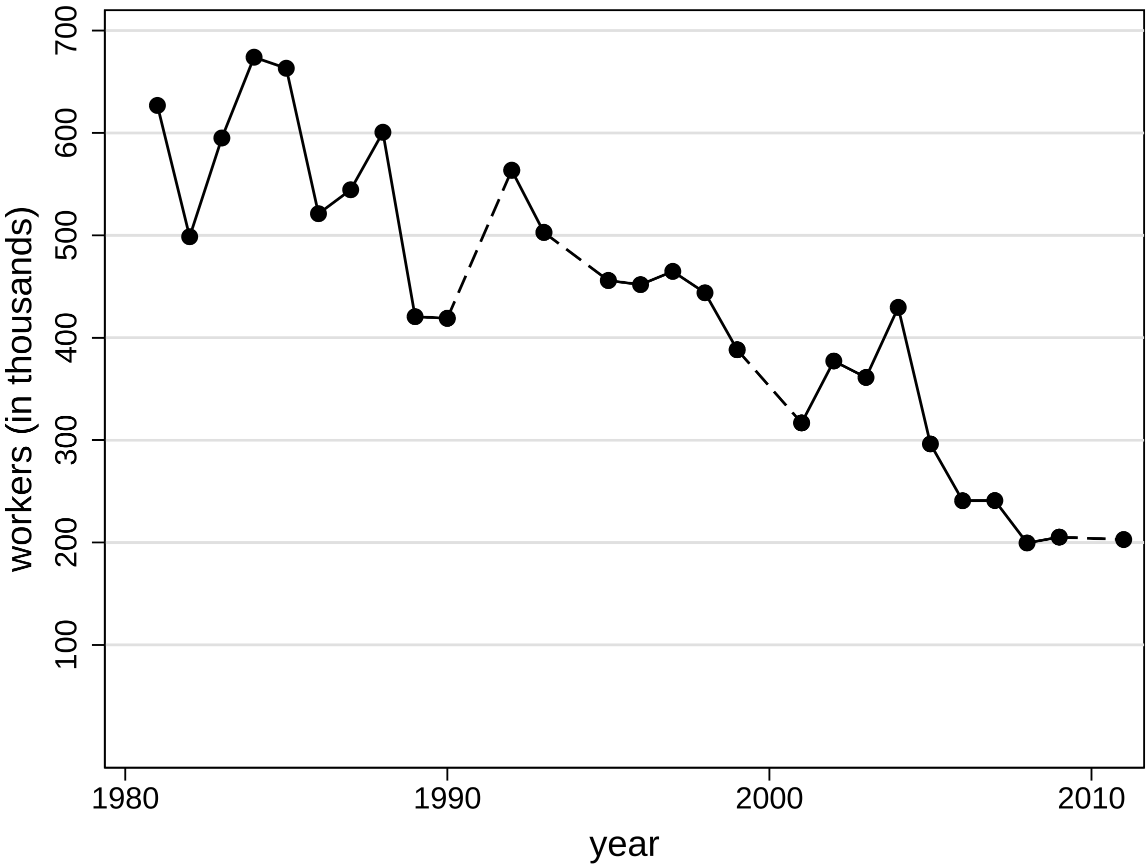

Notes: The Figure depicts the evolution between 1980 and 2011 of the total number of workers employed in soy production (expressed in thousands) in Brazil. Data source is PNAD, a national household survey representative at state level carried out yearly by the IBGE (the survey was not carried out in 1994 and in the census years: 1991, 2000 and 2010). Since the PNAD coverage changed over time, to harmonize the sample across years we exclude: (i) workers located in the states of: Rondonia, Acre, Amazonas, Roraima, Pará and Amapá (North macro-region) because only urban areas (and not rural areas) of these states were covered until 2004; (ii) workers located in the states of: Tocantins, Mato Grosso do Sul, Goias and the Distrito Federal because the sample of households in these states is not complete in the years from 1992 to 1997. The soy sector is identified separately from other agricultural sectors during the whole period: it is identified by code 021 in the sector classification used over the years 1980-2001 and by code 01107 in the sector classification used during the years 2002-2011.

Figure 7
Area planted with maize (1980-2010)

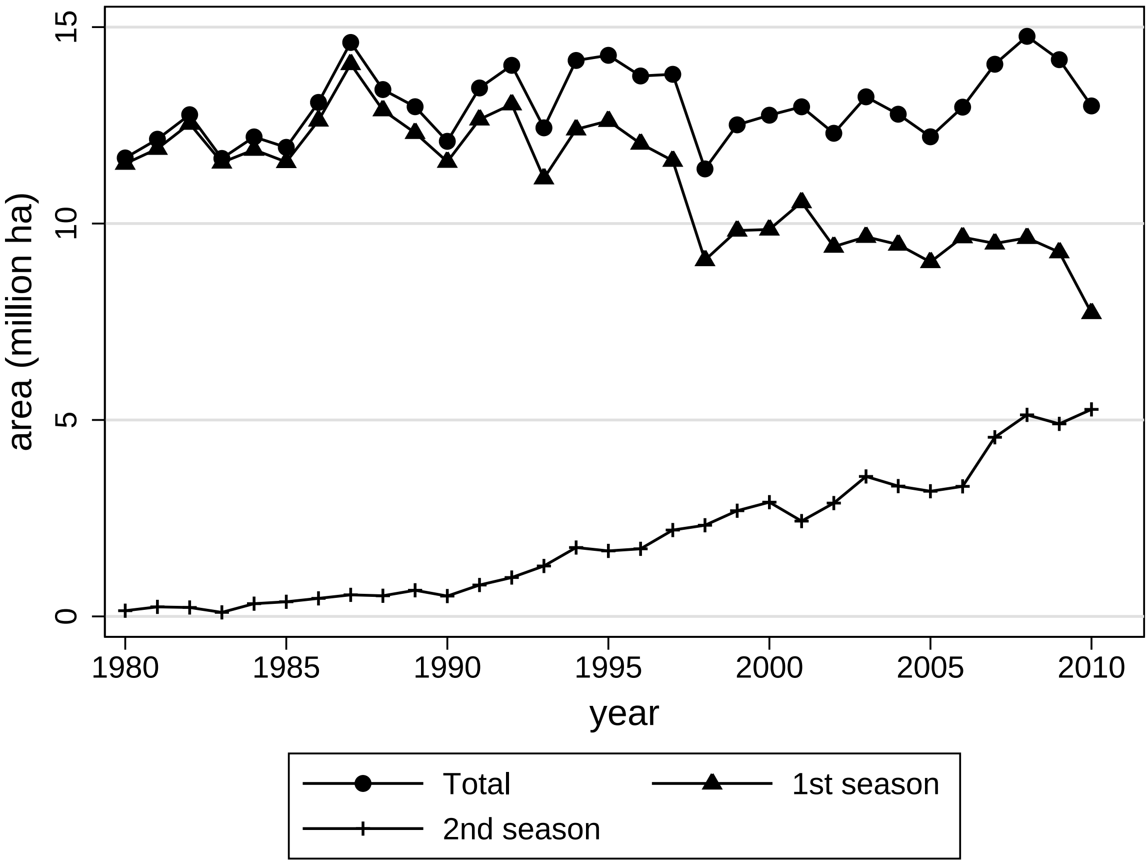

Notes: The Figure depicts the evolution between 1980 and 2010 of the area planted with maize in Brazil (expressed in million hectares). The series show the total area planted with maize as well as the breakdown by the season of harvest (1st or 2nd season). Data source is CONAB, Companhia Nacional de Abastecimento, which is an agency within the Brazilian Ministry of Agriculture. CONAB carries out monthly surveys to monitor the evolution of the harvest of all major crops in Brazil: the surveys are representative at state level and are constructed by interviewing on the ground farmers, agronomists and financial agents in the main cities of the country. All data can be downloaded at: http://www.conab.gov.br/conteudos.php?a=1252&t=..

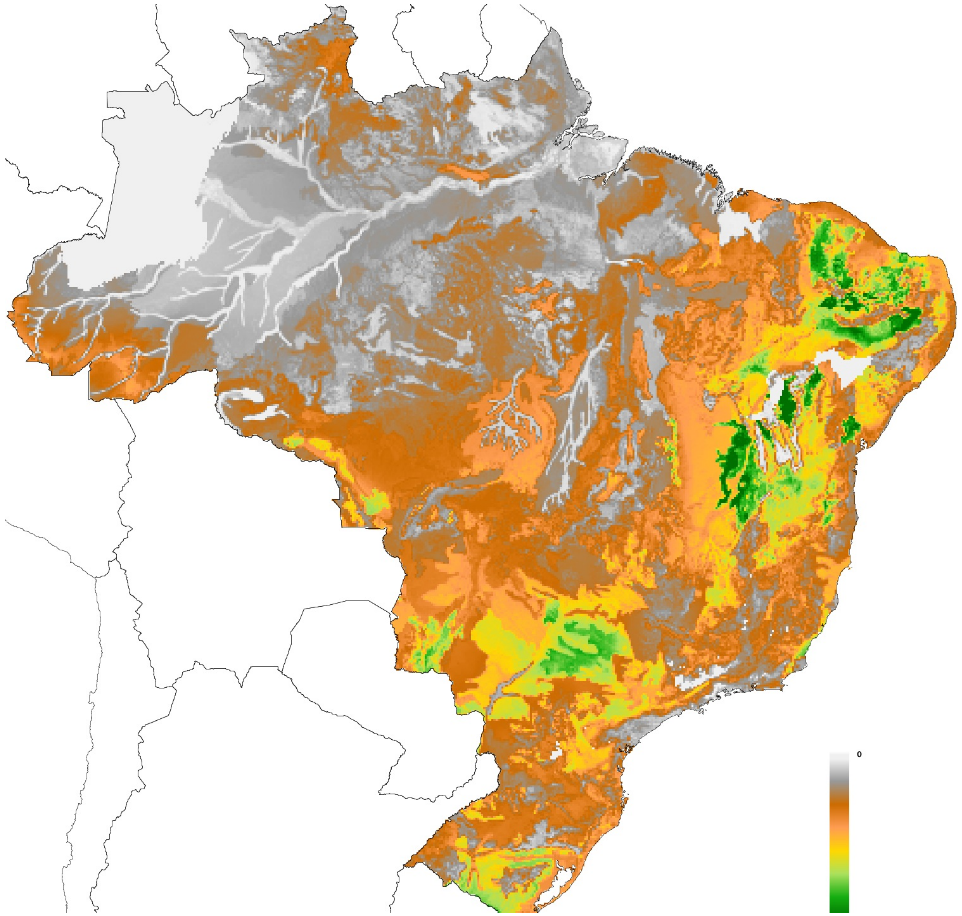

Figure 8

Potential soy yield under low agricultural technology

Notes: Data source is FAO-GAEZ. See section 4 for details.

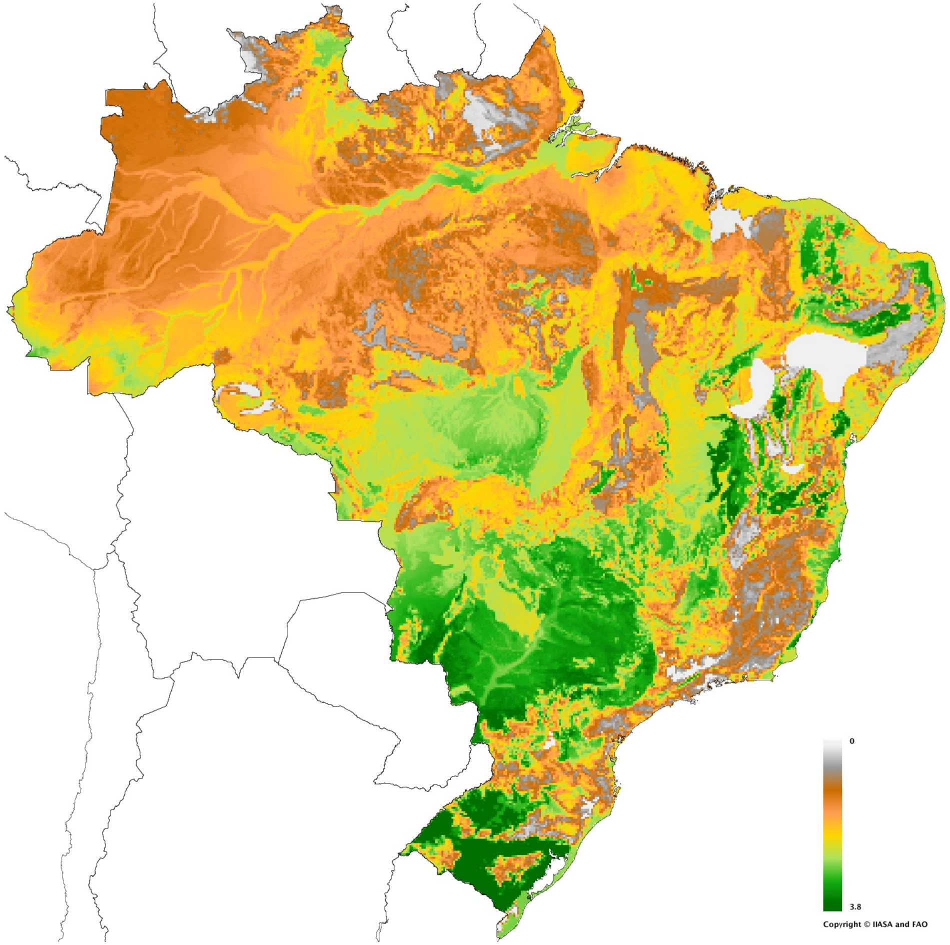

Figure 9

Potential soy yield under high agricultural technology

Notes: Data source is FAO-GAEZ. See section 4 for details.

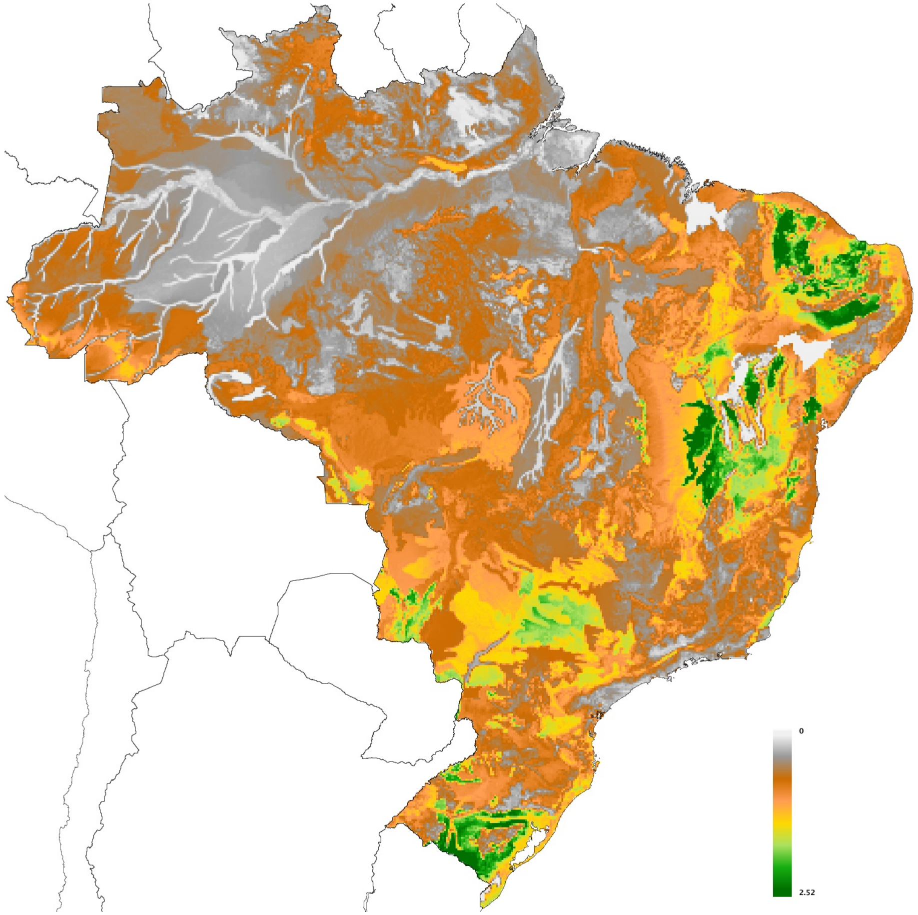

Figure 10

Potential maize yield under low agricultural technology

Notes: Data source is FAO-GAEZ. See section 4 for details.

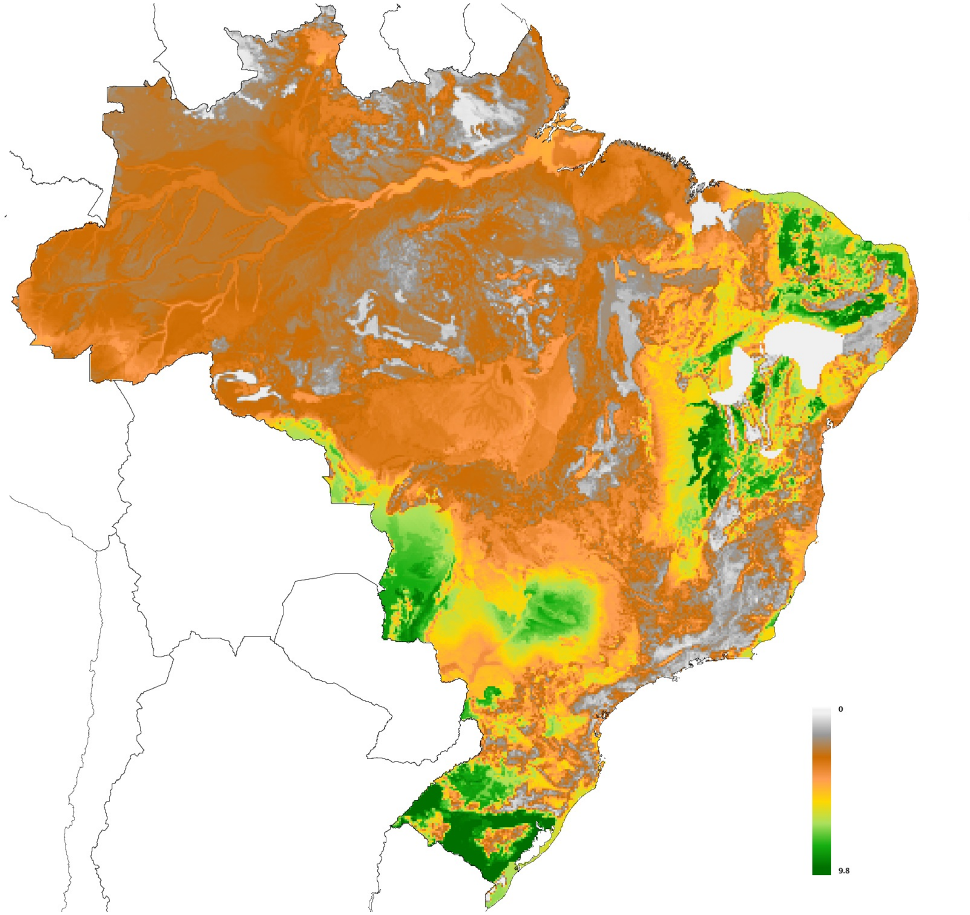

Figure 11

Potential maize yield under high agricultural technology

Notes: Data source is FAO-GAEZ. See section 4 for details.

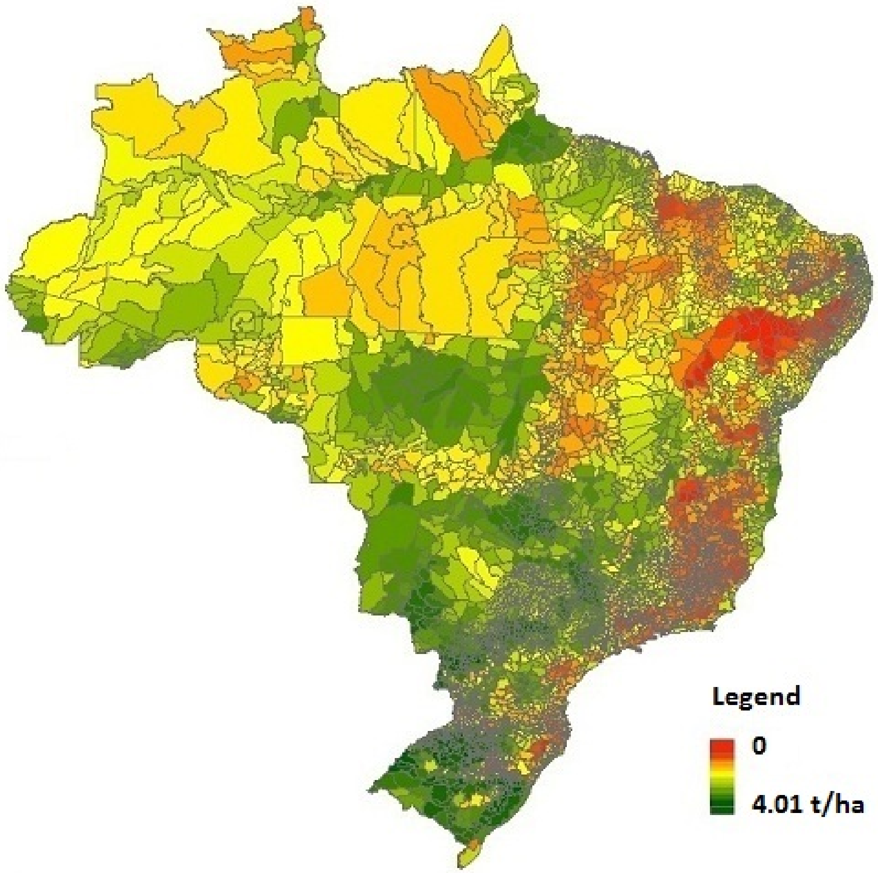

Figure 12
Technological change in soy
Potential yield under high technology minus potential yield under low technology
Municipalities

Notes: Authors' calculations from FAO-GAEZ data. See section 4 for details.

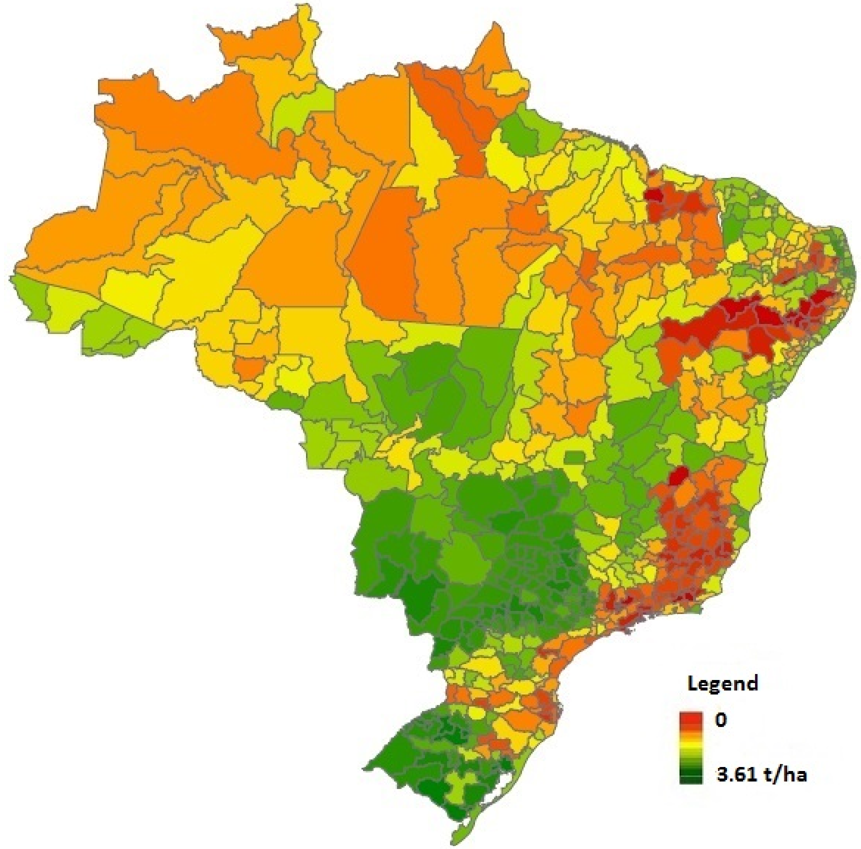

Figure 13
Technological change in soy
Potential yield under high technology minus potential yield under low technology
Micro-regions

Notes: Authors' calculations from FAO-GAEZ data. See section 4 for details.

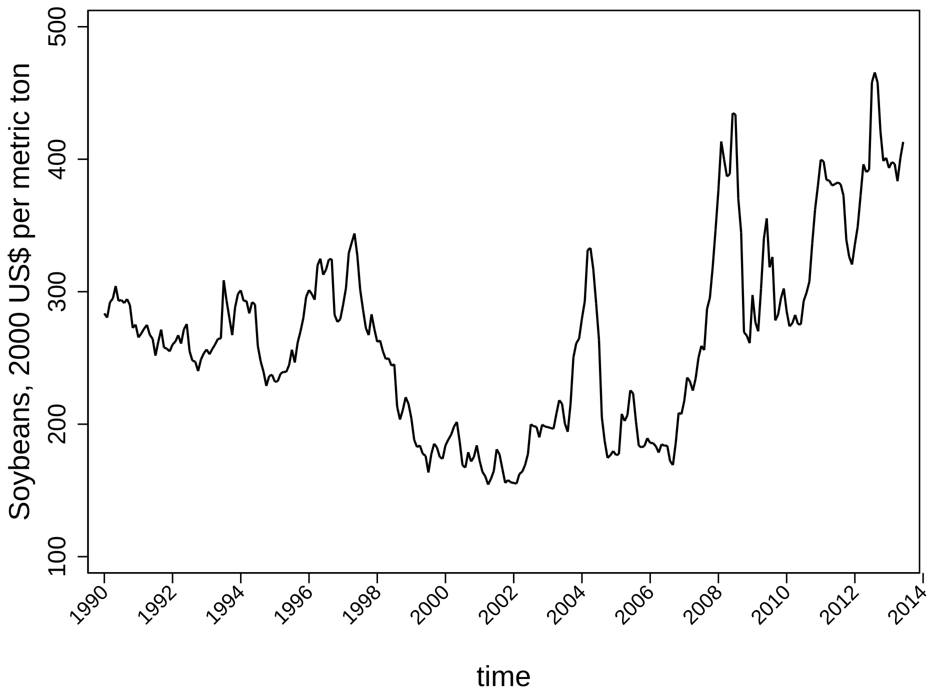

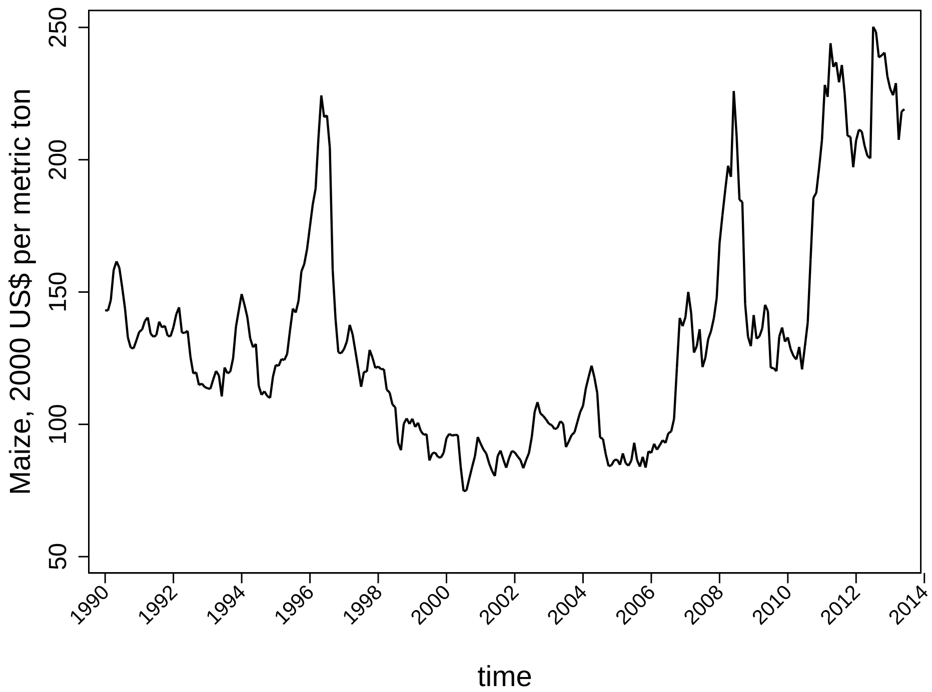

Figure 14
Evolution of soy price (1990-2013)

Notes: The Figure shows the monthly evolution of soy real price between 1990 and 2013. Data are from the IMF Primary Commodity Prices database, series code: *PSOYB_USD*, expressed in nominal US$ per metric ton. We deflate the series using the US *Consumer Price Index for All Urban Consumers: All Items*, source: Federal Reserve St. Louis, series code: *CPIAUCNS*, rescaled so that 2000 is the base year.

Figure 15
Evolution of maize price (1990-2013)

Notes: The Figure shows the monthly evolution of maize real price between 1990 and 2013. Data are from the IMF Primary Commodity Prices database, series code: *PMAIZMT_US*, expressed in nominal US$ per metric ton. We deflate the series using the US *Consumer Price Index for All Urban Consumers: All Items*, source: Federal Reserve St. Louis, series code: *CPIAUCNS*, rescaled so that 2000 is the base year.

Figure 16
Agricultural Frontier

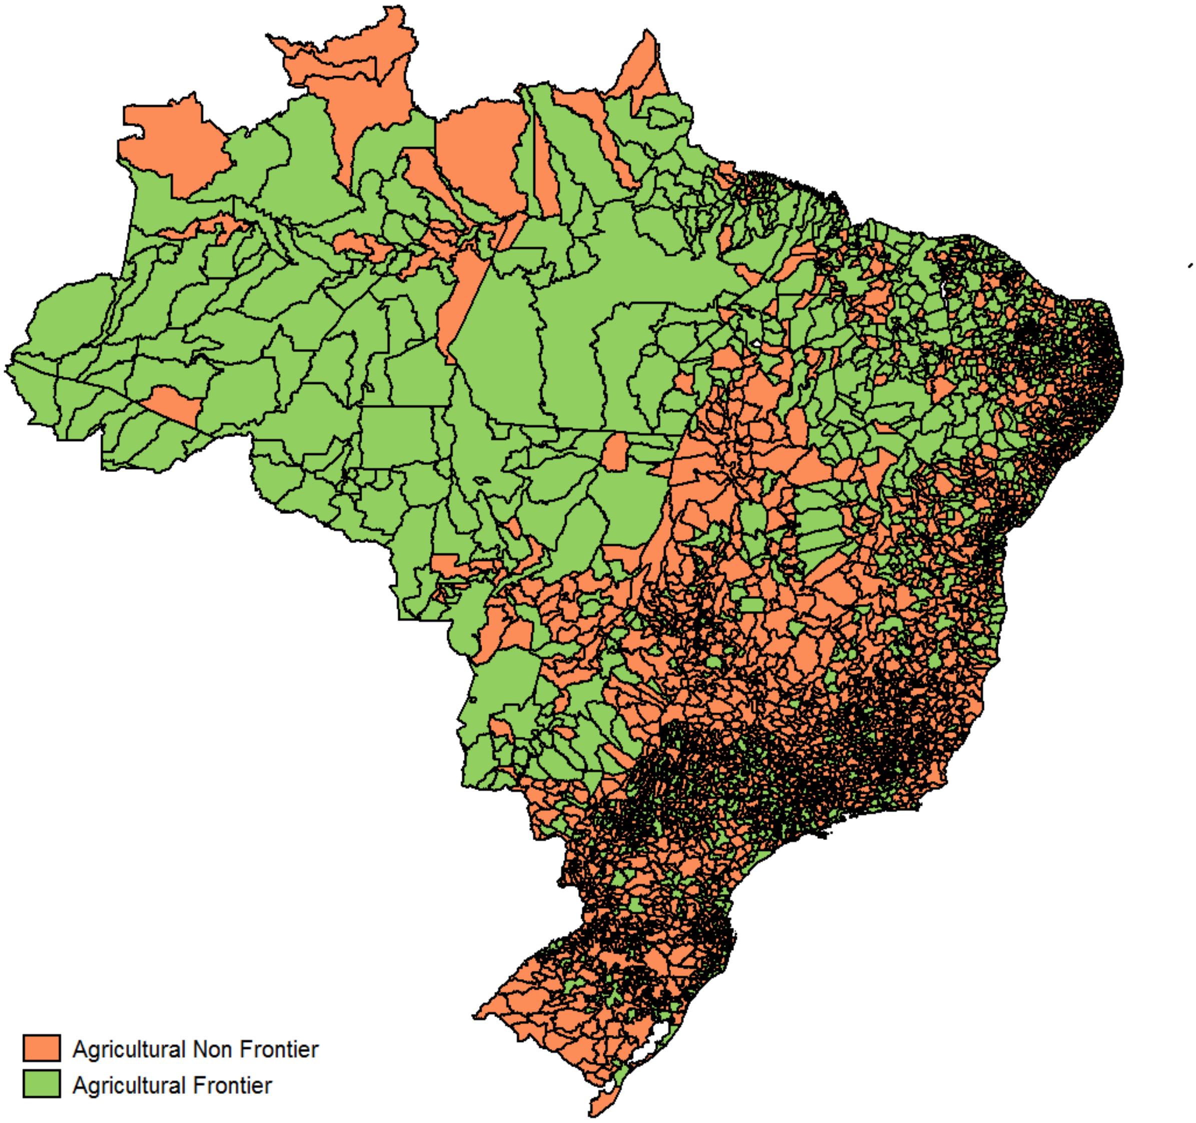

Notes: The Figure shows our definition of agricultural frontier and agricultural non frontier in Brazil. We define municipalities that are part of the agricultural frontier as those that, between 1996 and 2006, experienced an increase in the use of agricultural land for the cultivation of permanent crops, seasonal crops, and cattle ranching. Municipalities that experienced no increase (or a negative change) in used agricultural land between 1996 and 2006 are defined as agricultural non frontier. Data sources are the Brazilian Agricultural Censi of 1996 and 2006, IBGE.

Table 1
Land use (million 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 million 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
Labor intensity in Brazilian agriculture (1996-2006)

	Labor intensity		Change in labor intensity	
Principal activity:	1996	2006	Absolute	Relative
Seasonal crops	107.6	83.7	-23.9	-22%
soy	*28.6*	*17.9*	*-10.7*	*-37%*
all cereals	*92.4*	*76.8*	*-15.6*	*-17%*
other	*159.2*	*145.4*	*-13.8*	*-9%*
Permanent crops	126.8	127.4	0.6	0%
Cattle ranching	22.6	30.6	8.1	36%
Forest	33.9	46.1	12.2	36%

Note: The table reports labor intensity in agriculture by principal activity of the farm. Labor intensity is computed as number of workers per 1000 hectares. 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. 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.

Table 3
Summary statistics of main variables at municipality level

Variable Name			1996		1996-2006 Change		
PANEL A: Agricultural Census			mean	st.dev.	mean	st.dev.	obs.
Log agri labor productivity			7.689	1.191	0.561	0.763	4,159
Log labor intensity in agriculture			-2.586	1.048	-0.026	0.551	4,159
Soy area share			0.027	0.097	0.013	0.062	3,652
Maize area share			0.049	0.068	0.010	0.093	3,652
GE soy area share			0.000	0.000	0.015	0.075	3,652
			2000		2000-2010 Change		
PANEL B: Population Census			mean	st.dev.	mean	st.dev.	obs.
Employment shares:							
Agriculture			0.411	0.196	-0.057	0.073	4,159
Manufacturing			0.099	0.086	0.011	0.054	4,159
Services			0.350	0.139	0.028	0.054	4,159
Public sector			0.140	0.051	0.018	0.036	4,159
Log employment in manufacturing			5.994	1.561	0.216	0.601	4,159
Log wage in manufacturing			5.519	0.509	0.305	0.374	4,159
	Low inputs		High inputs		Difference		
PANEL C: FAO GAEZ	mean	st.dev.	mean	st.dev.	mean	st.dev.	obs.
Potential yield in soy	0.302	0.154	2.111	0.938	1.809	0.851	4,159
Potential yield in maize	0.992	0.494	4.063	2.197	3.071	1.812	4,159

Note: The table reports summary statistics of the main outcomes and independent variables used in the empirical section. In Panel A we report the main variables extracted from the Agricultural Census. In Panel B we report the main variables extracted from the Population Census. For each variable we report mean and standard deviation of its level in 1996 (2000) and of its change between 1996 (2000) and 2006 (2010). In Panel C we report sumary statistics of potential yields of soy and maize under low inputs, under high inputs, and of the difference in yields between high and low inputs. Data is from the FAO GAEZ dataset. For detailed variable definitions see Appendix B.

Table 4
Basic correlations in the data: agriculture
Productivity, labor intensity and employment share

VARIABLES	(1) Δ Log value per worker $_{2006-1996}$	(2) Δ Log labor intensity $_{2006-1996}$	(3) Δ Employment share $_{2010-2000}$
Δ Soy area share $_{2006-1996}$	0.566**	-0.483***	-0.051*
	(0.233)	(0.154)	(0.027)
Δ Maize area share $_{2006-1996}$	1.589***	0.741***	-0.027
	(0.182)	(0.119)	(0.019)
Observations	3,775	3,775	3,775
R-squared	0.023	0.009	0.002

Note: The table reports the OLS estimated coefficients of equation 8 in the text. The independent variables are defined as the share of farm land reaped with soy and maize. The dependent variables are reported on top of the respective columns. Value per worker is defined as total value of output of all agricultural activities divided by the total number of workers employed in agriculture. Labor intensity is the total number of workers employed in agriculture divided by total area in farms. Share of workers employed in agriculture is defined as total number of workers in agriculture divided by total number of workers in all sectors. The source of data for the independent variables and the dependent variables reported in columns 1 and 2 are the agricultural censi of 1996 and 2006. The source for the employment share reported in column 3 are the population censi of 2000 and 2010. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 5
Basic correlations in the data: manufacturing
Employment share, employment and wages

VARIABLES	(1) Δ Employment share $_{2010-2000}$	(2) Δ Log employment $_{2010-2000}$	(3) Δ Log wage $_{2010-2000}$
Δ Soy area share $_{2006-1996}$	0.085***	0.986***	0.046
	(0.021)	(0.224)	(0.121)
Δ Maize area share $_{2006-1996}$	0.003	-0.018	-0.041
	(0.012)	(0.142)	(0.085)
Observations	3,775	3,775	3,775
R-squared	0.005	0.005	0.000

Note: The table reports the OLS estimates of the coefficients in equation 8 in the text. The independent variables are defined as the share of farm land reaped with soy and maize. The dependent variables are reported on top of the respective columns. Employment share in manufacturing is defined as number of people employed in the manufacturing sector (CNAE codes between 15 and 37) divided by total number of people employed in all sectors. Employment in manufacturing is the natural logarithm of people employed in the manufacturing sector. Wage is calculated as the logarithm of the average wage of manufacturing workers in 2000 Reais. The source of data for the independent variables are the agricultural censi of 1996 and 2006. The source for the dependent variables are the population censi of 2000 and 2010. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 6
Comparing Municipalities below/above median increase in potential soy yield

	(1) Below ΔA^{soy} median	(2) Above ΔA^{soy} median	(3) Difference
Agricultural Employment Share $_{1991}$	0.526	0.465	-0.061***
			[0.007]
Manufacturing Employment Share $_{1991}$	0.077	0.094	0.016***
			[0.003]
Share Rural Pop $_{1991}$	0.516	0.404	-0.112***
			[0.007]
Log Income per Capita $_{1991}$	4.388	4.656	0.267***
			[0.018]
Log Pop Density $_{1991}$	3.153	3.216	0.062
			[0.041]
Literacy rate $_{1991}$	0.688	0.745	0.057***
			[0.005]
Observations	2,080	2,079	

Note: The table reports average values of observable characteristics of municipalities that rank below and above the median of ΔA^{soy}, the exogenous measure of technological change in soy. Column (3) reports the difference between columns (2) and (1), along with the standard error and significance level of the difference. Significance levels: *** p<0.01, ** p<0.05, * p<0.1. The observable characteristics are: agricultural employment share, manfuacturing employment share, share of rural adult population, income per capita (in logs), population density (in logs), and literacy rate. All observable characteristics are from the Population Census 1991.

Table 7
The effect of technological change on agriculture
GE soy adoption

	(1)	(2)	(3)	(4)
VARIABLES	Δ GE Soy area share $_{2006-1996}$		Δ Non-GE Soy area share $_{2006-1996}$	
ΔA^{soy}	0.021***	0.019***	-0.009***	-0.009***
	(0.002)	(0.002)	(0.002)	(0.002)
Share Rural Pop $_{1991}$	0.039***	0.085***	-0.017***	-0.044***
	(0.005)	(0.008)	(0.004)	(0.007)
Log Income per Capita $_{1991}$		-0.000		0.001
		(0.003)		(0.003)
Log Pop Density $_{1991}$		0.003***		-0.005***
		(0.001)		(0.001)
Literacy Rate $_{1991}$		0.114***		-0.048***
		(0.011)		(0.010)
Observations	3,652	3,652	3,652	3,652
R-squared	0.083	0.162	0.019	0.044

Note: The table reports the OLS estimates of the coefficients in equation 11 in the text where the dependent variable is defined as the share of farm land reaped with GE soy (columns 1 and 2) and non-GE soy (columns 3 and 4). ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. The source of data for the independent variables are the agricultural censi of 1996 and 2006. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 8
The effect of technological change on agriculture
Soy and maize expansion

	(1)	(2)	(3)	(4)
VARIABLES	Δ Soy area share $_{2006-1996}$		Δ Maize area share $_{2006-1996}$	
ΔA^{soy}	0.013***	0.013***		0.001
	(0.001)	(0.002)		(0.003)
ΔA^{maize}		-0.001	0.003***	0.003***
		(0.001)	(0.001)	(0.001)
Share Rural Pop $_{1991}$	0.020***	0.039***	0.011**	0.010
	(0.003)	(0.005)	(0.004)	(0.007)
Log Income per Capita $_{1991}$		0.001		-0.005
		(0.002)		(0.004)
Log Pop Density $_{1991}$		-0.002***		0.004***
		(0.000)		(0.001)
Literacy Rate $_{1991}$		0.064***		-0.006
		(0.007)		(0.012)
Observations	3,652	3,652	3,652	3,652
R-squared	0.067	0.124	0.009	0.015

Note: The table reports the OLS estimates of the coefficients in equation 11 in the text. Dependent variables – reported on top of the respective columns – are defined as the share of farm land reaped with soy and maize. ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source of data for the dependent variables are the agricultural censi of 1996 and 2006. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 9
The effect of technological change on agriculture
Productivity, labor intensity and employment share

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Δ Log value per worker $_{2006-1996}$		Δ Log labor intensity $_{2006-1996}$		Δ Employment share $_{2010-2000}$	
ΔA^{soy}	0.116***	0.132***	-0.058***	-0.064***	-0.013***	-0.018***
	(0.024)	(0.026)	(0.018)	(0.021)	(0.002)	(0.002)
ΔA^{maize}	-0.025**	-0.033***	0.031***	0.033***	0.003***	0.005***
	(0.011)	(0.011)	(0.008)	(0.009)	(0.001)	(0.001)
Share Rural Pop $_{1991}$	0.257***	0.125*	-0.134***	-0.175***	-0.075***	-0.060***
	(0.057)	(0.070)	(0.048)	(0.051)	(0.005)	(0.007)
Log Income per Capita $_{1991}$		-0.007		0.028		0.015***
		(0.045)		(0.039)		(0.004)
Log Pop Density $_{1991}$		-0.015		-0.017		-0.002*
		(0.011)		(0.010)		(0.001)
Literacy Rate $_{1991}$		-0.278**		-0.119		-0.004
		(0.139)		(0.116)		(0.014)
Observations	4,159	4,159	4,159	4,159	4,159	4,159
R-squared	0.009	0.012	0.005	0.007	0.047	0.055

Note: The table reports the OLS estimates of the coefficients in equation 11 in the text. The dependent variables are reported on top of the respective columns. Value per worker is defined as total value of output of all agricultural activities divided by the total number of workers employed in agriculture. Labor intensity is the total number of workers employed in agriculture divided by total area in farms. Share of workers employed in agriculture is defined as total number of workers in agriculture divided by total number of workers in all sectors. ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source of data for the dependent variables reported in columns 1 and 2 are the agricultural censi of 1996 and 2006. Thus, changes are calculated over the years 1996 and 2006. The source for the employment share reported in column 3 are the population censi of 2000 and 2010. In this case, changes in the dependent variable are calculated over the years 2000 and 2010. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 10

The effect of agricultural technological change on manufacturing

Employment share, employment and wages

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Δ Employment share $_{2010-2000}$		Δ Log employment $_{2010-2000}$		Δ Log wage $_{2010-2000}$	
ΔA^{soy}	0.021***	0.019***	0.210***	0.174***	-0.035***	-0.028**
	(0.002)	(0.002)	(0.018)	(0.019)	(0.012)	(0.013)
ΔA^{maize}	-0.005***	-0.004***	-0.054***	-0.038***	0.018***	0.014**
	(0.001)	(0.001)	(0.008)	(0.009)	(0.005)	(0.006)
Share Rural Pop $_{1991}$	-0.008**	0.005	-0.231***	0.024	0.201***	-0.006
	(0.004)	(0.005)	(0.043)	(0.056)	(0.026)	(0.035)
Log Income per Capita $_{1991}$		0.001		0.088**		-0.107***
		(0.003)		(0.037)		(0.025)
Log Pop Density $_{1991}$		0.002**		0.021***		-0.035***
		(0.001)		(0.008)		(0.005)
Literacy Rate $_{1991}$		0.028***		0.270**		0.101
		(0.010)		(0.116)		(0.074)
Observations	4,159	4,159	4,159	4,159	4,159	4,159
R-squared	0.059	0.066	0.059	0.075	0.022	0.043

Note: The table reports the OLS estimates of the coefficients in equation 11 in the text. The dependent variables are reported on top of the respective columns. Employment share in manufacturing is defined as number of people employed in the manufacturing sector (CNAE codes between 15 and 37) divided by total number of people employed in all sectors. Employment in manufacturing is the natural logarithm of people employed in the manufacturing sector. Wage is calculated as the logarithm of the average wage of manufacturing workers in 2000 Reais. ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source for the dependent variables are the population censi of 2000 and 2010. In this case, changes in the dependent variable are calculated over the years 2000 and 2010. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 11
The effect of agricultural technological change on sectoral employment shares
Employment shares

VARIABLES	(1) Δ Agriculture Employment share $_{2010-2000}$	(2) Δ Manufacturing Employment share $_{2010-2000}$	(3) Δ Services Employment share $_{2010-2000}$	(4) Δ Public Sector Employment share $_{2010-2000}$
ΔA^{soy}	-0.018***	0.019***	-0.002	0.001
	(0.002)	(0.002)	(0.002)	(0.001)
ΔA^{maize}	0.005***	-0.004***	0.000	-0.001*
	(0.001)	(0.001)	(0.001)	(0.001)
Share Rural Pop $_{1991}$	-0.060***	0.005	0.039***	0.016***
	(0.007)	(0.005)	(0.005)	(0.003)
Log Income per Capita $_{1991}$	0.015***	0.001	-0.013***	-0.003
	(0.004)	(0.003)	(0.003)	(0.002)
Log Pop Density $_{1991}$	-0.002*	0.002**	0.001	-0.001*
	(0.001)	(0.001)	(0.001)	(0.001)
Literacy Rate $_{1991}$	-0.004	0.028***	-0.010	-0.014**
	(0.014)	(0.010)	(0.010)	(0.007)
Observations	4,159	4,159	4,159	4,159
R-squared	0.055	0.066	0.089	0.039

Note: The table reports the OLS estimates of the coefficients in equation 11 in the text. The dependent variables are reported on top of the respective columns. Employment shares are defined as number of people employed in each sector divided by total number of people employed in all sectors. ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source for the dependent variables are the Population Censi of 2000 and 2010. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 12
Variable Factor Endowment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Migration rate $_{2010-2000}$		Δ Agri Empl. Share $_{2010-2000}$		Δ Manuf Empl. Share $_{2010-2000}$	
VARIABLES	All	Agri Non-Frontier	Agri Frontier	Agri Non-Frontier	Agri Frontier	Agri Non-Frontier	Agri Frontier
ΔA^{soy}	-0.012***	-0.013**	-0.013**	-0.021***	-0.015***	0.021***	0.018***
	(0.004)	(0.005)	(0.006)	(0.003)	(0.004)	(0.002)	(0.004)
ΔA^{maize}	0.006***	0.007***	0.003	0.007***	0.001	-0.005***	-0.003*
	(0.002)	(0.002)	(0.003)	(0.001)	(0.002)	(0.001)	(0.001)
Share Rural Pop $_{1991}$	-0.104***	-0.123***	-0.056***	-0.065***	-0.046***	0.012**	-0.008
	(0.012)	(0.014)	(0.021)	(0.008)	(0.012)	(0.006)	(0.009)
Log Income per Capita $_{1991}$	0.071***	0.071***	0.065***	0.018***	0.009	0.005	-0.004
	(0.008)	(0.009)	(0.013)	(0.005)	(0.007)	(0.004)	(0.005)
Log Pop Density $_{1991}$	-0.002	0.004	-0.005*	-0.003**	-0.001	0.001	0.001
	(0.002)	(0.003)	(0.003)	(0.001)	(0.002)	(0.001)	(0.001)
Literacy Rate $_{1991}$	-0.013	-0.004	0.065*	-0.018	0.041*	0.011	0.032**
	(0.023)	(0.028)	(0.038)	(0.017)	(0.024)	(0.012)	(0.016)
Observations	4,159	2,625	1,534	2,625	1,534	2,625	1,534
R-squared	0.172	0.193	0.174	0.060	0.062	0.069	0.061

Note: The table reports the OLS estimates of the coefficients in equation 11 in the text. The dependent variables are reported on top of the respective columns. Migration rate is defined as number of net migrants in a municipality divided by the total number of people living in that municipality in the initial year. To construct net migration in a municipality we use the standard cohort average method (see Appendix B for details). Coefficients reported in columns 3, 5 and 7 are estimated using only municipalities that are defined as agricultural frontier. We define municipalities that are part of the agricultural frontier as those that, between 1996 and 2006, experienced an increase in agricultural land used for the cultivation of permanent crops, seasonal crops, and cattle ranching. Coefficients reported in columns 2, 4 and 6 are estimated using only municipalities that are defined as Agricultural Non-Frontier. These are municipalities that experienced no increase, or a negative change, in used agricultural land between 1996 and 2006. ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source for the dependent variables are the Population Censi of 2000 and 2010. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 13
The effect of agricultural technological change on manufacturing and migration

Manufacturing employment, manufacturing wages and net migration

Robustness check of results reported in Tables 10 and 12: controlling for pre-existing trends

VARIABLES	(1) Δ Log manufacturing employment $_t$	(2) Δ Log manufacturing wage $_t$	(3) Migration rate $_t$
$\Delta A^{soy} \times After_t$	0.243***	-0.130***	-0.014**
	(0.025)	(0.018)	(0.006)
$\Delta A^{maize} \times After_t$	-0.062***	0.052***	0.011***
	(0.012)	(0.008)	(0.003)
ΔA^{soy}	-0.009	0.068***	-0.004
	(0.020)	(0.014)	(0.005)
ΔA^{maize}	0.001	-0.025***	-0.004
	(0.009)	(0.007)	(0.003)
Share Rural Pop $_{t-1}$	0.223***	0.010	-0.153***
	(0.042)	(0.028)	(0.014)
Log Income per Capita $_{t-1}$	0.022	-0.055***	0.064***
	(0.027)	(0.020)	(0.007)
Log Pop Density $_{t-1}$	-0.018***	-0.003	-0.002
	(0.006)	(0.004)	(0.002)
Literacy Rate $_{t-1}$	0.242***	0.190***	-0.002
	(0.082)	(0.054)	(0.019)
$After_t$	-0.226***	0.169***	-0.033***
	(0.032)	(0.023)	(0.007)
Observations	8,040	8,040	8,040
R-squared	0.032	0.019	0.140

Note: The dependent variables are reported on top of the respective columns. Employment in manufacturing is the natural logarithm of people employed in the manufacturing sector. Wage is calculated as the logarithm of the average wage of manufacturing workers in 2000 Reais. Migration rate is defined as number of net migrants in a municipality divided by the total number of people living in that municipality in the initial year. To construct net migration in a municipality we use the standard cohort average method (see Appendix B for details). ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source for the dependent variables are the Population Censi of 1991, 2000 and 2010. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 14
The effect of agricultural technological change on manufacturing
Employment share, employment and wages
Robustness of results reported in Table 10 to a larger unit of observation: micro-regions

VARIABLES	(1) Δ Employment share $_{2010-2000}$	(2) Δ Log employment $_{2010-2000}$	(3) Δ Log wage $_{2010-2000}$
ΔA^{soy}	0.015***	0.123***	-0.027
	(0.004)	(0.028)	(0.016)
ΔA^{maize}	-0.002	-0.031**	0.018***
	(0.001)	(0.013)	(0.007)
Share Rural Pop $_{1991}$	0.010	0.002	-0.108
	(0.012)	(0.120)	(0.087)
Log Income per Capita $_{1991}$	-0.002	0.078	-0.163**
	(0.007)	(0.087)	(0.069)
Log Pop Density $_{1991}$	0.004***	0.033***	-0.032***
	(0.001)	(0.011)	(0.007)
Literacy Rate $_{1991}$	0.014	0.023	0.085
	(0.021)	(0.255)	(0.169)
Observations	557	557	557
R-squared	0.093	0.128	0.249

Note: The table reports the OLS estimates of the coefficients in equation 11 in the text. The dependent variables are reported on top of the respective columns. Employment share in manufacturing is defined as number of people employed in the manufacturing sector (CNAE codes between 15 and 37) divided by total number of people employed in all sectors. Employment in manufacturing is the natural logarithm of people employed in the manufacturing sector. Wage is calculated as the logarithm of the average wage of manufacturing workers in 2000 Reais. ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source for the dependent variables are the population censi of 2000 and 2010. In this case, changes in the dependent variable are calculated over the years 2000 and 2010. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the micro-region. Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 15
The effect of agricultural technological change on manufacturing

Employment share, employment and wages

Robustness of results reported in Table 10 to excluding sectors directly linked to soy and maize

VARIABLES	(1) Δ Employment share $_{2010-2000}$	(2) Δ Log employment $_{2010-2000}$	(3) Δ Log wage $_{2010-2000}$
ΔA^{soy}	0.012***	0.154***	-0.004
	(0.001)	(0.020)	(0.015)
ΔA^{maize}	-0.004***	-0.052***	0.005
	(0.001)	(0.009)	(0.007)
Share Rural Pop $_{1991}$	0.008**	0.047	-0.003
	(0.004)	(0.057)	(0.042)
Log Income per Capita $_{1991}$	-0.002	0.082**	-0.132***
	(0.002)	(0.037)	(0.028)
Log Pop Density $_{1991}$	0.003***	0.036***	-0.033***
	(0.000)	(0.008)	(0.006)
Literacy Rate $_{1991}$	0.020***	0.130	0.212**
	(0.007)	(0.121)	(0.085)
Observations	4,159	4,151	4,080
R-squared	0.034	0.046	0.027

Note: Employment share in manufacturing is defined as number of people employed in the manufacturing sector (CNAE codes between 15 and 37) divided by total number of people employed in all sectors. Employment in manufacturing is the natural logarithm of people employed in the manufacturing sector. Wage is calculated as the logarithm of the average wage of manufacturing workers in 2000 Reais. ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source for the dependent variables are the population censi of 2000 and 2010. In this case, changes in the dependent variable are calculated over the years 2000 and 2010. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1. The number of observations in columns 2 and 3 is smaller because in some municipalities sectors directly linked to soy and maize account for the whole manufacturing sector.

Table 16

The effect of agricultural technological change on manufacturing

Data on employment and wages from the Annual Manufacturing Survey (PIA)

Robustness of results reported in Table 10 to controlling for commodity prices

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Log Total Employment $_t$			Log Wage $_t$		
A^{soy}	0.122***	0.122***	0.096***	-0.026*	-0.026*	-0.018
	(0.030)	(0.030)	(0.029)	(0.013)	(0.013)	(0.013)
A^{maize}	-0.031**	-0.031**	-0.027*	0.013**	0.013**	0.011*
	(0.015)	(0.015)	(0.015)	(0.006)	(0.006)	(0.006)
Share Rural Pop $_{1991} \times t$	0.052***	0.052***	0.019	0.030***	0.030***	0.025***
	(0.012)	(0.012)	(0.016)	(0.005)	(0.005)	(0.007)
$P^{soy}A^{soy}$		-0.001	-0.000		0.000	0.000
		(0.001)	(0.001)		(0.001)	(0.001)
$P^{maize}A^{maize}$		-0.001	-0.001		0.000	0.000
		(0.001)	(0.001)		(0.000)	(0.000)
Literacy Rate $_{1991} \times t$			-0.098***			0.014
			(0.037)			(0.017)
Log Pop Density $_{1991} \times t$			-0.010***			0.001
			(0.002)			(0.001)
Log Income per Capita $_{1991} \times t$			0.019*			-0.009**
			(0.010)			(0.004)
Observations	25,262	25,262	25,262	25,239	25,239	25,239
R-squared	0.923	0.923	0.923	0.778	0.778	0.778

Note: The table reports the OLS estimates of the coefficients in equation 12 in the text. The dependent variables are reported on top of the respective columns. Total employment is the natural logarithm of the total number of workers employed in manufacturing plants (CNAE 1.0 codes 15 to 37) owned by firms that employ at least 30 employees within an AMC. The average wage is computed from manufacturing plants (CNAE 1.0 codes 15 to 37) owned by firms that employ at least 30 employees. Wage is defined as the aggregate wage bill (in real terms) across firm within an AMC divided by total number of workers across the same firms within the same AMC. A^{soy} is defined as potential soy yield under high inputs for the years between 2003 and 2007, and the potential soy yield under low inputs for the years between 1996 and 2002. A^{maize} is defined as potential maize yield under high inputs for the years between 2003 and 2007, and potential maize yield under low inputs for the years between 1996 and 2002. P^zA^z *controls* stand for the interaction of the potential yield of soy and maize under low inputs interacted with price levels of these crops between 1996 and 2007. The source for the dependent variables is the plant-level supplement of the yearly industrial survey (PIA) for the years 1996 to 2007. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The unit of observation is the AMC (Área Mínima Comparável). Standard errors clustered at AMC level are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 17
The effect of technological change on agriculture

Soy and maize expansion

Robustness of results reported in Table 8 to correcting standard errors for spatial correlation

	(1)	(2)	(3)	(4)
VARIABLES	Δ Soy area share $_{2006-1996}$		Δ Maize area share $_{2006-1996}$	
ΔA^{soy}	0.013	0.013		0.001
Robust standard errors	(0.001)***	(0.002)***		(0.003)
Microregion-clustered standard errors	(0.002)***	(0.003)***		(0.004)
Mesoregion-clustered standard errors	(0.004)***	(0.005)**		(0.005)
ΔA^{maize}		-0.001	0.003	0.003
Robust standard errors		(0.001)	(0.001)***	(0.001)***
Microregion-clustered standard errors		(0.001)	(0.001)***	(0.002)*
Mesoregion-clustered standard errors		(0.003)	(0.002)*	(0.002)
Rural Pop Share	Y	Y	Y	Y
Controls	N	Y	N	Y
Observations	3,652	3,652	3,652	3,652
R-squared	0.067	0.124	0.009	0.015

Note: The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 18

The effect of technological change on agriculture

Productivity, labor intensity and employment share

Robustness of results reported in Table 9 to correcting standard errors for spatial correlation

VARIABLES	(1) Δ Log value per worker $_{2006}$	(2) Δ Log labor intensity $_{2006}$	(3) Δ Employment share $_{2010}$
ΔA^{soy}	0.132	-0.064	-0.018
Robust standard errors	(0.026)***	(0.021)***	(0.002)***
Microregion-clustered standard errors	(0.032)***	(0.026)**	(0.004)***
Mesoregion-clustered standard errors	(0.033)***	(0.030)**	(0.006)***
ΔA^{maize}	-0.033	0.033	0.005
Robust standard errors	(0.011)***	(0.009)***	(0.001)***
Microregion-clustered standard errors	(0.015)**	(0.012)***	(0.002)**
Mesoregion-clustered standard errors	(0.016)**	(0.015)**	(0.003)
Controls	Y	Y	Y
Observations	4,159	4,159	4,159
R-squared	0.012	0.007	0.055

Note: The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 19
The effect of technological change on manufacturing
Employment share, employment and wages
Robustness of results reported in Table 10 to correcting standard errors for spatial correlation

VARIABLES	(1) Δ Employment share $_{2010}$	(2) Δ Log employment $_{2010}$	(3) Δ Log wage $_{2010}$
ΔA^{soy}	0.019	0.174	-0.028
Robust standard errors	(0.002)***	(0.019)***	(0.013)**
Microregion-clustered standard errors	(0.004)***	(0.029)***	(0.014)**
Mesoregion-clustered standard errors	(0.006)***	(0.043)***	(0.019)
ΔA^{maize}	-0.004	-0.038	0.014
Robust standard errors	(0.001)***	(0.009)***	(0.006)**
Microregion-clustered standard errors	(0.002)**	(0.014)***	(0.006)**
Mesoregion-clustered standard errors	(0.003)	(0.026)	(0.008)*
Controls	Y	Y	Y
Observations	4,159	4,159	4,159
R-squared	0.059	0.075	0.043

Note: The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.