

Agricultural Productivity and Structural Transformation: Evidence from Brazil[†]

By PAULA BUSTOS, BRUNO CAPRETTINI, AND JACOPO PONTICELLI*

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 led to industrial growth, as predicted by the model. (JEL J43, O13, O14, O33, Q15, Q16)

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; Lewis 1954; 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

*Bustos: CEMFI, Casado del Alisal 5, 28014 Madrid (e-mail: paula.bustos@cemfi.es); Caprettini: University of Zurich, Schönberggasse 1, CH-8001 Zurich (e-mail: bruno.caprettini@gmail.com); Ponticelli: University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637 (e-mail: jacopo.ponticelli@chicagobooth.edu). We received valuable comments from David Atkin, Francisco Buera, Vasco Carvalho, Gino Gancia, Gene Grossman, Juan Carlos Hallak, Chang-Tai Hsieh, Joseph Kaboski, Nina Pavcnik, Joao Pinho de Melo, Andrés Rodríguez-Clare, Silvana Tenreyro, Jaume Ventura, Kei-Mu Yi, and participants at presentations held at Pontificia Universidade Católica de Rio de Janeiro, Fundação Getulio Vargas, CREI, UPF, LSE, Universidad Torcuato di Tella, Harvard University, Brown University, MIT, Chicago Booth, Federal Reserve Bank of New York, CEMFI, Toulouse School of Economics, Yale University, Columbia University, Stanford University, Dartmouth College, Northwestern University, Copenhagen Business School, University of Southern California, SED annual meeting, CEPR ESSIM, CEPR ERWIT, Barcelona GSE Summer Forum, Princeton IES Summer Workshop, AEA Annual Meeting, NBER SI International Trade and Investment, NBER SI Development and Entrepreneurship, NBER SI Development Economics, and LACEA. Part of this research was undertaken while Paula Bustos and Bruno Caprettini were visiting the International Economics Section (IES) at Princeton University. We thank the IES for its hospitality and support. We acknowledge financial support from the Private Enterprise Development in Low-Income Countries Project by the CEPR and UK Department for International Development and the Spanish Ministry of Science and Innovation (ECO-2011-25624); Paula Bustos acknowledges financial support from the Spanish Ministry of Economy and Competitiveness (RYC-2012-11979). The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <http://dx.doi.org/10.1257/aer.20131061> to visit the article page for additional materials and author disclosure statement(s).

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

structural transformation formalize their ideas by showing how productivity growth in agriculture can generate demand for manufacturing goods.² However, several scholars noted 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.³ Despite the richness of the theoretical literature, there is scarce empirical evidence testing the mechanisms proposed by these models.⁴

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, it can be characterized as labor-augmenting technical change. In addition, we study the effects of the introduction of a second harvesting season for maize (*milho safrinha*). This technique permits the growth of two crops per year, effectively increasing the land endowment. Thus, it can be characterized as land-augmenting technical change. The simultaneous expansion of these two crops allows us 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 as labor reallocates toward agriculture, as in Matsuyama (1992). Similar results are obtained when technical change is land-augmenting. However, if land and labor are strong complements in agricultural production, labor-augmenting technical change reduces labor demand in agriculture and causes workers to reallocate toward manufacturing. In sum, the model predicts that the effect of agricultural productivity on structural transformation in open economies depends 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 labor-augmenting agricultural technologies reduces labor demand in the agricultural sector and induces the reallocation of workers toward the industrial sector. However, causality could run in the opposite direction. For example: an increase in productivity in the industrial sector could raise labor demand and 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 commercially released in the United States in 1996, and

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

³See Mokyr (1976); Field (1978); Wright (1979); Corden and Neary (1982); Krugman (1987); and Matsuyama (1992).

⁴Empirical studies of structural transformation include Foster and Rosenzweig (2004, 2008); Nunn and Qian (2011); Michaels, Rauch, and Redding (2012); and Hornbeck and Keskin (2015). We discuss this literature in more detail below.

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.

In particular, we obtain an exogenous measure of technological change in agriculture by using estimates of potential crop yields across geographical areas of Brazil from the Food and Agriculture Organization (FAO)'s Global Agro-Ecological Zones (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 goods can move across geographical areas of Brazil, but 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, which are assumed to behave as the small open economy described in the model.

We find that municipalities where the new technology is predicted to generate a larger increase in potential yields of soy were indeed characterized by faster adoption of GE soy. In addition, these regions experienced an increase in the value of output per worker and a reduction in agricultural labor intensity. Besides, the local industrial sector was characterized by faster employment growth and a reduction in wages. Interestingly, the effects of technology adoption were different for maize. Regions where 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. However, they also experienced an increase in agricultural labor intensity, a reduction in industrial employment, and faster growth in wages.

The different 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. Land-augmenting technical change, the case of second-harvest maize, leads to an increase in the marginal product of labor in agriculture and a reduction in industrial employment. However, labor-augmenting technical change, the case of GE soy, leads to a reduction in the marginal product of labor in agriculture and employment growth of the industrial sector. Thus, in what follows we refer to labor-augmenting technical change as labor-saving.⁵

⁵A formal definition of labor-saving technical change is contained in Section II.

Our estimates can be used to quantify the effect of local labor-saving agricultural technical change on local structural transformation. In particular, we compute the elasticity of local sectoral employment shares to changes in agricultural productivity induced by soy technical change: a 1 percent increase in agricultural labor productivity leads to a 0.16 percentage point decrease in the agricultural employment share and an increase in the manufacturing employment share of a similar magnitude. These estimates can be used to understand to what extent the observed differences in the speed of structural transformation across Brazilian municipalities can be explained by labor-saving technical change in soy. In the year 2000, the average municipality had employment shares in agriculture and manufacturing of 38 and 10 percent, respectively. During the next decade, the degree of labor reallocation across sectors varied extensively across municipalities. Our estimates imply that labor-saving technical change in soy can explain 24 percent of the observed differences in the reduction of the agricultural employment share across Brazilian municipalities and 31 percent of the corresponding differences in the growth of the manufacturing employment share.

We assess the robustness of our estimates to a number of deviations from our baseline framework. First, estimates are stable when we augment our empirical specification to allow municipalities with different initial levels of development to be on different structural transformation trends. Second, we obtain similar estimates in the subsample of Brazilian municipalities where the agricultural frontier did not expand. Third, our estimates are not driven by preexisting trends in manufacturing employment nor migration flows. Fourth, our results are robust to using a larger unit of observation, microregions. Fifth, at least 60 percent 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, our estimates are not driven by contemporaneous changes in commodity prices. Seventh, our main estimates remain statistically significant when we correct standard errors to account for spatial correlation.

We complement our findings with an analysis of the service sector. For this purpose, we extend the theoretical model by incorporating nontraded services. A central feature of the analysis is the distinction between two effects of labor-saving technical change in agriculture: the supply effect and the demand effect. The supply effect is generated by the reduction in the marginal product of labor in the agricultural sector, which reduces agricultural employment. The demand effect is generated by higher income resulting from agricultural productivity growth, which leads to larger consumption and employment in the service sector. As a result, the net effect of agricultural technical change on industrialization depends on the relative strength of these opposing effects. In addition, the demand effect is driven by the increase in land rents, thus its strength depends on the extent to which landowners consume services in the region where their land is located. When we turn to the data, we find that local labor-saving technical change does not significantly affect local employment in the service sector. This finding is consistent with information from the Agricultural Census suggesting that the share of land owned by resident landlords is small. Note, however, that these findings do not necessarily imply that agricultural technical change did not have an effect on the demand for services

in the aggregate Brazilian economy. This is because the difference-in-differences empirical strategy is not suitable to identify aggregate demand effects when land-owners do not consume services in the regions where their land is located. Thus, a further investigation of the effect of agricultural technical change on the aggregate demand for services is left for future work.

Finally, we investigate the impact of agricultural technical change on migration flows. In our model labor is assumed to be immobile across municipalities, thus all the adjustment to labor-saving technological change occurs through a reallocation of labor toward the manufacturing sector. However, if workers could relocate to other municipalities, some of this adjustment would occur through out-migration. Indeed, we find that municipalities with larger increases in potential soy yields experienced a net outflow of migrants between 2000 and 2010. Our estimates imply that the presence of migration flows across municipalities dampens the effects of technical change on sectoral employment shares, as around one-third of the adjustment occurs through migration flows.

Related Literature.—There is a long tradition in economics of studying the links between agricultural productivity and industrial development. Nurkse (1953); Schultz (1953); and Rostow (1960) argued that agricultural productivity growth was an essential precondition for the industrial revolution. 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 raises income per capita, which generates demand for manufacturing goods if preferences are not homothetic. The higher relative demand for manufactures generates a reallocation of labor away from agriculture (Murphy, Shleifer, and Vishny 1989; Kongsamut, Rebelo, and Xie 2001; Gollin, Parente, and Rogerson 2002). 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 for agricultural goods does not grow as fast as productivity and labor reallocates toward manufacturing (Baumol 1967; Ngai and Pissarides 2007).⁶

The view that increases in 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 toward the comparative advantage sector (Mokyr 1976; Field 1978; and Wright 1979). Their ideas were formalized by Matsuyama (1992), who showed that 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 input to production, thus technical change is, by definition, Hicks-neutral. In our model there are two factors, land and labor, and the two are complements in agricultural production. Thus, technical

⁶Another 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; Acemoglu and Guerrieri 2008). For a recent survey of the structural transformation literature see Herrendorf, Rogerson, and Valentinyi (2013a).

change can be factor-biased. In this setting, a new prediction emerges: when technical change is labor-augmenting, an increase in agricultural productivity leads to a reallocation of labor toward the industrial sector even in open economies.⁷

Our work 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 maize, rice, sorghum, and wheat during the Green Revolution in India. To guide empirical work, they present a model in which agricultural and manufacturing goods are tradable and technical change is Hicks-neutral. Consistent with their model, they find that villages with larger improvements in crop yields experienced lower manufacturing growth. Our findings are in line with theirs in the case of maize, for which technical change is land-augmenting. However, we find the opposite effects in the case of soy, for which technical change is labor-saving. Thus, relative to theirs, our work highlights the importance of the factor-bias of technical change in shaping the relationship between agricultural productivity and industrial development in open economies.

Our treatment of services in the model follows the literature on the Dutch Disease. In particular, Corden and Neary (1982) consider a three-sector open economy model with nontraded goods. One of the traded sectors is extractive and experiences a boom, which leads to deindustrialization and an expansion of the service sector. We build on their distinction between two effects of the boom: the spending effect and the resource movement effect, which we call the demand and supply effects. Our setting differs in that we consider labor-saving technical change which reduces the marginal product of labor in the booming sector, agriculture. Thus, in our model the net effect of agricultural technical change on industrialization depends on the relative strength of these two opposing effects.

Our research also connects to the literature studying the role of manufacturing in economic development. This literature has shown that a reallocation of labor into manufacturing can increase aggregate productivity. First, when labor productivity is lower in agriculture than in the rest of the economy (Gollin, Parente, and Rogerson 2002; Lagakos and Waugh 2013; and Gollin, Lagakos, and Waugh 2014). Second, when the manufacturing sector is characterized by economies of scale generated by on-the-job accumulation of human capital such as learning-by-doing (Krugman 1987; Lucas 1988; and Matsuyama 1992).

Finally, our work is 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); the effects of agriculture on local economic activity (Hornbeck and Keskin 2015); and the role of out-migration from rural areas in favoring the adoption of capital-intensive agricultural technologies (Hornbeck and Naidu 2014).

The remaining of the paper is organized as follows. Section I gives background information on agriculture in Brazil. Section II presents the theoretical model.

⁷This prediction rests on the assumptions that land and labor are strong complements in agricultural production, and land is only used in the agricultural sector. This last assumption is not necessary to obtain the prediction. Refer to the general discussion of the effects of technical change in an open economy with two goods and two factors in Findlay and Grubert (1959).

⁸This literature is surveyed by Syrquin (1988) and Foster and Rosenzweig (2008).

Section III describes the data. Section IV presents the empirical strategy and results. Section V shows a set of robustness checks on our main results. Section VI concludes.

I. Agriculture in Brazil

In this section we provide background information on recent technological developments in Brazilian agriculture. In particular, we focus on two new agricultural technologies for the cultivation of soy and maize. The first is the use of genetically engineered (GE) seeds in soy cultivation. The second is the introduction of a second harvesting season for maize during the same agricultural year.

A. Technical Change in Soy: Genetically Engineered Seeds

The main advantage of GE soy seeds relative to traditional ones is that they are herbicide-resistant, which facilitates the use of no-tillage planting techniques.⁹ 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 selectively eliminates 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.¹⁰

The first generation of GE soy seeds, the Roundup Ready (RR) variety, was commercially released in the United States in 1996 by the agricultural biotechnology firm Monsanto. In 1998, the Comissão Técnica Nacional de Biossegurança (CTNBio) authorized Monsanto to field-test GE soy in Brazil for five years as a first step before commercialization. Finally, in 2003, the Brazilian government authorized the planting and commercialization of GE soy seeds.¹¹ Prior to legalization, smuggling of GE soy seeds from Argentina was detected since 2001 according to the Foreign Agricultural Service of the United States Department of Agriculture (USDA 2001).

The new technology was characterized by fast adoption rates: in 2006 GE seeds were planted in 46.4 percent of the area cultivated with soy in Brazil, according to

⁹Genetic engineering (GE) techniques allow a precise alteration of a plant's traits. This allows targeting 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 which was herbicide-resistant.

¹⁰GE soybeans seeds allow farmers to adopt a new set of techniques that lowers labor requirements 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. 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, Klotz-Ingram, and Jans 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 who 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 United States reports no improvements in yield with respect to conventional soybeans (Fernandez-Cornejo and Caswell 2006).

¹¹In 2003, the commercialization of GE soy was permitted for one harvesting season, requiring farmers to burn all unsold stocks after the harvest (law 10.688). This temporary measure was renewed in 2004. Finally, in 2005, a new Bio-Safety Law authorized production and commercialization of GE soy in its Roundup Ready variety (law 11.105, art. 35).




FIGURE 1. SOY AND MAIZE IN BRAZIL, 1980–2011

Notes: Data sources are CONAB and PNAD. In panels A–D we exclude the states of Rondonia, Acre, Amazonas, Roraima, Pará, Amapá, Tocantins, Mato Grosso do Sul, Goiás, and Distrito Federal due to incomplete coverage by PNAD in the early years of the sample. See online Appendix for details.

the last agricultural census (IBGE 2006). In the following years the technology continued spreading to the point that it covered 85 percent of the area planted with soy in Brazil by the 2011–2012 harvesting season, according to the Foreign Agricultural Service of the USDA (USDA 2012).

The timing of adoption of GE soy seeds coincides with an increase in labor productivity and a fast expansion in the area planted with soy in Brazil. Panel A of Figure 1 documents that soy labor productivity has been increasing in Brazil since the early 1990s, and accelerated sharply in the early 2000s: soy production per worker went from 100 tonnes per worker in 2003 to around 300 tonnes per worker in 2011. Labor productivity growth was accompanied by an expansion in the area planted with soy. Table 1 reports land use by agricultural activity according to the 1996 and 2006 agricultural censuses. It shows that the area cultivated with seasonal

TABLE 1—LAND USE AND LABOR INTENSITY BY AGRICULTURAL ACTIVITY

Principal activity	Land use (million ha)		Labor intensity (workers per 1,000 ha)	
	1996	2006	1996	2006
Permanent crops	7.5	11.7	127.2	126.7
Seasonal crops	34.3	44.6	110.8	83.7
Soy	9.2	17.9	28.6	17.1
Cereals	14.3	15.3	92.5	94.9
Other seasonal crops	10.8	11.4	169.6	129.8
Cattle ranching	177.7	168.4	25.9	30.6
Forestry	110.7	91.7	34.4	42.6
Not usable	15.2	8.3	NA	NA
Other uses	8.3	9.0	NA	NA
Total	353.6	333.7	55.6	49.7

Notes: Cereals include rice, wheat, and maize (among others). Other seasonal crops include cotton, sugarcane, tobacco, cassava, and beans (among others). Permanent crops include coffee and cocoa (among others). Not usable land includes lakes and areas that are not suitable for either crop cultivation or cattle ranching. Data source is the agricultural census. See online Appendix for details.

crops increased by 10.3 million hectares between 1996 and 2006.¹² Out of these, 8.7 million hectares were converted to soy cultivation. Similarly, panel B of Figure 1 shows that the area planted with soy has been growing since the 1980s, and experienced a sharp acceleration in the early 2000s.¹³

The adoption of GE soy can affect labor demand in the agricultural sector through two channels: the within-crop and the across-crop effects. The first effect is due to a reduction in the amount of agricultural workers per hectare required to cultivate soy: labor intensity of soy production fell from 29 workers per 1,000 hectares in 1996 to 17 workers per 1,000 hectares in 2006 (Table 1). The timing of this change in labor intensity is illustrated by panel C, which shows a sharp increase in the area planted per worker in soy production in the early 2000s.¹⁴ This reduction in labor intensity was strong enough to entirely offset the potential increase in labor demand for soy due to the expansion in the area planted. As a result, employment in soy production experienced a constant decrease during the period under study (panel D).

The second channel through which the adoption of GE soy can affect labor demand is the across-crop effect. This effect is due to the expansion of soy cultivation over areas previously devoted to other crops. This effect reduces the labor intensity of production in the agricultural sector because soy production is one of

¹² Seasonal crops are those produced from plants that need to be replanted after each harvest, such as soy and maize.

¹³ Yearly data on area planted are sourced from the surveys conducted by CONAB (Companhia Nacional de Abastecimento, an agency within the Brazilian Ministry of Agriculture). These surveys of farmers and agronomists monitor the annual harvests of major crops in Brazil and are representative at the state level. Because our unit of analysis is the municipality, we only use data from the CONAB survey to illustrate the timing of the evolution of aggregate agricultural outcomes during the period under study. In the empirical analysis, instead, we use data from the agricultural census which covers all farms in the country.

¹⁴ Notice that the decrease in labor intensity in soy production between 1996 and 2006 implied by panel C of Figure 1 is larger than the one reported in the text and Table 1. This is because the data sources are different. Panel C displays yearly data on area planted with soy from the CONAB survey and yearly data on employment in soy production from the PNAD survey. Table 1 instead is based on data on area planted and employment from the agricultural censuses of 1996 and 2006. See the online Appendix for a more detailed discussion of data sources for panel C.

TABLE 2—SUMMARY STATISTICS OF MAIN VARIABLES AT MUNICIPALITY LEVEL

	1996		1996–2006			
	Mean	SD	Mean	SD	Observations	
<i>Panel A. Agricultural census</i>						
log output per worker	7.690	1.192	0.561	0.762	4,149	
log labor intensity	–2.585	1.048	–0.027	0.551	4,149	
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			
	Mean	SD	Mean	SD	Observations	
<i>Panel B. Population census</i>						
Employment shares:						
Agriculture	0.383	0.189	–0.064	0.074	4,149	
Manufacturing	0.104	0.090	0.014	0.057	4,149	
Services	0.362	0.136	0.032	0.057	4,149	
Other sectors	0.151	0.054	0.018	0.038	4,149	
log employment in manufacturing	5.885	1.580	0.221	0.608	4,149	
log wage in manufacturing	5.541	0.500	0.287	0.365	4,149	
	1991–2000			2000–2010		
	Mean	SD	Observations	Mean	SD	Observations
<i>Panel C. Migration</i>						
Net migration rate	–0.036	0.181	3,992	–0.024	0.124	4,149
	Low inputs		High inputs		Difference	
	Mean	SD	Mean	SD	Mean	SD
<i>Panel D. FAO-GAEZ</i>						
Potential yield in soy	0.302	0.154	2.113	0.938	1.811	0.851
Potential yield in maize	0.992	0.494	4.066	2.197	3.073	1.811

Note: See online Appendix for a detailed description of each variable.

the least labor-intensive agricultural activities: its production required 17 workers per 1,000 hectares while seasonal crops and permanent crops require 84 and 127, respectively (Table 1).

B. 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 during the spring season, between August and December. At the beginning of the 1980s a few farmers in the South-East region of Brazil started producing maize after the summer harvest, between March and July. This second season of maize cultivation spread across Brazil, where it is known as *milho safrinha* (small-harvest maize). Panel E of Figure 1 shows that the area devoted to second season maize has expanded steadily since the beginning of the 1990s, although the total area devoted to maize has increased only slightly.¹⁵

Cultivation of a second season of maize requires the use of modern cultivation techniques for the following reasons. First, more intensive land use removes nitrogen

¹⁵Data on area cultivated with maize broken down by the season of harvest are publicly available only at the aggregate level. For this reason in Section IV, when we study municipality-level data, we are not able to distinguish between maize cultivation in different seasons.

from the soil, which needs to be replaced by fertilizers. Second, the planting of a second crop requires careful timing, as yields drop considerably due to late planting. Third, herbicides are used to remove residuals from the first harvest on time to plant the second crop. Finally, the second season crop needs to be planted one month faster than the first, which usually requires higher mechanization.¹⁶

The introduction of a second harvesting season for maize can affect labor demand in the agricultural sector through the within-crop and across-crop effects described above. The first effect is directly due to the introduction of a second harvest which raises labor demand relative to the benchmark of a single maize harvest. The second effect is due to the expansion of maize over areas previously dedicated to less-labor intensive activities, which also tends to increase labor demand. According to the 1996 agricultural census, maize cultivation is more labor intensive than the main agricultural activities in Brazil. In this year, labor intensity in maize production was 100 workers per 1,000 hectares, above the labor intensity of soy, other cereals, and cattle ranching.¹⁷

II. 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 region that behaves as a small open economy in the sense that goods are freely tradable across regions but production factors are immobile. There are two sectors, agriculture and manufacturing, and two production factors, land and labor.

A. Setup

This small open economy has L residents, each endowed with one unit of labor. There are two sectors, manufacturing and agriculture, both of which produce tradable goods. Production of the manufactured good requires only labor and labor productivity in manufacturing is A_m . As a result, $Q_m = A_m L_m$, 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:

$$(1) \quad Q_a = A_N \left[\gamma (A_L L_a)^{\frac{\sigma-1}{\sigma}} + (1 - \gamma) (A_T T_a)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

where Q_a denotes production of the agricultural good, the two production factors are labor (L_a) and land (T_a), A_N is Hicks-neutral technical change, A_L is labor-augmenting technical change, and A_T is land-augmenting technical change. The parameter $\sigma > 0$

¹⁶For a more detailed discussion, see EMBRAPA (2006) and CONAB (2012).

¹⁷Labor intensity of each agricultural activity is reported in Table 1. Information on the area and number of workers employed in farms whose main activity is maize production is publicly available only for the agricultural census of 1996. In Table 1 we therefore report labor intensity for the “all cereals” category, which we also observe in 2006 and includes rice, wheat, maize, and other cereals. For a measure of maize labor intensity under advanced cultivation techniques, we refer to data for the United States. The USDA Agricultural Resources Management Survey (ARMS) reports that maize is more labor intensive than soy: 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.

captures the elasticity of substitution between land and labor, and $\gamma \in (0, 1)$. The production function described by equation (1) implies the following marginal product of labor:

$$(2) \quad MPL_a = A_N A_L \gamma \left[\gamma + (1 - \gamma) \left(\frac{A_T T}{A_L L_a} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}}.$$

This expression shows that Hicks-neutral and land-augmenting technical change increase the marginal product of labor. However, labor-augmenting technical change generates two opposing effects on the marginal product of labor. First, increases in A_L imply that each worker is more productive, as can be seen in the first term of the equation. Second, a larger A_L generates a reduction in the amount of land per unit of labor in efficiency units ($A_T T / A_L L_a$), which tends to reduce the marginal product of labor. This second effect is larger when land and labor are poor substitutes. Thus, the relative strength of the two opposing effects depends on the value of the parameter σ . In particular, $\partial MPL_a / \partial A_L < 0$ when the elasticity of substitution is smaller than the land share of output, $\sigma < 1 - \Gamma \equiv T_a MPL_a / Q_a$, as shown in the online Appendix. In what follows, we say that technical change is strongly labor-saving when this condition is satisfied.^{18, 19}

B. Equilibrium

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

$$(3) \quad P_a MPL_a = w = P_m MPL_m.$$

As a result, in equilibrium, the marginal product of labor in agriculture is determined by international prices and manufacturing productivity: $MPL_a = (P_m / P_a)^* A_m$. This condition and the land market clearing condition ($T_a = T$) determine the equilibrium allocation of labor,

$$(4) \quad L_a^* = \frac{A_T T}{A_L} \left\{ \frac{\gamma}{1 - \gamma} \frac{1 - \Gamma^*}{\Gamma^*} \right\}^{\frac{\sigma}{1 - \sigma}},$$

where the equilibrium labor share is $\Gamma^* = \gamma^\sigma (P_m A_m / P_a A_N A_L)^{1 - \sigma}$. In turn, the equilibrium level of employment in manufacturing, L_m^* , can be obtained using the labor market clearing condition, $L_m + L_a = L$. Once L_m^* and L_a^* are determined, output

¹⁸Note that, because the production function takes the constant elasticity of substitution (CES) form, the land share of output is a function of the equilibrium level of employment in agriculture. In particular, in the relevant case where $\sigma < 1$ the land share is increasing on the level of agricultural employment. As a result, this condition is more likely to be satisfied when the equilibrium level of agricultural employment is high.

¹⁹See Neary (1981) and Acemoglu (2010) for more general discussions of the properties of technical change that reduce the marginal product of labor. We follow Acemoglu in using the term *strongly* labor-saving.

in each sector can be found using the production functions described in Section IIA. See the online Appendix for detailed derivations.

C. Technological Change and Structural Transformation

In this section we assess the response of agricultural and manufacturing employment to three types of technological change: labor-augmenting, land-augmenting, and Hicks-neutral.

Labor-Augmenting Technical Change.—The effect of labor-augmenting technical change on agricultural employment depends on whether the elasticity of substitution is smaller than the equilibrium land share of agricultural production ($\sigma < 1 - \Gamma^*$). When this condition is satisfied, we say that land and labor are strong complements. We consider the two possible parameter configurations below.

- (i) Land and labor are strong complements, $\frac{\partial L_a^*}{\partial A_L} < 0$ and $\frac{\partial L_m^*}{\partial A_L} > 0$.

An increase in A_L generates a reallocation of labor from agriculture to manufacturing. This is because when the elasticity of substitution between land and labor is smaller than the land share of output, labor-augmenting technical change reduces the marginal product of labor in agriculture. In equilibrium, the marginal product of labor is given by international prices and manufacturing productivity, thus it must stay constant when A_L increases. Thus, employment in agriculture must fall to keep the marginal product of labor at the equilibrium level.²⁰

PROOF:

See online Appendix.

- (ii) Land and labor are not strong complements, $\frac{\partial L_a^*}{\partial A_L} > 0$ and $\frac{\partial L_m^*}{\partial A_L} < 0$.

An increase in A_L generates a reallocation of labor from manufacturing to agriculture. This is because when the elasticity of substitution is larger than the land share of output, labor-augmenting technical change induces an increase in the marginal product of labor in agriculture.

Land-Augmenting Technical Change, $\frac{\partial L_a^*}{\partial A_T} > 0$ and $\frac{\partial L_m^*}{\partial A_T} < 0$.

An increase in A_T generates a reallocation of labor from manufacturing to agriculture. To see why this is the case, note that land-augmenting technical change rises the marginal product of labor in agriculture (see equation (2)).

Hicks-Neutral Technical Change, $\frac{\partial L_a^*}{\partial A_N} > 0$ and $\frac{\partial L_m^*}{\partial A_N} < 0$.

²⁰See Figure A2 in the online Appendix for a graphical representation of these effects.

An increase in A_N generates a reallocation of labor from manufacturing to agriculture. This is because a Hicks-neutral increase in agricultural productivity rises the marginal product of labor in agriculture (see equation (2)).

D. Empirical Predictions

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. In the case of maize, farmers started introducing advanced cultivation techniques and inputs which permit to grow two crops a year, effectively increasing the land endowment. Thus, this new technology can be characterized as land-augmenting technical change. In our empirical analysis, we quantify the effects of these two types of technical change on observable variables in the agricultural and manufacturing sector and test whether they display the sign patterns predicted by the model.

We analyze data aggregated at the municipality level, which is our unit of analysis. As a result, we interpret the production functions in the model as describing the aggregate level of agricultural and manufacturing production (Q_a and Q_m) in a given municipality. In addition, the agricultural census reports information on employment aggregated across agricultural activities. Thus, we interpret equation (1) as describing the aggregate production function for the agricultural sector, where $P_a Q_a$ is the value of agricultural output, L_a is agricultural employment, and T_a is land in agricultural establishments. We trace the effects of the two new agricultural technologies on these directly observed variables to test the following predictions of the model.

PREDICTION 1: *If land and labor are strong complements in production, labor-augmenting technical change in agriculture (A_L):*

- (i) *increases the value of output per worker, $\frac{P_a^* Q_a^*}{L_a^*}$;*
- (ii) *reduces the labor intensity of production, $\frac{L_a^*}{T}$;*
- (iii) *reduces the employment share of agriculture, $\frac{L_a^*}{L}$;*
- (iv) *increases the employment share of manufacturing, $\frac{L_m^*}{L}$.*

PROOF:

See online Appendix.

PREDICTION 2: *Land-augmenting technical change in agriculture (A_T):*

- (i) *does not change the value of output per worker;*

- (ii) *increases the labor intensity of production;*
- (iii) *increases the employment share of agriculture;*
- (iv) *reduces the employment share of manufacturing.*

PROOF:

See online Appendix.

E. Services

In this section we extend the model by including a third sector which produces nontraded services. The purpose of this extension is to understand to what extent the predictions of the model discussed above are modified by the presence of nontraded goods. A detailed analysis of the model with services is contained in the online Appendix.

We assume that the production function for services uses only labor and displays constant returns to scale. As a result, $Q_s = A_s L_s$, where Q_s denotes production of services and L_s denotes labor allocated to the service sector. Note that because services are nontradable, production can no longer be determined independently of consumption and income. Thus, we specify preferences and factor ownership. Consumers have the following Cobb-Douglas preferences over the three goods,

$$(5) \quad U(c_a, c_m, c_s) = c_a^{\alpha_a} c_m^{\alpha_m} c_s^{\alpha_s},$$

where $\alpha_a + \alpha_m + \alpha_s = 1$.²¹ There are two types of agents in the economy: L workers, each endowed with one unit of labor; and T landowners, each endowed with one unit of land. We assume that workers reside in the same region where they work. In contrast, landowners can reside in any region. We denote by θ the share of landowners residing in the same region where their land is located. Then, aggregate service consumption in a region is $C_s = c_{s,L}L + c_{s,T}\theta T$, where $c_{s,L}$ is the consumption of workers and $c_{s,T}$ the consumption of landowners.^{22,23}

²¹Our use of a homothetic utility function follows the findings in Herrendorf, Rogerson, and Valentinyi (2013b). They show that a homothetic utility function where the elasticity of substitution across sectors is smaller than one provides the best fit to the postwar US data when sectoral consumption data is measured in terms of value-added. Because we use data on employment to measure structural transformation, our analysis tracks value-added better than final goods consumption. As a result we use a homothetic utility function. However, we assume that the elasticity of substitution across sectors is equal to 1 to make the model simpler. We discuss below how the predictions of our model would be modified if this elasticity was smaller than 1.

²²Our treatment of landowners nests the two standard assumptions in the regional economics literature as discussed by Fujita (1989). The first is that land income accrues to absentee landowners and is thus not spent within the region, which corresponds to $\theta = 0$. The second is that land income is redistributed lump-sum to workers, which corresponds to $\theta = 1$, because preferences are homothetic. Note that this treatment implicitly assumes that absentee landowners reside outside the country. This is because we do not take into account the local consumption of landowners who reside in the region under consideration but own land in other regions. In the online Appendix, we also consider an alternative scenario where all landowners reside within the country but not necessarily in the region where they own land.

²³Note that θ is the share of services consumption of landowners that is spent locally. Thus, an alternative interpretation is that landowners reside in the region where they own land but buy some services in other regions.

In this setting, equilibrium employment in agriculture is the same as in the model without nontraded services, given by equation (4). This is because wages are set by the value of the marginal product of labor in manufacturing. Thus, the effects of agricultural technical change on agricultural employment are identical to the ones in the model without services. We call them the supply-side effects of technical change: $\frac{\partial L_a^*}{\partial A_i}$ for $i = N, T, L$.

In turn, equilibrium employment in services can be written as

$$(6) \quad L_s^* = \alpha_s L + \alpha_s \theta \frac{r^*}{w^*} T,$$

where r^* is the equilibrium land rent.²⁴ Note that workers spend a constant share of their labor endowment on services ($\alpha_s L$). This is because the service sector uses only labor for production. Thus, any increase in wages has both an income and substitution effect on the demand for services by workers. The income effect increases their demand for services as their labor endowment is more valuable. The substitution effect reduces the demand for services as their price, the wage, increases. When preferences are Cobb-Douglas, both effects have the same magnitude and cancel out.²⁵ As a result, agricultural technical change can only affect the demand for services through its effect on the consumption of landowners: $\alpha_s \theta \frac{r^*}{w^*}$. In turn, agricultural technical change always increases land rents. Thus, the demand for services and employment in the service sector increase. We call this the demand side effects of technical change: $\frac{\partial L_s^*}{\partial A_i}$ for $i = N, T, L$.

When technical change is Hicks-neutral or land-augmenting, both the supply-side and demand-side effects reduce manufacturing employment. However, when technical change is strongly labor-saving each effect moves manufacturing employment in opposite directions. On the one hand, the supply-side effect releases labor from agriculture, increasing the labor supply for manufacturing. On the other hand, the demand-side effect increases labor demand in services, reducing the supply of labor for manufacturing. Therefore, the net effect on manufacturing employment depends on the relative strength of each effect. In the online Appendix, we show that the supply-side effect dominates as long as $\sigma < (1 - \Gamma^*)(1 - \alpha_s \theta)$. Note that because $1 - \alpha_s \theta < 1$, this condition is stronger than the condition required for agricultural technical change to be strongly labor-saving: $\sigma < 1 - \Gamma^*$. Thus, it is satisfied as long as landowners' consumption share of local services ($\alpha_s \theta$) is not too large.

III. Data

The main data sources are the agricultural census, the population census, and the FAO Global Agro-Ecological Zones database. To perform robustness checks we

²⁴ See the online Appendix for detailed derivations and closed-form solutions for r^* and L_s^* .

²⁵ If, instead of Cobb-Douglas, preferences were homothetic with an elasticity of substitution smaller than 1, as suggested by Herrendorf, Rogerson, and Valentinyi (2013b), the income effect would dominate. Thus, the demand for services from workers would be increasing in wages.

also use manufacturing plant-level data from the Brazilian Annual Industrial Survey (PIA).²⁶

The agricultural census is released at intervals of ten years by the Instituto Brasileiro de Geografia e Estatística (IBGE), the Brazilian National Statistical Institute. The empirical analysis focuses on the last two rounds of the census which have been carried out in 1996 and in 2006. The agricultural census data are collected through direct interviews with the managers of each agricultural establishment and are made available online by the IBGE aggregated at municipality level.²⁷ The agricultural variables of interest are the share of land planted with soy and maize, the value of production per worker, and labor intensity.²⁸ The last two variables are aggregated across all agricultural activities. This is because the unit of observation in the census is the agricultural establishment, and these tend to perform several activities. As a result, it is not possible to obtain a measure of employment by crop.

We use the Brazilian population census to construct measures of the sectoral composition of employment and average wages. The population census is conducted every ten years and it covers the entire Brazilian population. We use data from the last two rounds of the census (2000 and 2010): this allows us to observe the variables of interest before and after the legalization of the GE soy seeds.²⁹ Data on the sector of employment are collected through a special survey that is administered to a representative sample of the Brazilian population within narrow cells defined by geographical district, sex, age, and urban or rural residence. The variables we focus on are the sector in which the person was working during the previous week and its wage.³⁰ For each municipality, we compute employment shares as the number of workers in each sector divided by total employment.³¹ Table 2 contains summary statistics for the main variables of interest.

We obtain an exogenous measure of technological change in agriculture by using estimates of potential soy and maize 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. In addition, the database reports potential yields under different technologies or input combinations. Yields under the low technology are described as those obtained planting traditional seeds, with no use of chemicals nor mechanization. Yields under the high technology are obtained using improved high

²⁶In this section we briefly discuss the main data sources and variables of interest. For detailed variable definitions, see the online Appendix.

²⁷Borders of municipalities often change, thus, to make them comparable over time, IBGE has defined Área Mínima Comparável (AMC), smallest comparable areas, which we use as our unit of observation. The average size of an AMC in terms of population is 39,858 inhabitants, while the average size of a municipality is 30,833 inhabitants (data from the 2000 population census). In terms of area, the average AMC has an area of around 2,000 square kilometers, while the average municipality has an area of 1,500 square kilometers.

²⁸The measure of agricultural employment used to construct the value of production per worker and labor intensity includes: employees, family members employed in farm activities, sharecroppers, and people who reside in the farm and perform agricultural activities without a formal contract. There are two potential problems with this definition. The first is potential double counting of seasonal workers who work in more than one farm during the same calendar year. The second is that this variable does not include employees hired by service provider companies who are contracted by the farm to perform agricultural tasks. See the online Appendix for a detailed description of this variable.

²⁹To perform some of the robustness checks we also use the 1980 and 1991 population censuses.

³⁰The sector classification is comparable across the censuses 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.

³¹We restrict the sample to workers between 16 and 55 years old.

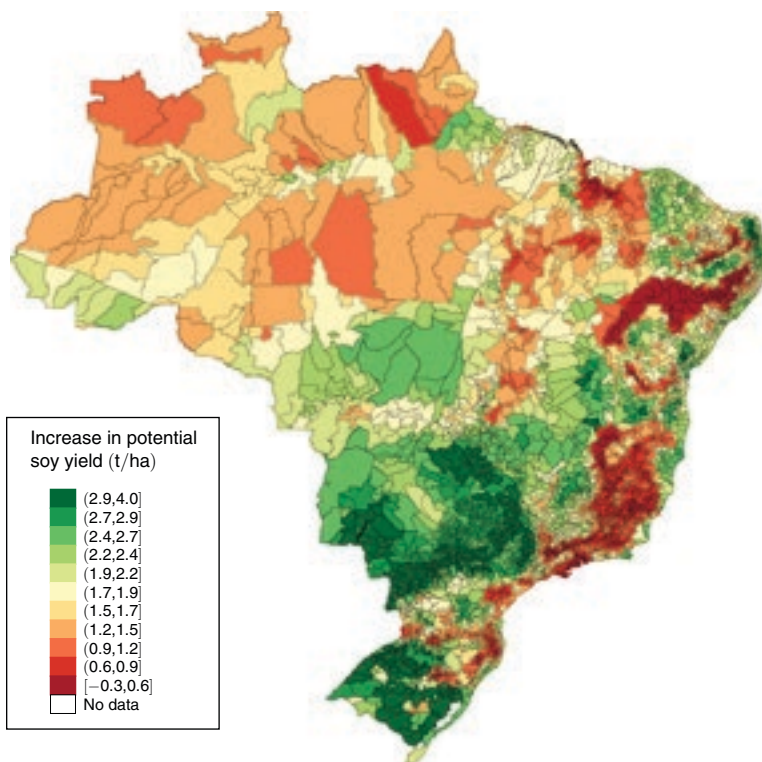


FIGURE 2. TECHNOLOGICAL CHANGE IN SOY: MUNICIPALITIES

Notes: Authors' calculations from FAO-GAEZ data. Technical change in soy production for each municipality is computed by deducting the average potential yield under low inputs from the average potential yield under high inputs.

yielding varieties, optimum application of fertilizers and herbicides, and mechanization.³² Maps displaying the resulting measures of potential yields for soy and maize under each technology are contained in online Appendix Figures A3 to A6.

We construct a measure of technical change in soy or maize production for each municipality by deducting the average potential yield under low inputs from the average potential yield under high inputs. Figure 2 illustrates the resulting measure of technical change in soy at the municipality level, while Figure 3 shows the same measure at the microregion level.

Finally, we use data from the Pesquisa Industrial Anual (PIA), the annual industrial survey conducted by the IBGE. We focus on firms operating in the manufacturing sector³³ and use yearly data from 1996 to 2007. All firms with more than five employees registered in the national firm registry (Cadastro Central de Empresas (CEMPRE)) are eligible for this survey. 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 construct measures of

³² See the online Appendix for a detailed definition of potential yields under different input combinations.

³³ Identified by the CNAE sector codes 15 to 37.

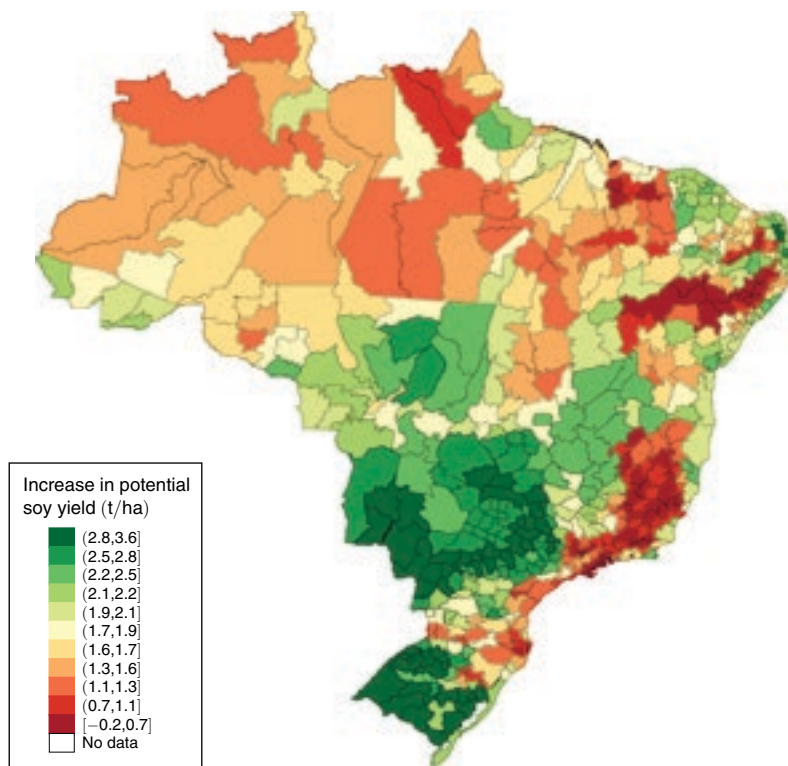


FIGURE 3. TECHNOLOGICAL CHANGE IN SOY: MICROREGIONS

Notes: Authors' calculations from FAO-GAEZ data. Technical change in soy production for each microregion is computed by deducting the average potential yield under low inputs from the average potential yield under high inputs.

total employment and average wages that are representative at municipality level by focusing on firms with 30 or more employees.

IV. 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 technical change on the allocation of labor across sectors.

In Section IVA we report simple correlations between the expansion of the area planted with soy (maize) and labor market outcomes for the agricultural and industrial sectors in each municipality. As discussed above, these correlations are not informative about the causal relation between these variables. Thus, in sections IVB to IVD, we present and implement an empirical strategy that attempts to establish the direction of causality by exploiting the timing of legalization and the differential impact of the new technology on potential yields across geographical areas.

A. 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? First, we present a set of OLS estimates of equations relating agricultural outcomes to the percentage of farm land cultivated with soy and maize. Second, we present the corresponding estimates for manufacturing outcomes. The basic form of the equations to be estimated in this section is

$$(7) \quad y_{jt} = \delta_j + \delta_t + \pi^{soy} (\text{Soy share})_{jt} + \pi^{maize} (\text{Maize share})_{jt} + \varepsilon_{jt},$$

where j indexes municipalities, t indexes time, δ_j are municipality fixed effects, δ_t are time fixed effects, y_{jt} is an outcome that varies across municipalities and time, and *Soy (Maize) share* is the total area reaped with soy (maize) divided by total farm land.³⁴ We observe agricultural outcomes for the census years 1996 and 2006. Because fixed effects and first difference estimates are identical when considering only two periods, we estimate (7) in first differences:

$$(8) \quad \Delta y_j = \Delta \delta + \pi^{soy} \Delta \text{Soy share}_j + \pi^{maize} \Delta \text{Maize share}_j + \Delta \varepsilon_j.$$

Agricultural Outcomes: Productivity, Labor Intensity, and Employment Share.—Table 3 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. 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.

The first two columns of Table 3 show that in areas where soy cultivation expanded, the value of agricultural production per worker increased and labor intensity in agriculture decreased. These empirical findings are consistent with the characterization of soy technical change as strongly labor-saving. The estimated coefficients imply that a 1 percentage point increase in soy area share corresponds to a 0.58 percent increase in labor productivity, and a 0.48 percent reduction in labor intensity. In contrast, in areas where maize cultivation expanded labor intensity increased. This evidence is consistent with our characterization of technical change in maize as land-augmenting. The estimated coefficients imply that a 1 percentage point increase in maize area share corresponds to a 1.6 percent increase in labor productivity, and a 0.74 percent 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

³⁴Total farm land includes areas devoted to crop cultivation (both permanent and seasonal crops), animal breeding, and logging.

TABLE 3—BASIC CORRELATIONS IN THE DATA: AGRICULTURE
(Productivity, Labor Intensity, and Employment Share)

	Δ log output per worker (1)	Δ log labor intensity (2)	Δ Employment share (3)
Δ Soy area share	0.583 (0.232)	-0.479 (0.154)	-0.090 (0.027)
Δ Maize area share	1.597 (0.184)	0.737 (0.119)	-0.014 (0.019)
Observations	3,765	3,765	3,765
R^2	0.023	0.008	0.003

Notes: Changes in dependent variables are calculated over the years 1996 and 2006 when the data sources are the agricultural censuses of 1996 and 2006 (columns 1 and 2), and over the years 2000 and 2010 when the data sources are the population censuses of 2000 and 2010 (column 3). Changes in explanatory variables are calculated over the years 1996 and 2006. The unit of observation is the municipality. Robust standard errors reported in parentheses.

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 3 shows that the employment share of agriculture decreased in places where soy expanded while estimates for maize are not statistically significant. The estimated coefficient implies that a 1 percentage point increase in soy area share corresponds to a 0.09 percentage point reduction in the agricultural employment share.

The finding that the agricultural employment share fell in areas where soy expanded suggests that soy technical change is not only labor-augmenting but also strongly labor-saving. In this case, our model predicts that technology adoption reduces labor demand in agriculture.

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 4 reports OLS estimates of equation (8) for three manufacturing sector outcomes: employment share, level of employment, and average wage.

The first column of Table 4 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 0.11 percentage point increase in manufacturing employment share and a 1.05 percent increase in manufacturing employment.

B. Empirical Strategy

In what follows 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

TABLE 4—BASIC CORRELATIONS IN THE DATA: MANUFACTURING
(*Employment Share, Employment, and Wages*)

	Δ Employment share (1)	Δ log employment (2)	Δ log wage (3)
Δ Soy area share	0.106 (0.022)	1.053 (0.226)	0.150 (0.113)
Δ Maize area share	0.001 (0.013)	0.018 (0.147)	-0.039 (0.080)
Observations	3,765	3,765	3,765
R^2	0.007	0.006	0.000

Notes: Changes in dependent variables are calculated over the years 2000 and 2010. Changes in explanatory variables are calculated over the years 1996 and 2006. The unit of observation is the municipality. Robust standard errors reported in parentheses.

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 labor-saving agricultural technologies reduces labor demand in the agricultural sector and induces a reallocation of labor toward the industrial sector. However, these correlations are not informative about the direction of causality. For example, they 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 the remainder of 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 show that our estimates are robust to using a larger unit of observation: microregions. Figures 2 and 3 contain maps of Brazil displaying both levels of aggregation.³⁵

We propose to identify the causal effect of the new technologies on structural transformation by exploiting the timing of adoption and their differential impact 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 United States in 1996, and legalized in Brazil in 2003. Given that the seeds were developed in the United States, their date of approval for commercialization in the United States, 1996, is arguably exogenous with respect to developments in the Brazilian economy. In

³⁵Microregions are groups of several municipalities created by the 1988 Brazilian Constitution and used for statistical purposes by IBGE.

contrast, the date of legalization, 2003, responded partly to pressure from Brazilian farmers. In addition, smuggling of GE soy 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.

Second, the new technology had 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, which we obtain from the FAO-GAEZ database. These potential yields are estimated using an agricultural model that predicts yields for each crop given climate and soil conditions. As potential yields are a function of weather and soil characteristics, not of actual yields in Brazil, they can be used as a source of exogenous variation in agricultural productivity across geographical areas. 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:

$$(9) \quad y_{jt} = \delta_j + \delta_t + \beta^{soy} A_{jt}^{soy} + \varepsilon_{jt},$$

where y_{jt} is an outcome that varies across municipalities (j) and time (t). δ_j are municipality fixed effects, δ_t are time fixed effects, and A_{jt}^{soy} is equal to the potential soy yield under high inputs from 2003 onward and to the potential soy yield. This variable takes the value corresponding to low inputs before 2003, and the value corresponding to high inputs afterwards. 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. Similarly, 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. We thus estimate a first-difference version of equation (9),

$$(10) \quad \Delta y_j = \Delta \delta + \beta^{soy} \Delta A_j^{soy} + \rho Rural_j,_{1991} + \Delta \varepsilon_j,$$

where the outcome of interest, Δy_j , is the change in outcome variables between the last two census years, and ΔA_j^{soy} is the potential yield of soy under the high technology minus the potential yield of soy under the low technology. Figure 2

contains a map of Brazilian municipalities displaying this measure of technical change. Additionally, we include a control for the share of rural population in 1991 to allow for differential trends for municipalities with different initial urbanization rates. This is important because, as mentioned above, coastal municipalities tend to have higher urbanization rates and there were migration flows from rural to urban areas during the period under study.³⁶

In the case of maize, we follow a similar empirical strategy. However, it is important to note that the cultivation techniques necessary to introduce a second harvesting season 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, the variation in adoption used in our empirical analysis is arguably exogenous to developments in the local industrial sector. As noted in Section I, the introduction of a second harvesting season for maize requires modern techniques that are intensive in the use of fertilizers, herbicides, and tractors. Then, we expect that the difference in FAO-GAEZ potential maize yields between the high and low technology captures the profitability of introducing a second harvesting season. Thus, we augment the equation described above to include the following variable: ΔA_j^{maize} , which is equal to the potential yield of maize under high inputs minus the potential yield of maize under low inputs. ΔA_j^{maize} can be thought of as the empirical counterpart of the land-augmenting technical change A_T presented in our model,

$$(11) \quad \Delta y_j = \Delta \delta + \beta^{soy} \Delta A_j^{soy} + \beta^{maize} \Delta A_j^{maize} + \rho Rural_j, 1991 + \Delta \varepsilon_j.$$

A potential concern with our identification strategy is that, although the soil and weather characteristics that drive the variation in ΔA_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, if municipalities with heterogeneous initial levels of development experienced 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 observable 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 5 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.³⁷ 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

³⁶The share of working age population residing in rural areas fell from 22 percent in 1991 to 14 percent in 2010.

³⁷Municipalities 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.

controls for differential trends across municipalities with heterogeneous initial characteristics in our baseline specification (11), as follows:

$$(12) \quad \Delta y_j = \Delta \delta + \beta^{soy} \Delta A_j^{soy} + \beta^{maize} \Delta A_j^{maize} \\ + \rho Rural_{j,1991} + \mathbf{X}'_{j,1991} \boldsymbol{\omega} + \Delta \varepsilon_j,$$

where the vector $\mathbf{X}_{j,1991}$ contains the set of municipality characteristics discussed above.

In our baseline specifications we report standard errors that are robust to heteroskedasticity. A potential concern is that these estimated standard errors may not be consistent if our exogenous measures of technical change are correlated across space. For this reason, in Section VF, we assess the robustness of our results to a set of spatial correlation patterns.

In the following sections we report estimates of the effects of technical change on agricultural production and the sectoral composition of employment. In particular, we report estimates of the effects of technical change on a set of agricultural outcomes in IVC; on manufacturing outcomes in IVD; and on service sector outcomes in IVE.

C. Agricultural Outcomes

Soy and Maize Expansion.—We start by documenting the relationship between technical change measured by the increase in the FAO-GAEZ potential yields of soy and maize, and the actual change in the share of agricultural land cultivated with each crop. The objective of this exercise is to check whether the change in potential yields is a good proxy for the profitability of adoption of the new agricultural technologies. If this is the case, we expect the increase in potential yield of a given crop to predict the actual expansion in the share of agricultural land 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 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 6: the increase in potential soy yield predicts the expansion in the share of agricultural area planted with GE soy 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 6 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

TABLE 5—COMPARING MUNICIPALITIES BELOW/ABOVE MEDIAN INCREASE IN POTENTIAL SOY YIELD

	Below ΔA^{soy} median (1)	Above ΔA^{soy} median (2)	Difference (3)
Agricultural employment share	0.500	0.443	-0.057 (0.007)
Manufacturing employment share	0.080	0.097	0.017 (0.003)
Share rural population	0.516	0.404	-0.112 (0.007)
log income per capita	4.389	4.656	0.267 (0.018)
log pop. density	3.155	3.219	0.064 (0.041)
Literacy rate	0.688	0.745	0.057 (0.005)
Observations	2,075	2,074	

Notes: Average values of observable characteristics of municipalities that rank below and above the median of ΔA^{soy} . All observable characteristics are from the population census of 1991. Column 3 reports the difference between columns 2 and 1, along with its standard error.

planted area with soy instead of GE soy.³⁸ 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 7. First, note that while soy technical change has a positive effect on the area 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 changes 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 1 standard deviation increase in potential soy yield corresponds to an increase in the soy share of agricultural land of 0.26 of a standard deviation. To understand the magnitude of our estimate, this is an increase of agricultural land devoted to soy by 877 hectares in response to a 0.85 tons per hectare increase in potential soy yield. The corresponding estimate for maize implies that a 1 standard deviation increase in potential maize yield corresponds to a 0.08 of a standard deviation increase in the maize share of agricultural land. This means that, in response to a 1.8 tons per hectare

³⁸In the case of maize, we can only focus on the broader measure of area planted with maize as the publicly available agricultural census data do not contain information on the season of planting of maize at the municipality level.

TABLE 6—THE EFFECT OF TECHNOLOGICAL CHANGE ON AGRICULTURE:
GE SOY ADOPTION

	Δ GE soy area share		Δ Non-GE soy area share	
	(1)	(2)	(3)	(4)
ΔA^{soy}	0.021 (0.002)	0.019 (0.002)	-0.009 (0.002)	-0.009 (0.002)
Share rural population	0.039 (0.005)	0.085 (0.008)	-0.017 (0.004)	-0.044 (0.007)
log income per capita		-0.000 (0.003)		0.001 (0.003)
log pop. density		0.003 (0.001)		-0.005 (0.001)
Literacy rate		0.114 (0.011)		-0.048 (0.010)
Observations	3,652	3,652	3,652	3,652
R^2	0.083	0.162	0.019	0.044

Notes: Changes in dependent variables are calculated over the years 1996 and 2006. All municipality controls are from the population census of 1991. The unit of observation is the municipality. Robust standard errors reported in parentheses.

increase in potential maize yield, agricultural land devoted to maize increases by 426 hectares.

Agricultural Productivity, Labor Intensity, and Employment Share.—Next, we study the effects of agricultural technical change on agricultural production and employment. Table 8 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.

Estimates reported in columns 1 and 3 indicate that areas where potential soy yields increased relatively more experienced a larger increase in the value of agricultural production per worker and a larger reduction in labor intensity between 1996 and 2006. Next, we study the effect of agricultural technical change in soy on the agricultural employment share. Estimates reported in column 5 indicate that areas with a larger increase in potential soy yield experienced a faster reduction in the agricultural employment share between 2000 and 2010. Note that estimated coefficients are stable or slightly larger when we control for lagged municipality characteristics in columns 2, 4, and 6. This finding indicates that our estimates are not capturing differential growth trends across municipalities. Because technical change in soy is characterized as labor-augmenting, these empirical findings are consistent with the predictions of the model for the case where land and labor are strong complements in agricultural production (see Prediction 1). Thus, the estimates of the effects of soy technical change reported in Table 8 imply that technical change in soy was strongly labor-saving and confirm the conclusions drawn from the simple correlations in the data reported in Table 3.

The estimates discussed above 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

TABLE 7—THE EFFECT OF TECHNOLOGICAL CHANGE ON AGRICULTURE
(Soy and Maize Expansion)

	Δ Soy area share		Δ Maize area share	
	(1)	(2)	(3)	(4)
ΔA^{soy}	0.013 (0.001)	0.013 (0.002)		0.001 (0.003)
ΔA^{maize}		-0.001 (0.001)	0.003 (0.001)	0.003 (0.001)
Share rural population	0.020 (0.003)	0.039 (0.005)	0.011 (0.004)	0.010 (0.007)
log income per capita		0.001 (0.002)		-0.005 (0.004)
log pop. density		-0.002 (0.000)		0.004 (0.001)
Literacy rate		0.064 (0.007)		-0.006 (0.012)
Observations	3,652	3,652	3,652	3,652
R^2	0.067	0.124	0.009	0.015

Notes: Changes in dependent variables are calculated over the years 1996 and 2006. All municipality controls are from the population census of 1991. The unit of observation is the municipality. Robust standard errors reported in parentheses.

TABLE 8—THE EFFECT OF TECHNOLOGICAL CHANGE ON AGRICULTURE
(Productivity, Labor Intensity, and Employment Share)

	Δ log output per worker		Δ log labor intensity		Δ Employment share	
	(1)	(2)	(3)	(4)	(5)	(6)
ΔA^{soy}	0.115 (0.024)	0.131 (0.026)	-0.057 (0.018)	-0.064 (0.021)	-0.018 (0.002)	-0.021 (0.002)
ΔA^{maize}	-0.025 (0.011)	-0.033 (0.011)	0.031 (0.008)	0.033 (0.009)	0.005 (0.001)	0.006 (0.001)
Share rural population	0.258 (0.057)	0.125 (0.070)	-0.136 (0.048)	-0.177 (0.051)	-0.091 (0.005)	-0.076 (0.007)
log income per capita		-0.010 (0.045)		0.029 (0.039)		0.014 (0.004)
log pop. density		-0.016 (0.011)		-0.017 (0.011)		-0.000 (0.001)
Literacy rate		-0.270 (0.139)		-0.124 (0.116)		-0.012 (0.014)
Observations	4,149	4,149	4,149	4,149	4,149	4,149
R^2	0.009	0.012	0.005	0.007	0.068	0.073

Notes: Changes in dependent variables are calculated over the years 1996 and 2006 when the data sources are the agricultural census of 1996 and 2006 (columns 1 to 4), and over the years 2000 and 2010 when the data sources are the population census of 2000 and 2010 (columns 5 and 6). All municipality controls are from the population census of 1991. The unit of observation is the municipality. Robust standard errors reported in parentheses.

coefficient on ΔA_j^{soy} when the outcome is agricultural labor productivity.³⁹ Using our more conservative estimates, namely those that include all municipality controls

³⁹ Due to the different timing of the agricultural and population censuses, agricultural labor productivity changes are measured over the period 1996–2006 while employment share changes are measured over the period

in columns 2 and 6, this ratio is equal to $-0.021/0.134 = -0.155$.⁴⁰ The size of this elasticity implies that a 1 percent increase in agricultural labor productivity corresponds to a 0.155 percentage point decrease in the 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 be explained by technical change in soy, as follows. Consider the average Brazilian municipality, which in the year 2000 had employment shares in agriculture and manufacturing of 38 and 10 percent, respectively. If this municipality experienced an increase in potential soy yields equivalent to a 1 standard deviation from the average increase due to soy technical change, agricultural labor productivity would rise 11 percent, and the agricultural employment share would fall 1.76 percentage points.⁴¹ This estimate corresponds to 24 percent of a standard deviation in the change of the agricultural employment share between 2000 and 2010 (7.4 percentage points, see Table 2).⁴²

In the case of maize, the estimated coefficients reported in columns 3 and 5 indicate that areas with higher increase in potential maize yield experienced a larger increase in labor intensity and the agricultural employment share during the period under study. These findings are consistent with the predictions of the model for the effects of land-augmenting technical change (see Prediction 2). In addition, column 1 shows that areas where maize yields increased relatively more experienced a smaller increase in the value of agricultural output per worker. Our model is too stylized to capture this feature in the data, which is likely driven by the across-crop effect of technical change: reallocation of labor toward maize production reduces the value of output per worker in agriculture. This is because maize production is more labor-intensive than soy production, thus the value of the average product of labor is lower for maize.⁴³

2000–2010. Thus, the elasticity estimates correspond to the effect of four-year lagged agricultural productivity changes on employment shares.

⁴⁰We 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.

⁴¹The first number is computed multiplying 1 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): $0.851 \times 0.134 = 0.114$. The second number is computed multiplying the predicted increase in agricultural labor productivity for 1 standard deviation in ΔA_j^{soy} by the elasticity of agricultural employment share to agricultural labor productivity: $0.114 \times (-0.155) = -0.0176$.

⁴²The reported estimates are representative for the average Brazilian municipality and not the aggregate Brazilian economy, which only had a 17 percent employment share in agriculture in the year 2000. Elasticity estimates that are representative for a municipality that has the same sectoral distribution of employment as the aggregate economy can be obtained by weighting each observation by the aggregate employment share of each municipality. We report such estimates in online Appendix Table A1. These estimates imply that the elasticity of the agricultural employment share to agricultural labor productivity is -0.053 , around one-third of the estimate for the average municipality discussed here.

⁴³A more formal explanation of the effect of labor reallocation toward maize on the value of agricultural output per worker follows. Suppose that there are only two crops, soy and maize, and two production factors, land and labor. In addition, maize production is more labor-intensive than soy. The value of output per worker in agriculture is defined as $\frac{PQ}{L} \equiv \frac{P_{mze}Q_{mze} + P_{soy}Q_{soy}}{L} = \frac{P_{mze}Q_{mze}}{L_{mze}} \frac{L_{mze}}{L} + \frac{P_{soy}Q_{soy}}{L_{soy}} \frac{L_{soy}}{L}$. In this case, a reallocation of labor toward maize production reduces the value of output per worker in agriculture. This is because if soy production is more land-intensive than maize production ($\frac{T_{soy}}{L_{soy}} > \frac{T_{mze}}{L_{mze}}$), the value of the average product of labor is higher

To sum up, the results presented in Table 8 suggest that the introduction of new agricultural technologies in Brazil had a sizable impact on agricultural labor markets. Areas where the potential profitability 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 the predictions of the model for the effects of strongly labor-saving technical change. In the case of maize, areas where the potential profitability of the introduction of a second harvesting season was higher experienced an increase in labor intensity and in the employment share of agriculture. These findings are consistent with the predictions of the model for the effects of land-augmenting technical change.

D. Manufacturing Outcomes

In this section we study the effect of agricultural technical change on manufacturing employment and wages. Table 9 reports the results of estimating equation (12) where the dependent variables are three manufacturing outcomes: the employment share of manufacturing, the level of manufacturing employment, and the average wage in manufacturing.

The estimates indicate that areas where potential soy yields increased relatively more, experienced a larger increase in the manufacturing employment share between 2000 and 2010. 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 report that the absolute level of manufacturing employment increased, not only its employment share. Finally, columns 5 and 6 show that manufacturing wages fell. These estimates are consistent with the empirical predictions of the model: because technical change in soy is strongly labor-saving, it reduces labor demand in agriculture generating a reduction in wages and a reallocation of labor toward the manufacturing sector.

The estimates discussed above 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 IVC: 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. This ratio is equal to $0.021/0.134 = 0.157$ in the estimation including controls for lagged municipality characteristics. This elasticity implies that a 1 percent increase in agricultural labor productivity corresponds to a 0.157 percentage points increase in the 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 1 standard deviation increase in potential soy yield experienced an increase in agricultural labor productivity of 11 percent, and a corresponding 1.79 percentage points increase in

for soy $\left(\frac{P_{soy}Q_{soy}}{L_{soy}} > \frac{P_{mze}Q_{mze}}{L_{mze}} \right)$. To see why this is the case, note that the zero profit conditions for maize and soy ($P_iQ_i = rT_i + wL_i$ for $i = soy, mze$) imply $\frac{P_iQ_i}{L_i} = r \frac{T_i}{L_i} + w$.

TABLE 9—THE EFFECT OF AGRICULTURAL TECHNOLOGICAL CHANGE ON MANUFACTURING
(*Employment Share, Employment, and Wages*)

	Δ Employment share		Δ log employment		Δ log wage	
	(1)	(2)	(3)	(4)	(5)	(6)
ΔA^{soy}	0.023 (0.002)	0.021 (0.002)	0.218 (0.018)	0.186 (0.020)	-0.032 (0.012)	-0.024 (0.012)
ΔA^{maize}	-0.005 (0.001)	-0.004 (0.001)	-0.057 (0.009)	-0.043 (0.009)	0.018 (0.005)	0.014 (0.005)
Share rural population	-0.006 (0.004)	0.011 (0.005)	-0.186 (0.044)	0.051 (0.056)	0.197 (0.026)	-0.014 (0.035)
log income per capita		0.002 (0.003)		0.093 (0.037)		-0.107 (0.026)
log pop. density		0.002 (0.001)		0.020 (0.008)		-0.035 (0.005)
Literacy rate		0.034 (0.010)		0.197 (0.117)		0.093 (0.075)
Observations	4,149	4,149	4,149	4,149	4,149	4,149
R^2	0.063	0.073	0.056	0.068	0.022	0.045

Notes: Changes in dependent variables are calculated over the years 2000 and 2010. All municipality controls are from the population census of 1991. The unit of observation is the municipality. Robust standard errors reported in parentheses.

manufacturing employment share.⁴⁴ This estimate corresponds to 31 percent of a standard deviation in the change of the manufacturing employment share between 2000 and 2010 (5.7 percentage points, see Table 2).

In the case of maize, the estimates reported in columns 1 and 2 of Table 9 indicate that areas where potential maize yields increased relatively more experienced a smaller increases in the 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: because technical change in maize is land-augmenting, it increases labor demand in agriculture, generating an increase in wages and a reallocation of labor away from the manufacturing sector.

E. Services and Other Sectors

In this section we complement our empirical findings with an analysis of the service sector. For this purpose, we reproduce the estimates of the effects of technical change on the agricultural and manufacturing employment shares in Table 10, where we also include estimates for the service sector.⁴⁵ The point estimates of the effect

⁴⁴This number is computed multiplying the predicted increase in agricultural labor productivity for 1 standard deviation in ΔA_j^{soy} by the elasticity of manufacturing employment share to agricultural labor productivity: $0.114 \times 0.157 = 0.0179$.

⁴⁵The services sector includes: construction, commerce, lodging and restaurants, transport, finance, housing services, domestic workers, and other personal services. Other sectors include: public administration, education, health, international organizations, extraction, and public utilities.

of soy technical change on the agriculture and manufacturing employment shares have the same size: they are -0.021 and 0.021 , respectively, both with a standard error of 0.002 . At the same time, the estimates of the effects on the service and other sectors are very small and not statistically different from zero. This implies that local technical change in soy induced a reallocation of labor from agriculture to the local manufacturing sector but not toward local services.

To interpret these findings, we turn to the model with nontraded services where we identified two effects of labor-saving technical change in agriculture: the supply effect and the demand effect. The supply effect is generated by the reduction in the marginal product of labor in the agricultural sector, which reduces agricultural employment. The demand effect is generated by higher income resulting from agricultural productivity growth which leads to increased consumption of services. As a result, the net effect of agricultural technical change on industrialization depends on the relative strength of the supply and demand effects. In addition, the demand effect is driven by the increase in land rents, thus its strength depends on the extent to which landowners consume services in the region where their land is located. Then, the model can explain the absence of an effect of local technical change on employment in the local service sector if the share of land rents that accrue to landlords consuming services in the same municipality where they own land (θ) is small. We use information from the agricultural census about the presence of small family farms in each municipality to show that this is the case. Family farms covered only 15 percent of cultivated area of farms whose main activity is soy production in 2006.⁴⁶ These data can be used as a proxy for θ under the assumption that landlords owning large estates are less likely to reside locally or consume local services. We thus use this information to study whether the effect of soy technical change on the employment share of services was larger in areas with a higher presence of small family farms. We report the results from this analysis in Tables A2 and A3 in the online Appendix. First, we find that in areas characterized by a large presence of family farms, the expansion of the soy area is associated with increases in income per capita and the employment share of services.⁴⁷ Second, we implement our identification strategy that uses potential soy yields as a measure of technical change. We find that municipalities with a higher increase in potential soy yields experienced a larger increase in income per capita. Consistently with the simple correlations in the data, this increase in income per capita was larger in municipalities characterized by a higher presence of family farms. However, it did not lead to an increase in the services employment share.

⁴⁶Brazilian law 11.326 defines family farms as those satisfying all the following conditions: area below four fiscal units; substantial amount of labor force provided by the family; agricultural production as main source of family income; farm management by the family itself. They represented 76 percent of farms whose main activity is soy production but covered only 15 percent of their area. In the case of maize, they represent 88 percent of farms whose main activity is maize production and 55 percent of cultivated area.

⁴⁷These results are shown in columns 3 and 4 of Table A2 in the online Appendix. In particular, the size of the estimates reported in column 3 implies that an expansion in soy area is positively correlated with an increase in income per capita only in municipalities where family farms represent more than 54 percent of soy-producing farms. This condition is satisfied in 48 percent of soy-producing municipalities. Similarly, the size of the estimates reported in column 4 implies that an expansion in soy area is positively correlated with an increase in services employment share only in municipalities where family farms are more than 78 percent of soy-producing farms. This condition is satisfied only in 26 percent of soy-producing municipalities.

TABLE 10—THE EFFECT OF AGRICULTURAL TECHNOLOGICAL CHANGE ON EMPLOYMENT SHARES

	Δ Employment share			
	Agriculture (1)	Manufacturing (2)	Services (3)	Other sectors (4)
ΔA^{soy}	-0.021 (0.002)	0.021 (0.002)	-0.002 (0.002)	0.001 (0.001)
ΔA^{maize}	0.006 (0.001)	-0.004 (0.001)	-0.000 (0.001)	-0.001 (0.001)
Share rural population	-0.076 (0.007)	0.011 (0.005)	0.043 (0.005)	0.023 (0.004)
log income per capita	0.014 (0.004)	0.002 (0.003)	-0.015 (0.003)	-0.001 (0.002)
log pop. density	-0.000 (0.001)	0.002 (0.001)	0.000 (0.001)	-0.002 (0.001)
Literacy rate	-0.012 (0.014)	0.034 (0.010)	-0.009 (0.010)	-0.013 (0.007)
Observations	4,149	4,149	4,149	4,149
R^2	0.073	0.073	0.103	0.045

Notes: Changes in dependent variables are calculated over the years 2000 and 2010. All municipality controls are from the population census of 1991. Services include: construction, commerce, lodging and restaurants, transport, finance, housing services, domestic workers, and other personal services. Other sectors include: public administration, education, health, international organizations, extraction, and public utilities. The unit of observation is the municipality. Robust standard errors reported in parentheses.

Taken together, our empirical findings indicate that local technical change does not significantly affect local employment in services in the average Brazilian municipality. Note, however, that these findings do not imply that agricultural technical change did not have an effect on the demand for services in the aggregate economy. To clarify this point, we extend the model to analyze the simple case where the residence distribution of landowners across the country is identical to that of workers.⁴⁸ In this case the effect of local technical change on the local demand for services is small because the consumption of local absentee landowners is spread across all municipalities in proportion to their workers population. However, if several regions experience technical change at the same time, the aggregate demand effect of technical change might be large. Still, the difference-in-differences empirical strategy can not identify it because the effect has the same value for all municipalities. Thus, a further investigation of the effect of agricultural technical change on the aggregate demand for services is left for future work.

F. Variable Factor Endowments

The model presented in Section II 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

⁴⁸ See Section A.3.5 in the online Appendix for detailed derivations. This case is also equivalent to one where all land income is taxed away and redistributed lump-sum to workers. This case is relevant for Brazil because income taxes are collected by the federal government and partly redistributed across municipalities based on population.

economy: goods can be traded across municipalities but labor markets are local and there is a fixed supply of land. However, the period under study is characterized by significant internal migration flows: 16 percent of the population between 16 and 55 years old had moved to their 2010 municipality of residence during the previous 10 years. In addition, Brazil has vast areas of underutilized land, which were in part converted to agricultural activities during the period under study. Between 1996 and 2006 the land used for cultivation or cattle ranching increased by 7 percent to 154 million hectares in the regions of the North, North-East, and Center-West. Thus, in this section, we investigate the role of migration and the expansion in the agricultural frontier.

Labor.—We first investigate the impact of agricultural technical change on migration flows. The model predicts that municipalities more affected by labor-saving technical change (GE soy) experience a larger contraction in labor demand in the agricultural sector. Because labor is assumed to be immobile across municipalities, all the adjustment to technological change occurs through a reallocation of labor toward the manufacturing sector. However, if workers could relocate to other municipalities, some of this adjustment would occur through out-migration. To test this prediction, we construct net migration rates for every municipality between 2000 and 2010 using data from the population census.⁴⁹ Next, we estimate the baseline specification described by equation (12) using the net migration rate in each municipality as dependent variable. Estimation results are presented in the first column of Table 11. The estimated coefficient on the change in soy potential yields is negative and statistically significant, indicating that municipalities with larger increases in potential soy yields experienced a net outflow of migrants between 2000 and 2010. These estimates can be used to assess the relative importance of the two adjustment mechanisms mentioned above: labor reallocation toward other sectors and out-migration. For this purpose, we can first compute the elasticity of migration flows to changes in agricultural labor productivity due to GE soy adoption: a 1 percent increase in agricultural labor productivity corresponds to a 0.094 percentage points decrease in the migration rate.⁵⁰ This amounts to roughly one-third (0.37) of the reduction in the employment share of the agricultural sector.⁵¹ Finally, 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, as expected.

⁴⁹Net migration rates are defined as the number of net migrants in a municipality divided by its population. A detailed explanation of how net migration rates are constructed is contained in the online Appendix.

⁵⁰We compute this elasticity as in Section IVC: we divide the estimated coefficient on ΔA_j^{soy} when the outcome is the migration rate by the estimated coefficient on ΔA_j^{soy} when the outcome is agricultural labor productivity. When we estimate the specification including controls for municipality characteristics, this ratio is equal to $-0.013/0.134 = -0.094$.

⁵¹To compare the migration rate estimates with the reduction in the employment share of agriculture we need to take into account that the migration rate is computed relative to the overall population aged 16 to 55 years old in 2000, while employment shares are computed relative to workers only. Thus, we multiply the elasticity of migration rate to changes in agricultural labor productivity for the overall population aged 16 to 55 years old in 2000 (-0.094) by the share of active population in the age group 16–55 in 2000 (0.71) and the employment rate for that same age group (0.85). This adjusted elasticity is equal to -0.057 . Then, we divide this number by the estimated elasticity of agricultural employment share to changes in agricultural labor productivity (-0.155), obtaining a ratio of 0.37.

TABLE 11—VARIABLE FACTOR ENDOWMENT

	Migration rate			Δ Agriculture employment share		Δ Manufacturing employment share	
	All (1)	Nonfrontier (2)	Frontier (3)	Nonfrontier (4)	Frontier (5)	Nonfrontier (6)	Frontier (7)
ΔA^{soy}	-0.013 (0.004)	-0.015 (0.005)	-0.012 (0.006)	-0.023 (0.003)	-0.020 (0.004)	0.023 (0.002)	0.019 (0.004)
ΔA^{maize}	0.006 (0.002)	0.007 (0.002)	0.003 (0.003)	0.008 (0.001)	0.003 (0.002)	-0.005 (0.001)	-0.003 (0.002)
Share rural pop.	-0.078 (0.011)	-0.095 (0.014)	-0.035 (0.020)	-0.081 (0.008)	-0.061 (0.012)	0.019 (0.006)	-0.004 (0.009)
log income per capita	0.051 (0.008)	0.050 (0.009)	0.047 (0.013)	0.017 (0.005)	0.008 (0.007)	0.006 (0.004)	-0.003 (0.005)
log pop. density	-0.006 (0.002)	-0.002 (0.003)	-0.009 (0.003)	-0.001 (0.001)	0.001 (0.002)	0.001 (0.001)	0.001 (0.001)
Literacy rate	0.009 (0.023)	0.018 (0.027)	0.079 (0.038)	-0.026 (0.017)	0.032 (0.024)	0.018 (0.012)	0.038 (0.016)
Observations	4,149	2,617	1,532	2,617	1,532	2,617	1,532
R^2	0.104	0.119	0.113	0.080	0.076	0.076	0.066

Notes: Changes in dependent variables are calculated over the years 2000 and 2010. All municipality controls are from the population census of 1991. Municipalities that are part of the agricultural frontier are those that, between 1996 and 2006, experienced an increase in agricultural land used for the cultivation of permanent crops, seasonal crops, and cattle ranching. Municipalities that are part of the agricultural nonfrontier are those that experienced no increase, or a negative change, in used agricultural land between 1996 and 2006. The unit of observation is the municipality. Robust standard errors reported in parentheses.

The findings discussed above suggest that the presence of migration flows across municipalities dampens the effects of technical change on sectoral employment shares, as part of the adjustment occurs through migration flows. In particular, in our model, we can think of out-migration induced by labor-saving technical change as a reduction in the labor endowment, which would result in a reduction in the manufacturing employment share. This is because equilibrium agricultural employment is unaffected by a change in the labor endowment (see equation (4)). In turn, the equilibrium level of employment in manufacturing is determined by the labor market clearing condition, $L_m = L - L_a^*$. Thus, the manufacturing employment share must fall when the labor endowment falls. As a result, the presence of migration dampens the positive effects of soy technical change on the manufacturing employment share. A similar argument implies that the in-migration induced by land-augmenting technical change in maize would increase the manufacturing employment share and dampen the effects of maize technical change.

Land.—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 ecoregion in central Brazil where soils used to be too acidic and nutrient poor. Starting from the 1980s these soils were treated by the Empresa Brasileira de Pesquisa Agropecuária (EMBRAPA), which enabled agricultural activities to expand over these areas. The incorporation of forest or fallow land into agricultural activities can potentially affect our estimates of the effects of technical change. In the model, an expansion in the land endowment would have the same effects as land-augmenting technical change. Thus, differential

increases in the land endowment across regions could account for our finding that areas more affected by technical change in maize experienced an increase in the agricultural employment share or attenuate our findings for the effects of soy technical change.

To assess the extent to which our estimates are affected by expansions in the agricultural frontier we test the predictions of the model in a subsample of municipalities where the land endowment did not increase. In particular, 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 A7 in the online Appendix). Next, we estimate our baseline specification described by equation (12) separately for each 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 only slightly larger (smaller) in absolute value than estimates using the full sample, as shown in columns 4–7 of Table 11. This finding suggests that the expansion of the agricultural frontier does not significantly mitigate our baseline estimates. 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. In contrast, estimates are smaller and not statistically significant in the frontier. These findings suggest that introducing a second harvesting season for maize only had significant effects on labor demand in non-frontier municipalities.

Finally, we study whether migration patterns differ in frontier and non-frontier municipalities. Columns 2 and 3 of Table 11 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.

V. Robustness Checks

A. Additional Controls

A potential concern regarding our estimates is that municipalities that benefit the most from technical change in soy also have higher overall agricultural productivity. Thus, our estimates could be capturing differential structural transformation trends across municipalities that differ in their initial level of agricultural development. To address this concern, we report estimates of equation (12) including controls for three different measures of agricultural development: productivity, wages, and employment share.

Coefficient estimates are reported in Tables A4 and A5 of the online Appendix. The estimated effects of soy technical change on agricultural and manufacturing outcomes are robust to the inclusion of these controls. First, note that the sign of estimated coefficients remains the same and estimates remain significant at 1 percent. In terms of their absolute value, estimated coefficients are stable for the expansion of soy area, output per worker, labor intensity, and manufacturing wages. Estimates for the agricultural and manufacturing employment shares decrease 25 and 40 percent, respectively, when we include the control for agricultural labor productivity. The reason why estimates are affected by the inclusion of this control is that, to some

extent, places with higher initial soy yields benefited more from the new technology. As a result, the control for lagged overall agricultural productivity captures part of the variation we are interested in. Thus, we interpret our estimates of the effects of soy technical change conditional on the initial level of agricultural productivity as indicative that at least 60 percent of our estimated effects of technical change on sectoral employment shares are not driven by differential structural transformation trends across municipalities that differ in the initial level of agricultural productivity.⁵²

We obtain similar findings in the case of maize. Estimated coefficients are robust to including these additional controls. Estimates of the effect of maize technical change on agricultural labor intensity and manufacturing wages are stable and significant at 1 percent. In the case of the agricultural and manufacturing employment shares, estimates fall by 25 and 40 percent, respectively, when we control for lagged labor productivity in agriculture.

B. Preexisting Trends

In this section we show that our results are robust to controlling for preexisting trends. This exercise addresses the following concern: if municipalities that are better suited for adopting GE soy were already experiencing faster structural transformation before the legalization of this technology in Brazil, 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 preexisting 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:

$$(13) \quad \Delta y_{jt} = \delta_t + \beta_0^{soy} \Delta A_j^{soy} + \beta_1^{soy} \Delta A_j^{soy} After_t + \beta_0^{maize} \Delta A_j^{maize} + \beta_1^{maize} \Delta A_j^{maize} After_t + \mathbf{X}'_{jt-1} \boldsymbol{\omega} + \Delta \varepsilon_{jt},$$

where the outcome of interest, Δy_{jt} is the decadal change in outcome variables between the start of a period (year $t - 1$) and the end (year t). Each period spans a decade: 1991 to 2000 and 2000 to 2010. δ_t are time dummies for each decade and $After_t$ is a dummy equal to 1 if $t = 2010$. Thus, β_0^{soy} 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. In contrast, β_1^{soy} captures the differential effect of soy technical change after the introduction of GE soy seeds. Similarly, the coefficient β_1^{maize} captures the differential effect of maize technical change in the period 2000–2010. Finally, \mathbf{X}_{jt-1} is a vector containing a set of ten-year-lagged

⁵²Note that all coefficient estimates are stable when we only include the control for the lagged agricultural employment share, except for the estimated effect of technical change on employment shares themselves which tend to fall. Still, the estimated effect of technical change on the manufacturing employment share only falls from 0.021 to 0.014 and remains statistically significant at 1 percent. These results imply that our estimated coefficients are not capturing delayed responses to the trade liberalization that occurred at the beginning of the previous decade in areas with different initial agricultural specialization, studied by Dix-Carneiro and Kovak (2014).

municipality characteristics including the share of rural population, average income per capita, population density, and literacy rate.⁵³

Results for manufacturing employment are reported in column 1 of Table A6 of the online Appendix. Our estimate of β_0^{soy} , which 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, is very small and not statistically different from zero. This finding indicates that there are no pretrends in manufacturing employment. In addition, our estimate of β_1^{soy} , which estimates the differential effect of soy technical change on manufacturing employment after the introduction of GE soy seeds, is positive and precisely estimated. Similarly, in the case of maize, we do not find preexisting trends in manufacturing employment. Note that we perform this test for the level of manufacturing employment but not for the manufacturing and agricultural employment shares. This is because there were important changes in the definition of employment after the 1991 census, thus employment shares can not be measured in a consistent way across the 1991 and 2000 censuses.⁵⁴

Column 2 of Table A6 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 preexisting 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 9.

Finally, we check for preexisting trends in migration. A potential concern is that areas that 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 a higher incentive to adopt a labor-saving technology to cope with labor scarcity. Column 3 of Table A6 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 preexisting trends in migration for areas that have a higher increase in potential soy yields. Similarly, in the case of maize, we do not find preexisting trends in migration.⁵⁵

⁵³The 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.

⁵⁴Between the 1991 and 2000 population censuses 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 who 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 percent of agricultural workers reported zero income, against 34 percent in 2000 and 35 percent 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 who were not employed during the reference week were counted in the 1991 census but not 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 preexisting 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: there are very few zero-income workers in manufacturing (0.5 percent, 1.9 percent, and 1 percent of manufacturing workers declare zero income in 1991, 2000, and 2010, respectively); and manufacturing is less seasonal than other activities.

⁵⁵These results suggest that the migration flows generated by the expansion of the Brazilian road network in the years 1960–2000 that are studied by Morten and Oliveira (2014) are unlikely to be confounding our results.

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

C. Larger Unit of Observation: Microregions

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, especially if manufacturing activities 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: microregions. These regions are groups of territorially contiguous municipalities created, for statistical purposes, by the Brazilian Statistical Institute (IBGE). Table A7 reports the results of estimating equation (12) using microregions as a unit of observation. The outcome variables are the same as in Table 9: change in manufacturing employment share, change in manufacturing employment (in logs), and change in average manufacturing wage (in logs). The estimates are consistent and similar in magnitude to those reported in Table 9, both for soy and maize.

D. Input-Output Linkages

Our theoretical model predicts that agricultural technical change can have an effect on manufacturing employment through labor market forces only. 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 the area farmed with soy in a given municipality might 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 face high transport costs, there might be an incentive for them to locate in the same municipality in which agricultural production takes place. Therefore, the effect of agricultural technical change on manufacturing that we show in Table 9 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.

In order to identify input-output linkages in the data, we proceed as follows. We use the 2005 Brazilian input-output matrix (IBGE 2008) to identify manufacturing sectors that are providing inputs, or receiving outputs, from the soy and maize

sectors. On the input side, soy and maize are used as intermediate goods in only one manufacturing sector: the food and beverage sector, which in 2005 purchased around one-half of the total Brazilian production of both crops. On the output side the matrix is less detailed, thus we use information on goods purchased by agricultural and breeding farms in general. One-half of the inputs purchased by these farms are supplied by manufacturing sectors and four commodities account for 84 percent of the total value of inputs purchased: inorganic chemicals, fertilizers, diesel oil, and maize oil. These commodities are produced by the chemical industry, the oil refining industry, and the food and beverage industry. We use this information to construct measures of employment and wages in manufacturing that exclude those industries which are providing inputs, or receiving outputs, from the soy and maize sectors.

Table A8 reports estimates of our baseline specification described by equation (12) using as outcome variables measures of manufacturing employment and wages that exclude workers employed in sectors directly linked to soy and maize. Estimates of the effect of soy technical change on the manufacturing employment share and level are positive, precisely estimated and 38 to 10 percent smaller than our baseline estimates displayed in Table 9.⁵⁶ In turn, the effect of technical change in soy on manufacturing wages decreases substantially, and is not precisely estimated. In the case of maize, estimated coefficients are essentially unaffected by excluding workers in downstream and upstream 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 is not precisely estimated. Taken together, the results presented in this section imply that at least 62 percent 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 38 percent of the total estimated effect, which is an interesting avenue for further work.

E. 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 time fixed effects in equation (9). 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 A8 and A9 display the evolution of international prices of soy and maize, expressed in 2000 US\$. These figures show how the international prices of both

⁵⁶In our specification with all initial municipality controls, the point estimate on ΔA^{soy} when the outcome is manufacturing employment share goes from 0.021 to 0.013. 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 0.186 to 0.167. In this case, the two coefficients are not statistically different.

commodities have been in an upward trend starting from year 2007. This pattern is unlikely to affect our estimates when we use data from the last two agricultural censuses, 1996 and 2006. In particular, note that the international price for both soy and maize was lower in 2006 than in 1996. However, when we use data from the last two population censuses, which took place in 2000 and 2010, the end of period year is characterized by high international soy and maize prices with respect to the initial year. To address this concern, we assess the robustness of our findings for the manufacturing sector to controlling for changes in commodity prices.

The data from the population censuses do 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.⁵⁷ We estimate an equation of the following form:

$$(14) \quad y_{jt} = \delta_j + \delta_t + \beta^{soy} A_{jt}^{soy} + \beta^{maize} A_{jt}^{maize} + \sum_z \lambda^z P_t^z A_{j0}^z + t \mathbf{X}'_{j,1991} \boldsymbol{\omega} + \varepsilon_{jt},$$

where y_{jt} is total employment or average wage in a given municipality; A_{jt}^{soy} 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_{jt}^{maize}). 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 control for differential trends across municipalities with different initial levels of development by adding an interaction of the vector of initial municipality characteristics ($\mathbf{X}_{j,1991}$) and a time trend (t). In all specifications we control for both municipality and year fixed effects (δ_j and δ_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 A9 in the online Appendix): 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 9, 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.

⁵⁷The average wage is defined as the aggregate wage bill (in real terms) divided by the total number of workers employed in a municipality.

F. Spatial Correlation

The maps we present in Figures A3 to A6 suggest that the potential yields of soy and maize are correlated across space. Therefore, in this section we show that the estimates of the effect of agricultural technical change reported in Tables 8, 9, and 10 remain statistically significant when we correct standard errors to account for spatial correlation. First, we allow the residuals to be correlated within geographical areas larger than a single municipality. For this purpose, we compute standard errors clustered at two larger levels of aggregation: microregions and mesoregions.⁵⁸ Second, we calculate standard errors that correct for spatial dependence as suggested by Conley (1999). This procedure allows the correlation of residuals across municipalities to be a decaying function of distance until a fixed threshold, as explained below.

Tables A10, A11, and A12 in the online Appendix report our results. The first row below the coefficients reports baseline robust standard errors for comparison. The following two rows report standard errors clustered at micro- and mesoregion levels. Finally, the last three rows report the Conley standard errors calculated assuming errors are correlated within 50, 100, and 200 km. We report the significance level alongside each of these estimated standard errors. In the case of soy technical change, the tables show that although standard errors tend to increase after accounting for spatial correlation, most coefficient estimates remain statistically significant at 1 percent. In the case of maize, all estimates remain statistically significant except for the manufacturing employment share when clustering at the mesoregion level or considering the largest distance cutoff of 200 km.

G. Alternative Definition of Technical Change

In this section we discuss in more detail the measure of technical change obtained from the FAO-GAEZ dataset. We use the change in potential soy yields when switching from the low to the high technology as a source of exogenous variation in agricultural productivity. In particular, we use it as an instrument for agricultural labor productivity which permits to obtain the elasticity of sectoral employment shares to changes in labor productivity induced by technical change in soy. Ideally, we would like the measure of change in potential yields to capture only the effect of adopting GE soy, and not other changes in the production technology of soy. Our measure of technical change deviates from this ideal because the FAO-GAEZ dataset characterizes agricultural technologies as bundles of inputs, including seed quality, level of mechanization, and use of chemicals. Because all of these inputs change when switching from the low to the high technology, a potential concern is that our measure of technical change might capture other changes in the production technology of soy. We address this concern in two ways. First, in Table 6 we show that our measure of technical change in soy predicts the expansion in the area planted with GE soy but not the expansion in the area planted with non-GE soy. Second, we test the robustness of our results to an alternative definition of technical

⁵⁸ Both microregions and mesoregions are statistical divisions of Brazil proposed by the IBGE to facilitate the collection of data. There are 558 microregions and 137 mesoregions.

change that uses potential yields under an intermediate technology to capture the level of agricultural technology before the introduction of GE seeds. We discuss the results obtained below.

The FAO-GAEZ dataset characterizes the intermediate technology as using improved varieties of seeds, partial mechanization, and some use of chemicals. This technological level lies somewhere in between traditional and technologically advanced farming. We estimate equation (12) for the set of agricultural and manufacturing outcomes of interest, using the differences in potential yields in soy and maize between the high and the intermediate level of technological inputs to measure ΔA_j^{soy} and ΔA_j^{maize} . Table A13 presents the resulting estimates. A comparison with Tables 8, 9, and 10 shows that our main results are robust to this alternative definition of technical change in agriculture in the sense that point estimates and standard errors have a similar size. We can use the estimated coefficients under this alternative specification to compute the elasticity of agricultural and manufacturing employment shares to changes in agricultural labor productivity due to GE soy adoption in the same way as we do in Sections IVC and IVD. The elasticities obtained are 26 percent smaller in the case of the agricultural employment share and 45 percent smaller in the case of the manufacturing employment share.⁵⁹

In sum, our estimates of the effect of soy technical change on employment shares are smaller when we use the intermediate technology as a description of the situation before GE soy was adopted. Still, we prefer to use the difference between high and low level of inputs in our baseline specification for two reasons. First, it is a more precise measure of technical change in agriculture. This is because the high and low level of technical inputs are clearly defined, while intermediate inputs has a loose definition that could span different levels of agricultural technology. As a result, this measure might miss part of the variation that we are trying to capture. For example, improved seed varieties and herbicides which are described as part of the bundle of intermediate inputs can capture part of the effect of adopting GE seeds. Second, the main potential concern with using the difference between the high and the low technology as a measure of technical change is that it might capture changes in technology other than the adoption of GE soy, like mechanization. However, the finding that this measure of technical change is not positively correlated with the expansion of non-GE soy suggests that this concern is not important in practice.

VI. Final Remarks

This paper provides direct empirical evidence on the effects of agricultural productivity on structural transformation. We isolate these effects by studying the introduction of genetically engineered soy in Brazil. This technology allows farmers to employ fewer workers per unit of land to yield the same output, increasing labor productivity in agriculture. After its legalization in 2003, genetically engineered soy experienced a rapid and widespread adoption in Brazil. We exploit the

⁵⁹The elasticity of the agricultural employment share to agricultural labor productivity is -0.115 . As for manufacturing, we obtain an elasticity of employment share to agricultural labor productivity of 0.086 .

differential impact of this new technology on potential yields across geographical areas to estimate the causal effect of agricultural technical change on sectoral employment shares.

Our findings contribute to the debate on the effects of agricultural productivity on industrialization in open economies. We argue that these effects depend crucially on the factor-bias of technical change. We provide evidence that when technical change in agriculture is strongly labor-saving, as in the case of genetically engineered soy, it can foster industrialization. When, instead, technical change is land-augmenting, as in the case of the introduction of a second harvesting season in maize, agricultural productivity growth can retard industrialization.

REFERENCES

- Acemoglu, Daron.** 2010. "When Does Labor Scarcity Encourage Innovation?" *Journal of Political Economy* 118 (6): 1037–78.
- Acemoglu, Daron, and Veronica Guerrieri.** 2008. "Capital Deepening and Nonbalanced Economic Growth." *Journal of Political Economy* 116 (3): 467–98.
- Baumol, William J.** 1967. "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review* 57 (3): 415–26.
- Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli.** 2016. "Agricultural Productivity and Structural Transformation: Evidence from Brazil: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer.20131061>.
- Caselli, Francesco, and Wilbur J. Coleman II.** 2001. "The US Structural Transformation and Regional Convergence: A Reinterpretation." *Journal of Political Economy* 109 (3): 584–616.
- Clark, Colin.** 1940. *The Conditions of Economic Progress*. London: MacMillan.
- Companhia Nacional de Abastecimento (CONAB).** 2003–2012. "Levantamento De Avaliação Da Safra." Brasília.
- Conley, Timothy G.** 1999. "GMM Estimation with Cross Sectional Dependence." *Journal of Econometrics* 92 (1): 1–45.
- Corden, W. Max, and J. Peter Neary.** 1982. "Booming Sector and De-Industrialisation in a Small Open Economy." *Economic Journal* 92 (368): 825–48.
- Dix-Carneiro, Rafael, and Brian K. Kovak.** 2014. "Trade Reform and Regional Dynamics: Evidence from 25 Years of Brazilian Matched Employer-Employee Data." Unpublished.
- Duffy, Michael, and Darnell Smith.** 2001. "Estimated Costs of Crop Production in Iowa." Iowa State University Extension Service FM1712.
- Empresa Brasileira de Pesquisa Agropecuária (EMBRAPA).** 2006. "Cultivo do Milho." EMBRAPA, Sistemas de Produção.
- Fernandez-Cornejo, Jorge, and Margriet Caswell.** 2006. "The First Decade of Genetically Engineered Crops in the United States." *United States Department of Agriculture Economic Information Bulletin* 11.
- Fernandez-Cornejo, Jorge, Cassandra Klotz-Ingram, and Sharon Jans.** 2002. "Farm-Level Effects of Adopting Herbicide-Tolerant Soybeans in the USA." *Journal of Agricultural and Applied Economics* 34 (1): 149–63.
- Field, Alexander J.** 1978. "Sectoral Shifts in Antebellum Massachusetts: A Reconsideration." *Exploration in Economic History* 15 (2): 146–71.
- Findlay, Ronald, and Harry Grubert.** 1959. "Factor Intensities, Technological Progress, and the Terms of Trade." *Oxford Economic Papers* 11 (1): 111–21.
- Food and Agriculture Organization (FAO).** 2015. "Global Agro-Ecological Zones." Rome: Food and Agriculture Organization.
- Foster, Andrew D., and Mark R. Rosenzweig.** 2004. "Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India, 1970–2000." *Economic Development and Cultural Change* 52 (3): 509–42.
- Foster, Andrew D., and Mark R. Rosenzweig.** 2008. "Economic Development and the Decline of Agricultural Employment." *Handbook of Development Economics* 4: 3051–83.
- Fujita, Masahisa.** 1989. *Urban Economic Theory: Land Use and City Size*. Cambridge, UK: Cambridge University Press.

- Gollin, Douglas, David Lagakos, and Michael E. Waugh.** 2014. "The Agricultural Productivity Gap." *Quarterly Journal of Economics* 129 (2): 939–93.
- Gollin, Douglas, Stephen Parente, and Richard Rogerson.** 2002. "The Role of Agriculture in Development." *American Economic Review* 92 (2): 160–64.
- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi.** 2013a. "Growth and Structural Transformation." National Bureau of Economic Research Working Paper 18996.
- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi.** 2013b. "Two Perspectives on Preferences and Structural Transformation." *American Economic Review* 103 (7): 2752–89.
- Hornbeck, Richard, and Pinar Keskin.** 2015. "Does Agriculture Generate Local Economic Spillovers? Short-Run and Long-Run Evidence from the Ogallala Aquifer." *American Economic Journal: Economic Policy* 7 (2): 192–213.
- Hornbeck, Richard, and Suresh Naidu.** 2014. "When the Levee Breaks: Black Migration and Economic Development in the American South." *American Economic Review* 104 (3): 963–90.
- Huggins, David R., and John P. Reganold.** 2008. "No-Till: The Quiet Revolution?" *Scientific American* 299: 70–77.
- Instituto Brasileiro de Geografia e Estatística (IBGE).** 1996. Censo Agropecuário 1995–1996. Rio de Janeiro: IBGE.
- Instituto Brasileiro de Geografia e Estatística (IBGE).** 2003. "Metodologia do Censo Demográfico." *Série Relatórios Metodológicos* 25.
- Instituto Brasileiro de Geografia e Estatística (IBGE).** 2004. "Pesquisa Industrial Anual: Empresa." *Série Relatórios Metodológicos* 26.
- Instituto Brasileiro de Geografia e Estatística (IBGE).** 2006. Censo Agropecuário 2006. Rio de Janeiro: IBGE.
- Instituto Brasileiro de Geografia e Estatística (IBGE).** 2008. Matriz de Insumo-Produto, Vol. 23. Rio de Janeiro: IBGE.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie.** 2001. "Beyond Balanced Growth." *Review of Economic Studies* 68 (4): 869–82.
- Krugman, Paul.** 1987. "The Narrow Moving Band, the Dutch Disease, and the Competitive Consequences of Mrs. Thatcher." *Journal of Development Economics* 27 (1–2): 41–55.
- Kuznets, Simon.** 1957. "Quantitative Aspects of the Economic Growth of Nations: II: Industrial Distribution of National Product and Labor Force." *Economic Development and Cultural Change* 5 (4): 1–111.
- Lagakos, David, and Michael E. Waugh.** 2013. "Selection, Agriculture, and Cross-Country Productivity Differences." *American Economic Review* 103 (2): 948–80.
- Lewis, W. Arthur.** 1954. "Economic Development with Unlimited Supplies of Labour." *Manchester School* 22 (2): 139–91.
- Lucas, Robert E.** 1988. "On the Mechanics of Economic Development." *Journal of Monetary Economics* 22 (1): 3–42.
- Matsuyama, Kiminori.** 1992. "Agricultural Productivity, Comparative Advantage, and Economic Growth." *Journal of Economic Theory* 58 (2): 317–34.
- Michaels, Guy, Ferdinand Rauch, and Stephen J. Redding.** 2012. "Urbanization and Structural Transformation." *Quarterly Journal of Economics* 127 (2): 535–86.
- Mokyr, Joel.** 1976. *Industrialization in the Low Countries, 1795–1850*. New Haven, CT: Yale University Press.
- Morten, Melanie, and Jaqueline Oliveira.** 2014. "Migration, Roads and Labor Market Integration: Evidence from a Planned Capital City." Unpublished.
- Murphy, Kevin M., Andrei Shleifer, and Robert Vishny.** 1989. "Income Distribution, Market Size, and Industrialization." *Quarterly Journal of Economics* 104 (3): 537–64.
- Neary, J. Peter.** 1981. "On the Short-Run Effects of Technological Progress." *Oxford Economic Papers* 33 (2): 224–33.
- Ngai, L. Rachel, and Christopher A. Pissarides.** 2007. "Structural Change in a Multisector Model of Growth." *American Economic Review* 97 (1): 429–43.
- Nunn, Nathan, and Nancy Qian.** 2011. "The Potato's Contribution to Population and Urbanization: Evidence from A Historical Experiment." *Quarterly Journal of Economics* 126 (2): 593–650.
- Nurkse, Ragnar.** 1953. *Problems of Capital Formation in Underdeveloped Countries*. Oxford: Basil Blackwell.
- Rosenstein-Rodan, Paul N.** 1943. "Problems of Industrialisation of Eastern and South-Eastern Europe." *Economic Journal* 53 (210): 202–11.
- Rostow, W. W.** 1960. *The Stages of Economic Growth: A Non Communist Manifesto*. London: Cambridge University Press.
- Schultz, Theodore W.** 1953. *The Economic Organization of Agriculture*. New York: McGraw-Hill.

- Syrquin, Moshe.** 1988. "Patterns of Structural Change." In *Handbook of Development Economics*, Vol. 1, edited by Hollis Chenery and T. N. Srinivasan, 203–73. Amsterdam: Elsevier.
- United States Department of Agriculture (USDA).** 2001. "Agriculture in Brazil and Argentina: Developments and Prospects for Major Field Crops." United States Department of Agriculture Report WRS-01-3.
- United States Department of Agriculture (USDA).** 2012. "Brazil: Agricultural Biotechnology Annual." Washington, DC: United States Department of Agriculture, Economic Research Service.
- Wright, Gavin.** 1979. "Cheap Labor and Southern Textiles before 1880." *Journal of Economic History* 39 (3): 655–80.

Copyright of American Economic Review is the property of American Economic Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.