Essays in Development and Applied Economics

Bruno Caprettini

TESI DOCTORAL UPF / ANY 2015

DIRECTORS DE LA TESI

Prof. Antonio Ciccone (University of Mannheim and UPF)
Prof. Paula Bustos (Princeton University, CREI and UPF)

Universitat Pompeu Fabra Barcelona

Acknowledgements

My first thanks are for my two advisers Antonio Ciccone and Paula Bustos, without whose encouragement and guidance this thesis would not exist. Antonio guided me with patience and dedication through all the stages of the Ph.D., and he tried hard to teach me to be my harshest critic. His brilliant insights and his ability to get to the bottom of every issue have been invaluable not just to write this thesis but to shape the kind of researcher I am now.

Paula believed in me from the beginning, and always encouraged me to set the highest standard for myself. I owe her a lot and I learnt enormously by working with her. She also put me in contact with the best economists in the world and I am extremely grateful to her for this.

Special thanks to Joachim Voth and Alessandro Tarozzi: not only they gave me extremely valuable advices, but they were present in some of the most critical moments of the past 6 years. Several UPF and CREI professors took time to comment my ideas and to read my work. Thanks to Fernando Broner, Gino Gancia, Stephan Litschig, Giacomo Ponzetto, Luigi Pascali, Vasco Carvalho, Jaume Ventura, and Alessandra Bonfiglioli.

I wrote part of the thesis while I was a visiting student at the department of economics at Princeton University. Thanks to Gene Grossman, Stephen Redding and everybody in the International Trade Section for their very kind hospitality. In Princeton I also received very valuable advices from Stephen Redding, Esteban Rossi-Hansberg, Ezra Oberfield, Jan De Loecker, Eduardo Morales, Gene Grossman, Oleg Itskhoki, Harvey Rosen and Jean Grossman: I thank them all.

I received constructive comments in many seminars and conferences that I have attended. I thank all the participants of the IZA Summer School, ESSIM, BSF, the PEDL conference, the SSDEV in Ascea, as well as participants to the seminars in Princeton, IPEA, World Bank, IADB, CREI and UPF. Special thanks go to Jeffrey Wooldridge, Chris Woodruff and Kei Mu-Yi, who at some of these conferences have been incredibly generous with their time, advices and encouragement.

During my years as a Ph.D., I received financial support from Fernando Broner's ERC grant, Gino and Paula's Ramon Areces grant and from Paula's PEDL grant. I also received direct funding from the Private Enterprise Development in Low-Income Countries Project by the CEPR and UK Department for International Development. I am very grateful to all of them for their support. Some of this money went to fund research travel to Brazil, where I received the precious support of Carlos Lessa and Luis Pinto to access and work with Brazilian data: I thank them too.

Finally, many thanks to all the administrative staff at Pompeu and CREI. They say that an institution is only as good as its secretaries, and Marta Araque, Mariona Novoa, Laura Agustì and Anna Ricard are definitely one of the reasons why UPF and CREI are among the best economic departments in the world!

Of all the people outside of academia that I want to thank, my family comes first. My mum and my dad gave me unconditional support from the beginning, even though they are still so confused about what exactly I am doing at university after all these years that they never say that I "work": at most I am "studying" — as if I was still in school! Thank you also to my sister Valeria, who always manages to cheer me up.

A very big thank you to all the friends I met in Barcelona. First of all to Ciccio, Al and Vicky: it would have been worth to spend the last 6 years in Barcelona only to meet this bunch, and I am going to miss them enormously. Thank you to Cami and Peter for making me feel at home in Barcelona and for being there whenever I needed a break. Thank you to Jac for having been a guide and a friend from the beginning, and thank you to Matteo for sharing his experiences and his advices with me.

A huge thank you to all the other friends at Pompeu: to Andrea, Tanya, Uri, Kiz, Andrei, Gene, Felipe, Tom, Jag, Miguel, Jörg, Miguelito, Ricardo, Mapis, Cate and Emma: it has been great fun to be with them these years, and they have always been a source of support and understanding. Thanks to John who welcomed me in the U.S., and thanks to Marco, Francesco, Franco and Kari: I would have never managed to finish the Master without them. Thanks to everybody else at Pompeu who I may have forgot and thanks to Michael who sadly left us too soon.

Thanks also to the friends of a lifetime: Ado, Pepo, Gio, Jack, Michele, Tommaso, Giulia, Livia and the whole "intermet appreciation society". Even though it is hard to meet now that we live far away, they are always there when it matters the most.

Finally, the biggest thank you is for Zu. She has always been by my side, encouraging me and trying to remind me what are the most important things in life. I hope I won't forget her lessons now that we're going to live apart.

Abstract

The thesis is a collection of three self-contained essays. In the first chapter, I study the effects of the adoption of new agricultural technologies on structural transformation. I exploit the introduction of genetically engineered soybean seeds in Brazil to show that if the new agricultural technologies are strongly labor saving they promote industrial growth in even in an open economy. In the second chapter, I study the welfare gains of a tax reform that eliminated distortive turnover taxes in Brazil. I find that taking into account the sequential nature of production affects significantly the gains of the reform, that increase with the number of times firms exchange intermediate inputs. In the third chapter, I exploit the random assignment of teams to groups in the soccer's UEFA Champions League to provide causal evidence that hosting a sports team boosts the visibility of a city among tourists. I find that being drawn into the same group leads to between 5 and 8 percent more passengers flying for the three months following the end of the group stage.

Resum

La tesi és una col·lecció de tres treballs de recerca independents. En el primer treball analitzo els efectes de noves tecnologies agrícoles sobre el canvi estructural d'una economia. Utilitzo la introducció de llavors de soja genèticament modificades a Brasil per ensenyar que quan les noves tecnologies agrícoles són "*strongly labor saving*", promouen el creixement industrial també en economies obertes. En el segon treball, estudio una reforma que va eliminar una taxa distorsionadora a Brasil. Trobo que el fet de considerar processos productius seqüencials afecta significativament els beneficis de la reforma, i demostro que aquests beneficis augmenten quan les empreses intercanvien bens intermedis més sovint. En el tercer treball, utilitzo l'assignació casual dels equips de futbol als grups de la *UEFA Champions League* per ensenyar que tenir un equip esportiu a la pròpia ciutat estimula la visibilitat d'aquesta entre els turistes. Evidencio que les connexions aèries entre dues ciutats amb equips que hagin quedat en el mateix grup, porten entre un 5 i un 8 per cent més de passatgers durant els següents tres mesos després de la fi de la fase de grups de la *Champions League*.

Preface

The thesis is a collection of three self-contained essays. The thesis has no unifying theme, and it reflects the evolution of my interests and research style over the course of my doctoral studies. In all my projects I have strived to work with data and with theory to find credible answers to relevant economic questions, and in the two papers that make up the first and the second chapter of the thesis, I focus on different aspects of economic development.

In the first chapter of the thesis, co-authored with Paula Bustos 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

In the second chapter of the thesis, co-authored with Antonio Ciccone, we exploit a Brazilian tax reform to study the productivity losses caused by taxes on turnover, a type of tax that distorts transactions between firms and that is common in developing countries. We build a model in which many sectors exchange inputs sequentially before delivering final goods, and show that distortions to trade between firms are amplified by the "number of stages of production". We calibrate the model to the Brazilian economy and compute the productivity gains of the reform. We find that considering production processes that have 11 or more stages of production leads to estimated productivity gains of removing the turnover tax more than 4 times larger than considering production processes where firms exchange inputs only once. When production requires an infinite number of stages the gains of the reform are maximized and they are equal to around 0.05% of national income.

In the third chapter of the thesis, I ask whether hosting a sports team boost the visibility of a city among tourists. I answer to this question by looking at the effect of playing soccer's UEFA Champions League on air travel. I compare routes across cities that had their teams randomly drawn into the same group in the first phase of the competition to routes across cities hosting teams randomly allocated to different groups. The average effect of being drawn into the same group is between 5 and 8 percent more arrivals for the three months following the group stage, a period which coincides with a break in the competition.

Contents

List of Figures

List of Tables

Chapter 1

Agricultural productivity and structural transformation. Evidence from Brazil. *

1.1 Introduction

The early development literature documented that the growth path of most advanced economies was accompanied by a process of structural transformation. As economies develop, the share of agriculture in employment falls and workers migrate to cities to find employment in the industrial and service sectors (Clark (1940), Kuznets (1957)). These findings suggest that isolating the forces that can give rise to structural transformation is key to our understanding of the development process. In particular, scholars have argued that increases in agricultural productivity are an essential condition for economic development, based on the experience of England during the industrial revolution.[2] Classical models of structural transformation formalize their ideas by showing how productivity growth in agriculture can release labor or generate demand for manufacturing goods.[3] However, several scholars noted that the positive effects of agricultural productivity on industrialization occur only in closed economies, while in open

*Joint with Paula Bustos and Jacopo Ponticelli.

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

[3]See Baumol (1967), Murphy et al. (1989), Kongsamut et al. (2001), Gollin et al. (2002), Ngai and Pissarides (2007).

economies a comparative advantage in agriculture can slow down industrial growth.[4] Despite the richness of the theoretical literature, there is scarce direct empirical evidence testing the mechanisms proposed by these models.[5]

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 to grow two crops a year, effectively increasing the land endowment. Thus, it can be characterized as land-augmenting technical change. The simultaneous expansion of these two crops allows 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 towards agriculture, as in classical open economy models such as 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 towards manufacturing. In sum, the model predicts that the effects of agricultural productivity on structural transformation in open economies depend on the factor-bias of technical change.

In a first analysis of the data we find that regions where the area cultivated with soy expanded experienced an increase in agricultural output per worker, a reduction in labor intensity in agriculture and an expansion in industrial employment. These correlations are consistent with the theoretical prediction that the adoption of labor-augmenting agricultural technologies reduces labor demand in the agricultural sector and induces the

[4]See Mokyr (1976), Field (1978), Wright (1979),Corden and Neary (1982), Krugman (1987), and Matsuyama (1992).

[5]Empirical studies of structural transformation include Foster and Rosenzweig (2004, 2008), Nunn and Qian (2011), Michaels et al. (2012), Hornbeck and Keskin (2015). We discuss this literature in more detail below.

reallocation of workers towards the industrial sector. However, causality could run in the opposite direction. For example: an increase in productivity in the industrial sector could increase 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 invented in the U.S. in 1996, and legalized in Brazil in 2003, we use this last date as our source of variation across time. Second, as the new technology had a differential impact on yields depending on geographical and weather characteristics, we use differences in soil suitability across regions as our source of cross-sectional variation. Similarly, in the case of maize, we exploit the timing of expansion of second-harvest maize and cross-regional differences in soil suitability.

We obtain an exogenous measure of technological change in agriculture by using estimates of potential soil yields across geographical areas of Brazil from the FAO-GAEZ database. These yields are calculated by incorporating local soil and weather characteristics into a model that predicts the maximum attainable yields for each crop in a given area. Potential yields are a source of exogenous variation in agricultural productivity because they are a function of weather and soil characteristics, not of actual yields in Brazil. In addition, the database reports potential yields under traditional and new agricultural technologies. Thus, we exploit the predicted differential impact of the new technology on yields across geographical areas in Brazil as our source of cross-sectional variation in agricultural productivity. Note that this empirical strategy relies on the assumption that although goods can move across geographical areas of Brazil, labor markets are local due to limited labor mobility. This research design allows us to investigate whether exogenous shocks to local agricultural productivity lead to changes in the size of the local industrial sector. We use municipalities as our geographical unit of observation, 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 have a higher effect on potential yields of soy did experience a higher increase in the area planted with GE soy. In addition, these regions experienced increases in the value of agricultural output per worker and reductions in labor intensity measured as employment per hectare. These regions experienced faster employment growth and wage reductions in the industrial sector. Interestingly, the effects of technology adoption are different for maize. Regions where the FAO potential maize yields are predicted to increase the most when switching from the traditional to the new technology did indeed experience a higher increase in the area planted with maize and in the value of agricultural output. However, they also experienced increases in labor intensity, reductions in industrial employment and increases in wages.

The 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.

Our estimates can be used to quantify the effect of factor-biased agricultural technical change on structural transformation. In particular, we estimate the elasticity of sectoral employment shares to changes in agricultural productivity induced by soy technical change: 1 percent increase in agricultural labor productivity leads to 0.16 percentage points increase in the manufacturing employment share. 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-augmenting 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: the average increase in the the manufacturing employment share was 1.4 percentage points but the standard deviation was 5.7 percentage points. Our estimates imply that labor-augmenting technical change in soy can explain 31 percent of these observed differences in the growth of the manufacturing employment share across Brazilian municipalities.

Our estimates of the effect of technical change on sectoral employment shares imply that labor reallocated mostly from agriculture towards manufacturing. This is because estimated effects on the agricultural and manufacturing employment shares have a similar magnitude, while the estimate of its effect on the employment share of services is small and not statistically different from zero. The service sector is expected to behave differently from manufacturing because it is non traded, thus should respond to demand forces. In our setting, labor-augmenting technical change has a positive effect on the income of land-owners. As a result, the response of the demand for services depends on the propensity of land-owners to consume local services. Our ongoing empirical work suggests that the absence of an average effect on local services is related to the fact that in some areas of Brazil land owners do not reside locally. A further investigation of the effect of agricultural technical change on non-traded sectors is left for future work.

Finally, let us note that we perform the following robustness checks. First, we show that our estimates are stable when we allow municipalities with different initial levels of development to be on differential structural transformation trends. Second, we relax the assumption of constant factor endowments within municipalities. We show that we obtain similar estimates in the subsample of Brazilian municipalities where the land endowment did not increase. Similarly, we show that contemporaneous migration patterns are consistent with the predictions of the model: there is out (in) migration in areas more affected by labor augmenting (land augmenting) technical change. Third, we show that our estimates are not driven by pre-existing trends in manufacturing employment nor migration flows. Fourth, we show that our results are robust to using a larger unit of observation, micro-regions. Fifth, we show that at least 60% of our estimated effect of agricultural technical change on the manufacturing employment share is not driven by the processing of soy and maize in downstream industries nor larger agricultural sector demand for manufacturing inputs. Sixth, we show that our estimates are not driven by contemporaneous changes in commodity prices. Finally, we show that our main results remain statistically significant when we correct standard errors to account for spatial correlation.

Related Literature

There is a long tradition in economics of studying the links between agricultural productivity and industrial development. Nurkse (1953) and Rostow (1960) argued that agricultural productivity growth was an essential precondition for the industrial revolution. Schultz (1953) held the view that an agricultural surplus is a necessary condition for a country to start the development process. Classical models of structural transformation formalized their ideas by proposing two main mechanisms through which agricultural productivity can speed up industrial growth in closed economies. First, the demand channel: agricultural productivity growth rises income per capita, which generates demand for manufacturing goods if preferences are non-homothetic (Murphy et al. (1989), Kongsamut et al. (2001), Gollin et al. (2002)). The higher relative demand for manufactures generates a reallocation of labor away from agriculture. Second, the supply channel: if productivity growth in agriculture is faster than in manufacturing and these goods are complements in consumption, then the relative demand of agriculture does not grow as fast as productivity and labor reallocates towards manufacturing (Baumol (1967), Ngai and Pissarides (2007)).[6]

The view that agricultural productivity can generate manufacturing growth was challenged by scholars studying industrialization experiences in open economies. These scholars argued that high agricultural productivity can retard industrial growth as labor reallocates towards the comparative advantage sector (Mokyr (1976), Field (1978) and Wright (1979)). 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 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 towards the industrial sector even in open economies.[7]

Our work also builds on the empirical literature studying the links between agricultural productivity and economic development.[8] The closest precedent to our work is Foster and Rosenzweig (2004, 2008) who study the effects of the adoption of high-yielding-varieties (HYV) of corn, rice, sorghum, and wheat during the Green Revolution

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

[7]This 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. To see this, 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).

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

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 higher 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-augmenting. 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 paper is also related to the literature on the Dutch Disease. A classic reference is Corden and Neary (1982) who consider a three-sector open economy model with non-traded goods, where one of the traded sectors experiences a boom. A central feature of their analysis is a distinction between two effects of the boom: the labor effect and the demand effect.[9] The first effect is generated by the increase in the marginal product of labor in the booming sector, which draws workers out of other sectors. The second is generated by the higher income resulting from the boom which leads to increased demand for services. In the case of Hicks-neutral technical change both effects lead to de-industrialization. The setting we study differs from the standard framework in that technical change is labor-saving. Thus, the labor effect releases agricultural workers. As a result, the net effect of agricultural technical change on industrialization depends on the relative strength of the labor and demand effects. Our empirical analysis implies that in regions more affected by agricultural technical change labor reallocated from agriculture to manufacturing and not towards services. Our interpretation of these findings is that the differences-in-differences empirical strategy is well suited to identify the labor effect to the extent that labor markets are local. However, it might not be suitable to identify the demand effect if land owners do not reside locally and key services are provided by the federal government.

The result of a positive link between labor augmenting technical change in agriculture and manufacturing growth also speaks to the literature that has studied the role of manufacturing growth in promoting economic development. Economists have proposed two mechanisms through which reallocation of labor towards the manufacturing sector

[9]They refer to the labor effect as resource movement effect and the demand effect as spending effect. We name them differently to keep consistency with the literature discussed above.

may promote aggregate productivity and economic development. The first mechanism is static: when labor productivity is lower in agriculture than in the rest of the economy, reallocation of labor from agriculture to manufacturing will increase aggregate income (Gollin et al. (2002), Lagakos and Waugh (2013) and Gollin et al. (2014)). The second mechanism is dynamic: in growth models with dynamic economies of scale created by on-the-job accumulation of human capital (learning-by-doing), the reallocation of labor towards manufacturing can affect its productivity permanently (Krugman (1987) and Lucas (1988, 1993)). This is because a larger industrial sector can benefit from external scale economies created by learning-by-doing (Krugman (1987) 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 et al. (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 1.2 gives background information on agriculture in Brazil. Section 1.3 presents the theoretical model. Section 1.4 describes the data. Section 1.5 presents the empirical strategy and results. Section 1.6 shows the robustness of the results. Section 1.7 concludes.

1.2 Agriculture in Brazil

In this section we provide background information about recent developments in the Brazilian agricultural sector. As figure 1.1 shows, in the last decade, Brazilian labor force has been leaving the agricultural sector and moving towards in manufacturing and services. At the end of the 1990s, agriculture employed around 16 million workers, while manufacturing less than 8 million. By 2011, this gap was almost closed with agriculture and manufacturing employing, respectively, 12 and 10.5 million workers.

During the same period, agricultural productivity increased significantly, and productivity growth went hand-in-hand with an expansion in the area planted. Table 1.1 shows that the land cultivated with seasonal crops – i.e. crops produced from plants that need to be replanted after each harvest, such as soy and maize – increased by 10.4

million hectares between 1996 and 2006. Out of these 10.4 million, 6.2 million hectares were converted to soy cultivation.

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

1.2.1 Technical Change in Soy: Genetically Engineered Seeds

The first generation of GE soy seeds, the Roundup Ready (RR) variety, was commercially released in the U.S. in 1996 by the agricultural biotechnology firm Monsanto. However, cultivation of GE crops was forbidden in Brazil until 1998, when the Brazilian National Technical Commission on Biosecurity (CTNBio) authorized Monsanto to field-test GE soy in Brazil for 5-years as a first step before commercialization. Although cultivation of GE soy was prohibited, reports from the Foreign Agricultural Service of the United States Department of Agriculture (USDA) document that smuggling of GE soy seeds from Argentina – where they were approved for cultivation since 1996 – was already taking place from 2001 (USDA, 2001, p. 63). Eventually, pressure from soy farmers led the Brazilian government to legalize cultivation of GE soy seeds in 2003.[10]

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

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

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

crop residue, allowing farmers to save on production costs since less labor is required per unit of land to obtain the same output.[12]

The new technology spread quickly: in 2006 GE seeds were planted in 46.4% of the area cultivated with soy in Brazil, according to the last Agricultural Census (IBGE, 2006, p. 144). In the following years the technology continued spreading to the point that it covered 85% of the area planted with soy in Brazil in the 2011-2012 harvesting season, according to the Foreign Agricultural Service of the USDA (USDA (2012)). The timing of adoption of GE soy coincides with a fast expansion in the area planted with soy in Brazil. Figure 1.2 documents the evolution of the area planted with soy since 1980. The figure shows that this area grew steadily between 1980 and 1996, but experienced a fast expansion afterwards. In particular, note how the growth in area planted with soy accelerated after 2001, when the USDA documents that GE soy seeds started to be smuggled from Argentina.

The expansion of the area planted with soy can affect labor demand in the agricultural sector through two channels. First, by reducing the amount of agricultural workers per hectare required to cultivate soy (*within-crop effect*). As figure 1.3 shows, during the period under study the expansion of the area cultivated with soy was accompanied by an increase in the area planted per worker in soy production. As reported in table 1.2, labor intensity of soy production dropped from 28.6 workers per 1000 hectares in 1996 to 17.9 workers per 1000 hectares in 2006. Second, soy production is one of the least labor-intensive agricultural activities: table 1.2 shows that, in 2006, it required less than 20 workers per 1000 hectares against the 84 of the average seasonal crop and the 127 of the average permanent crop. As a result, the expansion of soy cultivation over areas previously devoted to other agricultural activities tended to reduce the labor intensity of

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

agricultural production (*across-crops effect*). The simultaneous expansion of the area cultivated with soy along with the decrease in labor intensity in correspondence with the introduction of GE soy seeds, explain the constant decrease in employment in soy production shown in figure 1.4.

1.2.2 Technical Change in Maize: Second Harvesting Season

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

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

Even with advanced cultivation techniques, maize is still more labor intensive than both soy and other agricultural activities like cattle ranching (see table 1.2). In the USDA Agricultural Resources Management Survey (ARMS) labor cost of maize cultivation in 2001 and 2005 were on average 1.8 and 1.4 times higher than the labor cost for soy cultivation.[13,14]

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

[14]In table 1.2 we do not report productivity for Brazilian farms whose main activity was maize cultivation because publicly available data on area in farms and number of workers by principal activity is available only for farms whose principal activity is either soy or cereals, a category that includes rice, wheat, maize and other cereals. The breakdown of area and number of worker by type of cereal cultivated is available in 1996, and in this year labor intensity in maize is 100.4 workers per 1000 ha of area cultivated, slightly above the labor intensity of all the farms that cultivate cereals (92.4, as reported in table 1.2).

Figure 1.6 documents the evolution of the area cultivated with maize since 1980. The figure shows that, although the total area devoted to maize has increased only slightly, the area devoted to second season maize has expanded steadily since the beginning of the 1990s.[15]

1.3 Model

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

1.3.1 Setup

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

$$Q_m = A_m L_m \tag{1.1}$$

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

$$Q_a = A_a \left[\gamma \left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}} + (1-\gamma)\left(A_T T_a\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{1.2}$$

where Q_a denotes production of the agricultural good, the two production factors are labor (L_a) and land (T_a), A_a is hicks-neutral technical change, A_L is labor-augmenting technical change and A_T is land-augmenting technical change. The parameter $\gamma \in (0,1)$, and the parameter $\sigma > 0$ captures the elasticity of substitution between land and labor. The production function described by equation (1.2) implies the following ratio

[15] Data on area cultivated with maize broken down by the season of harvest of maize are available only at the aggregate level. For this reason in section 1.5, when we study municipality-level data, we will not be able to distinguish between the two maize cultivation seasons.

of marginal product of land to marginal product of labor:

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

Thus, if land and labor are complements in production ($\sigma < 1$), labor-augmenting technical change is land-biased. That is, increases in A_L rise the marginal product of land relative to labor for a given amount of land per worker. Similarly, land-augmenting technical change is labor-biased. Finally, technical change is strongly labor-saving if improvements in technology reduce the marginal product of labor. In the case of labour-augmenting technical change, this requires $\frac{\partial MPL_a}{\partial A_L} < 0$, which imposes the stronger condition that the elasticity of substitution is smaller than the land share of output.[16] Note that this condition is more likely to be satisfied the more complementary are land and labor in production and the more important is land relative to labor in production.[17]

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

1.3.2 Equilibrium

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

$$P_a MPL_a = w = P_m MPL_m. \tag{1.3}$$

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

$$MPL_a = \left(\frac{P_m}{P_a}\right)^* A_m. \tag{1.4}$$

The equilibrium allocation of labor can be determined by substituting the land market clearing condition, $T_a = T$, in equation (1.4). In particular, with a C.E.S. production

[16]See Appendix A for a formal proof.

[17]See Neary (1981) and Acemoglu (2010) for general definitions of strongly labor-saving technical change.

function, the marginal product of labor takes the follwing form:

$$MPL_a = A_a A_L \left[\gamma^{\sigma} + (1-\gamma)\, \gamma^{\sigma-1} \left(\frac{A_T T}{A_L L_a} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}}. \tag{1.5}$$

A closed form solution for the equilibrium level of employment in agriculture L_a^{eq} can be obtained from equations (1.4) and (1.5) (see appendix A). In turn, the equilibrium level of employment in manufacturing, L_m^{eq}, can be determined using the labor market clearing condition, $L_m + L_a = L$. Once L_m^{eq} and L_a^{eq} are determined output in each sector can be found using the production functions described in equations (1.2) and (1.1). Equilibrium consumption is finally determined by: $\frac{\partial U/\partial C_a}{\partial U/\partial C_m} = \left(\frac{P_a}{P_m}\right)^*$ and the zero trade balance condition $(Q_a - C_a) = \left(\frac{P_m}{P_a}\right)^* (C_m - Q_m)$.

1.3.3 Technological Change and Structural Transformation

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

Hicks-neutral technical change

An increase in A_a generates a reallocation of labor from manufacturing to agriculture, that is $\frac{\partial L_a^{eq}}{\partial A_a} > 0$ and $\frac{\partial L_m^{eq}}{\partial A_a} < 0$. To see why this is the case, note that, in equlibrium, the marginal product of labor in agriculture is given by international prices and manufacturing productivity, thus it must stay constant when A_a increases. However, the increase in agricultural productivity rises the marginal product of labor in agriculture (see effect of rising A_a in equation (1.5)). Thus, employment in agriculture must increase to reduce the marginal product of labor to its equilibrium level.

Land-augmenting technical change

An increase in A_T generates a reallocation of labor from manufacturing to agriculture, that is $\frac{\partial L_a^{eq}}{\partial A_T} > 0$ and $\frac{\partial L_m^{eq}}{\partial A_T} < 0$. To see why this is the case, note that the land-augmenting technical change rises the marginal product of labor in agriculture (see effect of rising A_T in equation (1.5)). Thus, employment in agriculture must increase to bring the marginal product of labor back to its equilibrium level.

Labor-augmenting technical change

The effects of an increase in A_L on labor reallocation depends on the degree of sustituability between land and labor in agricultural production. This is because 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 equation (1.5). However, increases in A_L also generate a reduction in the amount of land per unit of labor in efficiency units $(A_T T / A_L L_a)$.This second effect tends to reduce the marginal product of labor, as can be seen in the second term of equation (1.5). Note that a reduction in the amount of land per worker generates a larger reduction in the marginal product of labor when these two factors are poor substituites. Thus, the relative strength of the two opposing effects depends on the value of the parameter σ.

A. *Strongly labor saving*

If the elasticity of substitution between land and labor is smaller than the land share of output, labor-augmenting technical change induces a reduction in the marginal product of labor in agriculture, that is $\frac{\partial MPL_a}{\partial A_L} < 0$. In equlibrium, the marginal product of labor in agriculture is given by international prices and manufacturing productivity, thus it must stay constant when A_L changes. Then, employment in agriculture must fall to bring the marginal product of labor back to its equilibrium level. Thus, an increase in A_L generates a reallocation of labor from agriculture to manufacturing, that is $\frac{\partial L_a^{eq}}{\partial A_L} < 0$ and $\frac{\partial L_m^{eq}}{\partial A_L} > 0$.

B. *Weakly labor saving*

If 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, that is $\frac{\partial MPL_a}{\partial A_L} > 0$. Then, agricultural employment must increase to bring the marginal product of labor back to its equilibrium level. Thus, an increase in A_L generates a reallocation of labor from manufacturing to agricluture, that is $\frac{\partial L_a^{eq}}{\partial A_L} > 0$ and $\frac{\partial L_m^{eq}}{\partial A_L} < 0$.

1.3.4 Empirical Predictions

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

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

1.4 Data

In this paper we use three main data sources:[18] the Agricultural Census for data on agriculture, the Population Census for data on the sectoral composition of employment and wages, and the FAO Global Agro-Ecological Zones database for potential yields

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

of soy and maize. To perform robustness checks we also use manufacturing plant-level data from the Brazilian Yearly Industrial Survey (PIA).

The Agricultural Census is released at intervals of 10 years by the *Instituto Brasileiro de Geografia e Estatística* (IBGE), the Brazilian National Statistical Office. We use data from the last two rounds of the census that have been carried out in 1996 and in 2006. This allows us to observe agricultural variables both before and after the introduction of genetically engineered soybean seeds, which were commercially released in the U.S. in 1996 and legalized in Brazil in 2003. The census data is collected through direct interviews with the managers of each agricultural establishment and is made available online by the IBGE aggregated at municipality level. The main variables we use from the census are: the value of agricultural production, the area devoted to agriculture in each municipality and the number of agricultural workers.[19] Out of the area devoted to agriculture in each municipality we are able to distinguish the area devoted to each crop in a given Census year. This allows us to monitor how land use has changed between 1996 and 2006.

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

[19]The number of agricultural workers 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 that work in more than one farm. The second is that this variable does not include employees hired by service provider companies that are contracted by the farm to perform agricultural activities. A detailed discussion of the variable "number of agricultural workers" in the Agricultural Census, and the potential problems associated with it, is reported in the data description section of the Appendix. Notice that these two problems do not affect the agricultural employment variables calculated using the Population Census.

[20]To perform some of the robustness checks we also use the 1980 and 1991 Population Censuses.

its wage. [21] For each municipality, we compute the employment share in manufacturing as the number of people between 16 and 55 years old working in CNAE sectors from 15 to 37 divided by the total number of people between 16 and 55 years old employed in that municipality.

Our third source of data is the Global Agro-Ecological Zones database produced by the FAO, which provides data on potential yields for soy and maize. Potential yields are the maximum yields attainable for a crop in a certain geographical area. They depend on the climate and soil conditions of that geographical area, and the level of technology available. The FAO-GAEZ database provides estimates of potential yields under different theoretical levels of technology. We focus on the two extreme levels of technology used in production: low and $high$. When the level of technology is assumed to be low, agriculture is not mechanized, it uses traditional cultivars and does not use nutrients or chemicals for pest and weed control. When the level of technology is $high$ instead, production is fully mechanized, it uses improved and high yielding varieties and "optimum" application of nutrients and chemical pest, disease and weed control.[22,23] The database reports potential yields for each crop under low and high technological levels for a worldwide grid at a resolution of 9.25 $\times$ 9.25 km. Figures 1.7 and 1.8 show the potential yields for soybean in Brazil under, respectively, low and high technology. Figure 1.9 and 1.10 show the correspondent maps for maize.

In order to match the potential yields data with agriculture and industry variables

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

[22] The description of each technology in the FAO-GAEZ dataset documentation is as follows. "Under the low input, traditional management assumption, the farming system is largely subsistence based and not necessarily market oriented. Production is based on the use of traditional cultivars (if improved cultivars are used, they are treated in the same way as local cultivars), labor intensive techniques, and no application of nutrients, no use of chemicals for pest and disease control and minimum conservation measures." "Under the high input, advanced management assumption, the farming system is mainly market oriented. Commercial production is a management objective. Production is based on improved high yielding varieties, is fully mechanized with low labor intensity and uses optimum applications of nutrients and chemical pest, disease and weed control."

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

we superimposed each of the potential yields' maps with political maps of Brazil reporting the boundaries of either municipality or micro-regions (a larger administrative unit of observation that encompass several municipalities). Next, we compute the average potential yield of all cells falling within the boundaries of every geographical unit. We repeated this operation for both soy and maize and for each of the two levels of technology. Our measure of technical change in soy or maize production within each municipality is obtained as the potential yield under high technology minus the potential yield under low technology. Figure 1.11 illustrates the resulting measure of technical change in soy at the municipality level, while Figure 1.12 shows the same measure at the micro-region level.

Finally, in order to perform some robustness checks, we use data from the *Pesquisa Industrial Anual* (PIA), the Yearly Industrial Survey carried out by the IBGE. This survey monitors the performance of Brazilian firms in the extractive and manufacturing sectors. We focus on the manufacturing sector as defined by CNAE 1.0 (sectors 15 to 37).[24] We use yearly data from 1996 to 2007. The population of firms eligible for the survey is composed by all firms with more than 5 employees registered in the national firm registry (CEMPRE, Cadastro Central de Empresas). The survey is constructed using two strata: the first includes a sample of firms having between 5 and 29 employees (*estrato amostrado*) and it is representative at the sector and state level. The second includes all firms having 30 or more employees (*estrato certo*). We focus on the sample of firms with 30 or more employees which is representative at municipality level. The variables we focus on are: total employment and average wages.

1.5 Empirics

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

[24]The broad category of CNAE 1.0 are identical to the broad categories of CNAE Domiciliar version 1.0 and of the ISIC classification version 3.1.

In the following section we report simple correlations between the expansion of the area planted with soy and maize and agricultural and industrial labor market outcomes in each municipality. Note that these correlations are not informative about the causal relation between these variables. In section 1.5.2, we present an empirical strategy that attempts to establish the direction of causality by exploiting the timing of adoption and the differential impact of the new technology on potential yields across geographical areas.

1.5.1 Basic Correlations in the Data

We start by documenting how the expansion of soy and maize cultivation during the 1996-2006 period relates to changes in agricultural production and industrial employment. These basic correlations in the data attempt to answer the following question: did areas where soy (maize) expanded experience faster (slower) structural transformation? In section 1.5.1.1 we present a set of OLS estimates of equations relating agricultural outcomes to the percentage of farm land cultivated with soy and maize. In section 1.5.1.2 we present a second set of OLS estimates of equations relating manufacturing outcomes to the percentage of farm land cultivated with soy and maize.

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

$$y_{jt} = \alpha_j + \alpha_t + \beta \left(\frac{\text{Soy Area}}{\text{Agricultural Area}} \right)_{jt} + \gamma \left(\frac{\text{Maize Area}}{\text{Agricultural Area}} \right)_{jt} + \varepsilon_{jt} \quad (1.6)$$

where j indexes municipalities, t indexes time, α_j are municipality fixed effects and α_t are time fixed effects. y_{jt} is an outcome that varies across municipalities and time and $\frac{\text{Soy (Maize) Area}}{\text{Agricultural Area}}$ is the total area reaped with soy (maize) divided by total farm land.[25,26] Our source for agricultural variables is the Agricultural Census, thus we observe them

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

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

for the years 1996 and 2006. Because fixed effects and first difference estimates are identical when considering only two periods, we estimate (1.6) in first differences:

$$\Delta y_j = \Delta\alpha + \beta\,\Delta\left(\frac{\text{Soy Area}}{\text{Agricultural Area}}\right)_j + \gamma\,\Delta\left(\frac{\text{Maize Area}}{\text{Agricultural Area}}\right)_j + \Delta\varepsilon_j \quad (1.7)$$

1.5.1.1 Agricultural Outcomes: Productivity, Labor Intensity and Employment Share

Table 1.4 reports OLS estimates of equation (1.7) for three agricultural outcomes. The first is labor productivity, measured as the value of output per worker in agriculture.[27] The second is labor intensity, measured as the number of workers per unit of land in agriculture. The third outcome is the employment share of agriculture, which attempts to capture the extent of structural transformation.[28]

The first two columns of table 1.4 show that in areas where soy cultivation expanded, the value of agricultural production per worker increased and labor intensity in agriculture decreased. In contrast, in areas where maize cultivation expanded the labor intensity decreased. This evidence is consistent with our characterization of technical change in soy as land-biased and technical change in maize as labor-biased. The estimated coefficients imply that a 1 percentage point increase in soy area share corresponds to a 0.58 percent increase in labor productivity, and a 0.48 percent reduction in labor intensity. In the case of maize, the estimated coefficients imply that a 1 percentage point increase in maize area share corresponds to a 1.6 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 (1.7) relates changes in employment shares between 2000 and 2010 to changes in the area planted with soy and maize between

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

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

1996 and 2006. In both cases the initial year precedes the timing of legalization of soybean seeds in Brazil (2003), as well as the first date in which smuggling of GE soy seeds was documented (2001). Column 3 of table 1.4 shows that the employment share of agriculture decreased in places where soy expanded while estimates for maize are not statsitically significant. The estimated coefficient implies that a 1 percentage point increase in soy area share corresponds to a 0.09 percentage point reduction in the agricultural employment share.

1.5.1.2 Manufacturing Outcomes: Employment Share, Total Employment and Wages

We now turn to the question of whether manufacturing employment expanded (contracted) in areas where soy (maize) expanded. Table 1.5 reports OLS estimates of equation (1.7) for three manufacturing outcomes: manufacturing employment share, the level of employment in manufacturing, and the average wage in the manufacturing sector.

The first column of table 1.5 shows that municipalities where soy expanded experienced a faster increase in the employment share in manufacturing. In contrast, this share remained unchanged in municipalities where maize expanded. Interestingly, in areas where soy expanded, not only the share but also the level of manufacturing employment increased, as shown in column 2. The estimated coefficient on the effect of the expansion of soy cultivation in manufacturing employment share indicates that municipalities experiencing a 1 percentage point increase in soy area share had a 0.11 percentage point increase in manufacturing employment share and a 1.05 percent increase in manufacturing employment.

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

1.5.2 The Effect of Agricultural Technological Change on Structural Transformation

In this section we provide empirical evidence on the causal effects of the adoption of new agricultural technologies on industrial development in Brazil. The basic correlations in the data reported in the previous section show that areas where soy expanded experienced an increase in output per worker and a reduction in labor intensity in agriculture while industrial employment expanded. These findings are consistent with the sequence of events predicted by the model, namely that the adoption of strongly labor saving agricultural technologies reduces labor demand in the agricultural sector and induces a reallocation of labor towards the industrial sector. However, these correlations are not informative about the direction of causality. For example, these correlations are consistent with the following alternative sequence of events: productivity growth in the industrial sector increases labor demand and wages, inducing agricultural firms to switch to less labor-intensive crops, like soy. In this section we attempt to establish the direction of causality.

Our empirical strategy relies on the assumption that goods can be traded across geographical areas of Brazil but labor markets are local. We investigate whether exogenous shocks to local agricultural productivity lead to changes in the size of the local industrial sector. Thus, our ideal unit of observation would be a region containing a city and its hinterland with limited migration across regions. We attempt to approximate this ideal using municipalities as our main level of geographical aggregation. This approach is adequate for municipalities in the interior of the country, which typically include both rural and urban areas. However, municipalities tend to be mostly urban in more densely populated coastal areas. To address this concern, we show that our estimates are robust to using a larger unit of observation: micro-regions.[29] Figures 1.11 and 1.12 contain maps of Brazil displaying both levels of aggregation.

We propose to identify the causal effect the new technologies on structural transformation by exploiting the timing of adoption and the differential impact of the new technology on potential yields across geographical areas. Let us first consider whether the timing of adoption is likely to be exogenous with respect to developments in the

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

Brazilian economy. GE soy seeds were commercially released in the U.S. in 1996, and legalized in Brazil in 2003. Given that the seeds were developed in the U.S., their date of approval for commercialization in the U.S., 1996, is arguably exogenous with respect to developments in the Brazilian economy. In contrast, the date of legalization, 2003, responded partly to pressure from Brazilian farmers. In addition, smuggling of GE soy 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 by FAO using an agricultural model that predicts yields for each crop given climate and soil conditions. As potential yields are a function of weather and soil characteristics, not of actual yields in Brazil, they can be used as a source of exogenous variation in agricultural productivity across geographical areas. Crucially for our analysis, the database reports potential yields under different technologies or input combinations. Yields under the low technology are described as those obtained using traditional seeds and no use of chemicals, while yields under the high technology are obtained using improved seeds, optimum application of fertilizers and herbicides and mechanization. Thus, the difference in yields between the high and low technology captures the effect of moving from traditional agriculture to a technology that uses improved seeds and optimum weed control, among other characteristics. We thus expect this increase in yields to be a good predictor of the profitability of adopting herbicide resistant GE soy seeds.

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

$$y_{jt} = \alpha_j + \alpha_t + \beta \, A_{jt}^{soy} + \varepsilon_{jt} \tag{1.8}$$

where y_{jt} is an outcome that varies across municipalities and time, j indexes municipalities, t indexes time, α_j are municipality fixed effects, α_t are time fixed effects and A^{soy}_{jt} is equal to the potential soy yield under high inputs from 2003 onwards and to the potential soy yield under low inputs in the years before 2003. A^{soy}_{jt} 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.[30] We thus estimate a first-difference version of equation (1.8):

$$\Delta y_j = \Delta\alpha + \beta\Delta A^{soy}_j + \Delta\varepsilon_{jt} \tag{1.9}$$

where the outcome of interest, Δy_j is the change in outcome variables between 1996 and 2006; ΔA^{soy}_j is the potential yield of soy under the high technology minus the potential yield of soy under the low technology. Figure 1.11 contains a map of Brazilian municipalities displaying this measure of technical change. Green municipalities are the ones where potential yields increase the most when switching to the high technology. As mentioned above, municipalities tend to have higher urbanization rates in more densely populated coastal areas. In addition, there were large migration flows from rural to urban areas during the period under study: the share of working age population residing in rural areas fell from 22% in 1991 to 14% in 2010. Since these movements have the potential to affect the process of structural transformation, we allow for differential trends for municipalities with different initial urbanization rates by including a control for the share of rural population in 1991 in equation (1.9).

In the case of maize, we follow a similar empirical strategy. Still, 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, we think it is reasonable argue that its diffusion is exogenous to developments in the local industrial sector. As noted in section 1.2,

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

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

$$\Delta y_j = \Delta\alpha + \beta\Delta A_j^{soy} + \gamma\Delta A_j^{maize} + \delta\ Rural_{j,1991} + \Delta\varepsilon_j \qquad (1.10)$$

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

A potential concern with our identification strategy is that, although the soil and weather characteristics that drive the variation in $\Delta A_j^{soy}$and ΔA_j^{maize} across geographical areas are exogenous, they might be correlated with initial levels of development across Brazilian municipalities. For example, to the extent that municipalities with heterogeneous initial levels of development experience different growth paths, our estimates could be capturing differential structural transformation trends across municipalities. To assess the extent of this potential concern we first compare 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 1.6 compares municipalities above and below the median change in potential soy yields (ΔA_j^{soy}) in terms of observable characteristics in 1991, before the introduction of GE soy.[31] 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

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

that our estimates are stable when we introduce controls for differential trends across municipalities with heterogeneous initial characteristics in our baseline specification (1.10), as follows:

$$\Delta y_j = \Delta\alpha + \beta\Delta A_j^{soy} + \gamma\Delta A_j^{maize} + \delta\ Rural_{j,1991} + \theta X_{j,1991} + \Delta\varepsilon_j \tag{1.11}$$

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

In the following subsections we report the results obtained using our measure of technical change to explain changes in agricultural production and in the sectoral composition of the economy. Section 1.5.2.1 reports the relationship between our measure of technical change and the expansion of soy and maize cultivation. Section 1.5.2.2 shows the relationship between this measure and agricultural outcomes. Finally, section 1.5.2.3 presents results using manufacturing outcomes.

1.5.2.1 Agricultural Outcomes: Soy and Maize Expansion

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

First, we expect that areas with a higher increase in potential soy yields when switching to the high technology are those adopting genetically engineered soy on a larger scale. Thus, we start by estimating equation (1.9) 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 1.7: the increase in potential soy yield predicts the expansion in GE soy area as a share of agricultural area between 1996 and 2006. The point estimate remains stable when controlling for initial municipality characteristics, as shown in column 2.

In columns 3 and 4 of table 1.7 we perform a falsification test by looking at whether our measure of technical change in soy explains the expansion in the area planted with non-GE soy. In this case, the coefficients are negative and significant. This finding supports our claim that the change in potential soy yield captures the benefits of adopting GE soy vis-à-vis traditional soy seeds.

Next, we jointly analyze the effects of technical change in soy and maize on the area planted with each crop. For this purpose, we use the broader measure of planted area with soy instead of GE soy.[32] This permits to control for municipality fixed effects by focusing on changes in area planted rather than levels. We start by estimating equation (1.11) 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 1.8. 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 the change in potential yields when switching to the high technology are good measures of crop-specific technical change in soy and maize during this period. In addition, both estimates are stable when we add controls for municipality characteristics. This finding suggests that the differential expansion of these crops across municipalities is not driven by differential trends across municipalities with different initial levels of development.

The size of the estimated coefficient on ΔA_j^{soy} implies that a one standard deviation increase in potential soy yield corresponds to a 1.1 percentage points increase in the share of soy in agricultural land (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 size of the estimated coefficient on ΔA_j^{mze} implies that a one standard deviation increase in potential maize yield corresponds to a 0.5 percentage points increase in the share of maize in agricultural land (8% of a standard deviation). This means that, in response to a 1.8 tons per hectare increase in potential maize yield, agricultural land devoted to maize increases by 426 hectares.

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

1.5.2.2 Agricultural Outcomes: Productivity, Labor Intensity and Employment Share

In this section we study the effects of agricultural technical change on productivity and employment in agriculture. Table 1.9 reports the results of estimating equation (1.11) when the dependent variables are three agricultural outcomes: the value of agricultural production per worker, labor intensity, and the share of workers employed in agriculture, all defined as in section 1.5.1.1.

The estimated coefficients on ΔA_j^{soy} indicate that areas with higher increase in potential soy yield experienced a larger increase in the value of agricultural production per worker and a larger reduction in labor intensity between 1996 and 2006 (columns 1 and 3). These findings confirm our characterization of technical change in soy as land-biased. Next, we study the effect of agricultural technical change in soy production on sectoral employment shares, starting from agriculture. Employment shares are calculated using the population Census. Therefore, our estimation of equation (1.11) relates changes in soy and maize potential yields to changes in employment shares between 2000 and 2010. The estimated coefficient on ΔA_j^{soy} indicates that areas with higher increase in potential soy yield experienced a reduction in agricultural employment share (column 5). This finding suggests that technical change in soy is not only land-biased but also strongly labor saving. Taken together, the results confirm the conclusions drawn from the simple correlations in the data reported in table 1.4. Finally, let us note that the estimated coefficients on ΔA_j^{soy} are larger in absolute value in the specifications that control for lagged municipality characteristics (columns 2, 4 and 6). This finding confirms that our estimates are not capturing differential growth trends across municipalities.

These estimates can be used to compute the elasticity of the agricultural employment share to changes in agricultural labor productivity due to GE soy adoption. We compute this elasticity as the ratio of the estimated coefficient on ΔA_j^{soy} when the outcome is agricultural employment share, and the estimated coefficient on ΔA_j^{soy} when the outcome is agricultural labor productivity.[33] Using our more conservative estimates

[33] 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 2000-2010. Thus, the elasticity estimates correspond to the effect of 4-year lagged agricultural productivity changes on employment shares.

– i.e. those that include all municipality controls in columns 2 and 6 – this ratio is equal to: $-0.021/0.131 = -0.158$.[34] The size of this elasticity implies that a 1 percent increase in agricultural labor productivity corresponds to a 0.158 percentage points decrease in agricultural employment share. To illustrate the magnitude of these estimates, we compute how much of the differences in the speed of structural transformation across Brazilian regions can soy technical change explain, as follows. Note that a municipality shocked with a one standard deviation increase in potential soy yield experienced an increase in agricultural labor productivity of 11 percent,[35] and a corresponding 1.76 percentage points decrease in agricultural employment share.[36] This estimate corresponds to 24% of a standard deviation in the change of agricultural employment share between 2000 and 2010 (7.4 percentage points, see table 1.3).

In the case of maize, the estimated coefficients on ΔA_j^{mze} indicate that areas with higher increase in potential maize yield experienced a smaller increase in the value of agricultural production per worker and a larger increase in labor intensity between 1996 and 2006 (columns 1 and 3). The first result is not consistent with the simple correlations in the data reported in table 1.4: municipalities where the area planted with maize expanded experienced an increase in the value of agricultural output per worker. This suggests that the coefficients on the expansion of the area planted in maize reported in tables 1.4 and the coefficient on the increase in potential maize yields reported in table 1.9 capture different variation. Note that both results are consistent with the predictions of the model. This is because our outcome of interest, the value of output per worker in agriculture, captures both the within and between crop effects of technical change in maize. The first effect is an increase in maize productivity and thus agricultural output per worker. However, the between crop effect of maize technical change generates a reduction in agricultural output per worker. This is because maize production is more labor-intensive than soy production, thus the value of the average product of labor is

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

[35]This number is computed multiplying one standard deviation in ΔA_j^{soy} by the estimated coefficient on ΔA_j^{soy} in our specification with municipality controls when the outcome is agricultural labor productivity (column 2 of table 1.9): $0.851 \times 0.131 = 0.111$

[36]This number is computed multiplying the predicted increase in agricultural labor productivity for one standard deviation in ΔA_j^{soy} by the elasticity of agricultural employment share to agricultural labor productivity: $0.111 \times -0.158 = -0.0176$.

lower for maize. Thus, a reallocation of labor towards maize reduces the value of output per worker in agriculture.[37] In turn, the finding that labor intensity increased in areas with higher increase in potential maize yields is consistent with the simple correlations in the data reported in table 1.4 and confirms our characterization of technical change in Maize as labor-biased. Finally, we find that areas more affected by technical change in maize production had a faster increase in the employment share of agriculture.

To sum up, the results presented in table 1.9 suggest that the introduction of new agricultural technologies in Brazil had a sizable impact on agricultural labor markets. Areas where the potential impact of GE soy adoption was higher experienced an increase in the value of agricultural production per worker, a reduction in the number of workers per unit of land, and a reduction in the employment share of agriculture. These findings are consistent with our characterization of the adoption of GE soy as a land-biased technical change. In the case of maize, areas where the potential impact of the introduction of a second harvesting season was higher experienced an increase in labor intensity and in the employment share of agriculture. This results are consistent with our characterization of the introduction of a second harvesting season as labor-biased technical change.

1.5.2.3 Manufacturing Outcomes: Employment Share, Employment and Wages

In this section we study the effect of agricultural technical change on manufacturing employment and wages. Table 1.10 reports the results of estimating equation (1.11) where the dependent variables are three manufacturing outcomes: the employment share of manufacturing, the level of manufacturing employment, and the average wage in manufacturing, all defined as in section 1.5.1.2. Manufacturing outcomes are calculated using the Population Census. Therefore, our estimation of equation (1.11) relates changes in soy and maize potential yields with changes in manufacturing outcomes between 2000

[37] A more formal explanation of the effect of labor reallocation towards 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 $\frac{PY}{L} \equiv \frac{P_m Y_m + P_s Y_s}{L} = \frac{P_m Y_m}{L_m}\frac{L_m}{L} + \frac{P_s Y_s}{L_s}\frac{L_s}{L}$. In this case, a reallocation of labor towards maize production reduces the value of output per worker in agriculture. This is because if soy production is more land-intensive than maize production ($\frac{T_s}{L_s} > \frac{T_m}{L_m}$), the value of the average product of labor is higher for soy ($\frac{P_s Y_s}{L_s} > \frac{P_m Y_m}{L_m}$). To see why this is the case, note that the zero profit conditions for maize and soy ($P_i Y_i = rT_i + wL_i$ for $i = s, m$) imply $\frac{P_i Y_i}{L_i} = r\frac{T_i}{L_i} + w$.

and 2010.

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

These estimates can be used to compute the elasticity of manufacturing employment share to changes in agricultural labor productivity due to GE soy adoption. We compute this elasticity as in section 1.5.2.2: we divide the estimated coefficient on ΔA_j^{soy} when the outcome is manufacturing employment share by the estimated coefficient on ΔA_j^{soy} when the outcome is agricultural labor productivity. When we estimate the specification including controls for lagged municipality characteristics, this ratio is equal to: $0.021/0.131 = 0.161$. This elasticity implies that a 1% increase in agricultural labor productivity corresponds to a 0.161 percentage points increase in manufacturing employment share. As in the previous section, we illustrate the magnitude of these estimates by computing how much of the differences in the speed of structural transformation across Brazilian regions can be explained by technical change in soy. Recall that a municipality shocked with a one standard deviation increase in potential soy yield experienced an increase in agricultural labor productivity of 11 percent,[38] and a corresponding 1.79 percentage points increase in manufacturing employment share.[39] This estimate corresponds to 31% of a standard deviation in the change of manufacturing employment share between 2000 and 2010 (5.7 percentage points, see table 1.3).

In the case of maize, the estimates reported in columns 1 and 2 of table 1.10 indicate

[38]This number is computed as before, by multiplying one standard deviation in ΔA_j^{soy} by the estimated coefficient on ΔA_j^{soy} in our specification with municipality controls when the outcome is agricultural labor productivity (column 2 of table 1.9): $0.851 \times 0.131 = 0.111$

[39]This number is computed multiplying the predicted increase in agricultural labor productivity for one standard deviation in ΔA_j^{soy} by the elasticity of manufacturing employment share to agricultural labor productivity: $0.111 \times 0.161 = -0.0179$.

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: technical change in maize is labor-biased thus increases labor demand in agriculture, generating an increase in wages and a reallocation of labor away from the manufacturing sector.

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

1.5.2.4 Services and Other Sectors

Note that the estimates of the effects of technical change on the agricultural and manufacturing employment shares discussed above have a similar magnitude. To make this point clearer, we reproduce them in table 1.11, where we also include estimates for the service and other sectors.[40] The point estimates of the effect 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 labor reallocated mostly from agriculture to manufacturing. The service sector is expected to behave differently from agriculture and manufacturing because it is mostly non traded, thus should respond to demand forces. Corden and Neary (1982) consider a three sector open economy model with non-traded goods, where higher agricultural productivity increases income and demand

[40]The 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.

for services. As non-traded goods need to be produced locally, some workers reallocate from agriculture towards the service sector. Our setting differs from this standard framework in that labor-saving technical change has an asymmetric effect on the income of workers and land-owners. Thus, the response of the demand for services depends on the propensity of land-owners to consume local services. Our ongoing empirical work suggests that the absence of an average effect on local services is related to the fact that in some areas of Brazil land owners do not reside locally. A further investigation of the effect of agricultural technical change on non-traded sectors is left for future work.

1.5.3 Variable Factor Endowments

The model presented in Section 1.3 describes a small open economy where goods can be freely traded but factor endowments are fixed. Our empirical strategy thus relies on the assumption that each unit of observation behaves as a small open economy: goods can be traded across municipalities but labor markets are local and there is a fixed supply of land. However, Brazil was characterized by significant internal migration flows during the last decade: in 2010 16% of the population aged between 16 and 55 years old had moved to their current municipality 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.[41] Thus, in this section, we investigate the role of migration and the expansion in the agricultural frontier.

1.5.3.1 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 towards the manufacturing sector. However, if workers could reallocate to other municipalities, some of this adjustment would occur through out-migration. To test this prediction, we estimate net migration rates for every municipality between 2000 and 2010 using data from the population cen-

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

sus. Net migration rates are defined as the number of (net) migrants in a municipality divided by its population.[42] Next, we estimate the baseline specification described by equation (1.11) using the 2000-2010 net migration rate in each municipality as dependent variable. Estimation results are presented in the first column of table 1.12. The estimated coefficient on the change in soy potential yields is negative and 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 towards 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. We compute this elasticity as in section 1.5.2.2: 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: $-0.013/0.131 = -0.097$. This elasticity implies that a 1% increase in agricultural labor productivity corresponds to a 0.097 percentage points decrease in the migration rate. This amounts to roughly a third (37%) of the reduction in the employment share of the agricultural sector.[43] Finally, let us note that 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.

The findings discussed above suggest that the presence of migration flows across municipalities dampen the effects of technical change on sectoral employment shares, as part of the adjustment occurs through migration flows. In particular, in our model, we can think of out-migration induced by labor augmenting technical change as a reduction in the labor endowment. In the model a reduction in the labor endowment results in a

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

[43] To 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 between 16 and 55 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 between 16 and 55 in 2000 (-0.097) 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 -0.059. Then, we divide this number by the estimated elasticity of agricultural employment share to changes in agricultural labor productivity (-0.158) obtaining a ratio of 0.37.

reduction in the manufacturing employment share. This is because equilibrium agricultural employment is given by the equality of the value of the marginal product of labor in agriculture and manufacturing which is unaffected by a change in the labor endowment (see equation (1.5)). In turn, the equilibrium level of employment in manufacturing is determined by the labor market clearing condition, $L_m + L_a = L$. Thus, the manufacturing employment share must fall when the labor endowment falls. Thus, 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.

1.5.3.2 Land

In this section we study the role of the expansion in the agricultural frontier. During this period the frontier expanded not only over the Amazon rainforest but also in the Cerrado. This is a tropical savanna eco-region in central Brazil where soils used to be too acidic and nutrient poor. Starting from the 1980s these soils were treated by the Brazilian Agricultural Research Corporation, EMBRAPA, which enabled agricultural activities to expand over these areas.

In the model, an expansion in the land endowment would have the same effects as land augmenting technical change. Thus, our findings that areas more affected by technical change in maize experienced an increase in the agricultural employment share could reflect differential increases in the land endowment. In contrast, our findings for the effect of soy technical change could be attenuated by an expansion of agricultural land. 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 1.13). Next, we estimate our baseline specification described by equation (1.11) 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 to 7 of table 1.12. 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 patters differ in frontier and non-frontier municipalities. Columns 2 and 3 of table 1.12 show that the effect of soy technical change on migration is similar for both samples. In contrast, the positive effect of maize technical change on migration is concentrated in non-frontier municipalities.

1.6 Robustness Checks

1.6.1 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 adress this concern, we report estimates of equation (1.11) including controls for the folowing variables: agricultural employment share, productivity, and wages.

Coefficient estimates are reported in tables 1.13 and 1.14. The estimated effects of soy technical change on agricultural and manufacturing outcomes are robust to the inclusion of these controls in the sense 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 relatively stable for the expansion of soy area and increase for output per worker, labor intensity and manufacturing wages. However, the absolute value of estimated coefficients decreases to around half for agricultural and manufacturing employment shares. The reason why estimates are affected by the inclusion of these controls 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 half 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.

In the case of maize, estimated coefficients are robust to including these additional controls except for employment shares. The estimated effect of maize technical change on the agricultural employment share falls by more than half and becomes significant only at 10 percent while the change in manufacturing employment share becomes small and not statistically significant. However, our estimates of the effect of maize technical change on agricultural labor intensity and manufacturing wages are stable and significant at 1 percent. Thus, we interpret our estimates of the effect of maize technical change as suggestive of the importance of the factor bias of technical change to determine the direction of sectoral labor reallocation.

1.6.2 Pre-Existing Trends

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

In order to test for the existence of pre-existing trends, we use data from the Population Censuses of 1980, 1991, 2000 and 2010. We thus estimate a model similar to the one presented in our baseline equation (1.11), but with an additional time period, as follows.

$$\begin{aligned}\Delta y_{jt} \quad = \quad & \alpha_t + \beta_1 \Delta A_j^{\text{soy}} + \beta_2 \Delta A_j^{\text{soy}} \times After2003_t + \gamma_1 \Delta A_j^{\text{maize}} + \\ & \gamma_2 \Delta A_j^{\text{maize}} \times After2003_t + \theta X_{jt-1} + \Delta \varepsilon_{jt} \qquad (1.12)\end{aligned}$$

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. α_t are time dummies for each decade and $After2003_t$ is a dummy equal to 1 if $t = 2010$. Thus, β_1 captures the effect of soy technical change that is common in the period before (1991-2000) and after (2000-2010) the adoption of GE soy seeds. In contrast, β_2 captures the differential effect of soy technical change after the introduction of GE soy seeds. Similarly, the coefficient γ_2 captures the differential effect of maize technical change in the period 2000-2010. Finally, X_{jt-1} are a set of ten-year-lagged municipality characteristics including the share of rural population, average income per capita, population density and literacy rate. [44]

We perform this test for two manufacturing outcomes: manufacturing employment level and manufacturing wages. We do not perform this test for manufacturing and agriculture employment shares since, due to important changes in the definition of employment introduced by the IBGE after the 1991 Census, we are unable to measure employment shares in a consistent way across the 1991 and 2000 Censuses.[45] Results for manufacturing employment are reported in column 1 of table 1.15. The coefficient on ΔA^{soy} estimates the effect of soy technical change that is common in the period before (1991-2000) and after (2000-2010) the adoption of GE soy seeds (β_1). That this coefficient is very small and not statistically different from zero indicates that there are

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

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

no pre-existing trends in manufacturing employment. In contrast, the coefficient on the interaction term between ΔA^{soy} and $After2003_t$, which estimates the differential effect of soy technical change on manufacturing employment after the introduction of GE soy seeds (β_2) is positive and precisely estimated. Similarly, in the case of maize, we do not find pre-existing trends in manufacturing employment.

Column 2 of table 1.15 shows the results of estimating equation (1.12) when the outcome variable is the average wage in manufacturing. In this case, ΔA^{soy} had an opposite effect on manufacturing wages between 1991 and 2000 with respect to the 2000-2010 period. Therefore, the existence of these pre-existing trends in manufacturing wages attenuates our baseline estimated effects of soy and maize technical change on wages in the period 2000-2010, presented in table 1.10.

Finally, we check for pre-existing 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 had a higher incentive to adopt a labor-saving technology to cope with labor scarcity. Column 3 of table 1.15 shows the results of estimating equation (1.12) when the outcome variable is net migration rate. The coefficient on ΔA^{soy} shows that there are no differential pre-existing trends in migration for areas that have a higher increase in potential soy yields. Similarly, in the case of maize, we do not find pre-existing trends in migration.

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

1.6.3 Larger Unit of Observation: Micro-Regions

In the empirical analysis performed so far we assumed that municipalities are a good approximation of the relevant labor market faced by Brazilian agricultural workers. A potential issue is that local labor market boundaries do not overlap with a municipality's administrative boundaries. In particular, some municipalities might be too small to properly capture labor flows between urban and rural areas, provided that manufacturing activities mostly take place in the former, and agricultural activities in the latter. In order to take into account this concern we aggregate our data at a larger unit of observation: micro-regions. These regions are groups of territorially contiguous municipalities created for statistical purposes by the Brazilian Statistical Institute (IBGE). Table 1.16 reports the results of estimating equation (1.11) using micro-regions as a unit of observation. The outcome variables are the same as in table 1.10: change in manufacturing employment share, change in manufacturing employment (in logs) and change in average manufacturing wage (in logs). The estimates are broadly consistent and similar in magnitude to those reported in table 1.10, both for soy and maize.

1.6.4 Input-Output Linkages

Our theoretical model predicts that agricultural technical change can have an effect on manufacturing employment through labor market forces. In the case of soy, for example, the adoption of new agricultural technologies releases agricultural workers that find employment in the manufacturing sector. In this section we investigate to which extent our findings reflect the strength of another channel through which agricultural technical change can affect manufacturing employment: input-output linkages. Soy and maize farming require inputs produced by other sectors, including manufacturing. Therefore, for example, an expansion of area farmed with soy in a given municipality might 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 1.10 could be explained by an increase in the agricultural demand for manufacturing inputs. A similar argument applies for manufacturing industries that use soy and maize as intermediate inputs, such as the food processing industry. In order to assess the contribution of these direct linkages on our estimates, we construct a measure of manufacturing employment that excludes the sectors directly linked to soy and maize production through input-output chains.

In order to identify input-output linkages in the data, we proceed as follows. First, we use the Brazilian input-output matrix (IBGE (2008)) to identify manufacturing sectors that are providing inputs, or receiving outputs, from the soy and maize sectors. Soy and maize are used directly as intermediate inputs in only one manufacturing sector: the food and beverage sector, that in 2005 purchased around half of the total Brazilian production of both crops. In the input-output matrix, the sector classification on the input side is less detailed than on the output side. Thus, instead of inputs used by farms whose main activity is soy production, we use information on inputs used by agricultural and breeding farms in general. Almost half of the inputs purchased by these farms are supplied by manufacturing sectors (49%) and four commodities account for 84% of the total value of inputs purchased by farms from the manufacturing sector. These commodities are: inorganic chemicals, fertilizers, diesel oil and maize oil, and they are entirely produced by the chemical industry, the oil refining industry and the food and beverage industry. Then, we use this information to construct measures of employment and wages in manufacturing that exclude those industries that are providing inputs, or receiving outputs, from the soy and maize sectors.

Table 1.17 reports estimates of our baseline specification described by equation (1.11) using as outcome variables this narrower definition of manufacturing employment and wages. When we exclude workers employed in sectors directly linked to soy and maize, the estimated coefficients on ΔA^{soy} tend to decrease in size with respect to our baseline estimates presented in table 1.10. However, the estimated coefficients on ΔA^{soy} still account for 62% of the effect on manufacturing employment share and for 90% of the effect on manufacturing employment shown in table 1.10.[46] The effect of technical change in soy on manufacturing wages, instead, decreases substantially, and it is not precisely estimated. In the case of maize, the estimated coefficients on ΔA^{maize} are essentially unaffected by excluding workers in downstream and upstream manufacturing sectors when the outcomes are manufacturing employment share and level. As in the case of soy, the effect on manufacturing wages decreases in size and it is not precisely estimated. Taken together, the results presented in this section imply that at least 60% 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

[46]In 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.

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% of the total estimated effect, which is an interesting avenue for further work.

1.6.5 Commodity Prices

In this section we show that our results are robust to controlling for international commodity prices. To the extent that variation in international prices of soy and maize affect agricultural outcomes in all Brazilian municipalities proportionally, their effects are captured by the constant term in equation (1.11). 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 1.14 and 1.15 display the evolution of international prices of soy and maize, expressed in 2000 US$. These figures show how the international prices of both commodities have been in an upward trend starting from year 2007. This pattern most likely does not affect our estimates when we use variation across the last two Agricultural Censuses, i.e. 1996 and 2006. If anything, for both soy and maize, their international price level in 2006 was lower than in 1996. However, when we use variation across the last two Population Censuses, i.e. between 2000 and 2010, the end of period year is characterized by high international soy and maize prices with respect to the initial year.

The data from the Population Censuses does not allow us to control for yearly variation in soy and maize prices. We therefore rely on an alternative source of data for manufacturing outcomes: the Annual Manufacturing Survey (PIA). The Annual Manufacturing Survey is carried out yearly, allowing us to both exclude years of high international commodity prices and fully control for price variation. It covers the universe of manufacturing firms with at least 30 employees in Brazil, and it is therefore representative at municipality level for this class of firms. We focus on two variables from this survey: manufacturing employment and average wages.[47]

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

We estimate an equation of the following form:

$$y_{jt} = \alpha_j + \alpha_t + \beta A_{jt}^{soy} + \gamma A_{jt}^{maize} + \sum_z \gamma_z P_t^z A_{j0}^z + \delta Rural_{j1991} \times t + \theta X_{jt1991} \times t + \varepsilon_{jt} \tag{1.13}$$

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 add as controls the share of rural population and the same set of initial municipality characteristics used in our main specification, all interacted with a time trend. In all specifications we control for both municipality and year fixed effects (α_j and α_t) and cluster standard errors at the municipality level to address potential serial correlation in the error term.

The results obtained using data from the Annual Manufacturing Survey are consistent with those obtained using the Population Census (see table 1.18): 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 1.10, and it loses statistical significance when we add all controls. Importantly, when we control for differential effects of international prices in columns 2 and 5, our point estimates do not change. In terms of magnitude, the point estimates we obtain with this specification for the coefficients on both ΔA^{soy} and ΔA^{maize} are similar to those obtained with the same outcomes using the Population Census data. To compare the magnitude of the estimated coefficients with those obtained using the Population Census data we calculate how much of the change in manufacturing employment and wages between the years 1996 and 2007 can be explained by the change in potential soy and maize yields in the same years. The estimated coefficients in column 2 and 5 in table 1.18 imply that a one standard deviation increase in potential soy yield corresponds to an increase in manufacturing employment of 8% of a standard deviation and a decrease in average manufacturing wages of 4% of a standard deviation. An increase in potential maize yield has instead the opposite effect on both manufacturing employment and wages: a one standard deviation increase in potential maize yield reduces manufacturing employment by 4% of a standard deviation and increases manufacturing wages by 4% of a standard deviation.[48]

1.6.6 Spatial Correlation

The maps we present in figures 1.7 through 1.12 suggest that the potential yield of both soy and maize are correlated across space. Therefore, in this section we show that our estimates remain significant when we allow the residuals to be correlated within geographical areas larger than a single municipality.

For the regressions in tables 1.8, 1.9 and 1.10, we compute standard errors clustered at two additional levels of aggregation: *micro-regions* and *meso-regions*.[49] Tables 1.19, 1.20 and 1.21 report the coefficients of A^{soy} and A^{maize} showed in tables 1.8, 1.9 and 1.10, respectively. The first row below the coefficients reports the baseline robust standard errors for comparison. The second and third row below these coefficients report the standard errors clustered at micro and meso-region level, along with their significance levels. Table 1.19 show how the effect of the change in potential soy yield on the expansion of soy cultivation becomes less precise after clustering at micro and meso-region level, but remains by and large significant at least at the 5% level. Moreover, tables 1.20 and 1.21 show that, in the case of soy, all our results for agriculture and manufacturing (with the sole exception of manufacturing wages when clustering at meso-region level) remain significant after clustering standard errors to allow for spatial correlation within larger units of observation.

[48]The fact that these effects are quantitatively smaller with respect to those obtained with the Population Census can be attributed to the following reasons: (1) the yearly industrial survey, differently from the Population Census, does not cover informal employment (2) the changes in manufacturing employment and wages in the Population Census are calculated between the years 2000 and 2010, while in the industrial survey they are calculated between the years 1996 and 2006 (i.e. this data covers a shorter period after the introduction of the new technology), (3) the industrial survey tends to over-represent larger firms that are more likely to hire high skilled workers with respect to low skilled workers, and are therefore less likely to be hiring former agricultural workers.

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

1.6.7 Alternative Definition of Technical Change

The measure of technical change proposed in section 1.5.2 is defined as the difference in potential yields between high and low level of technological inputs. These two levels of inputs map into specific farming techniques. The low inputs corresponds to traditional farming, using traditional seeds and no chemicals, while the high inputs corresponds to technologically advanced farming, characterized by high-yielding varieties of seeds, full mechanization and optimal use of chemicals. The FAO-GAEZ dataset also provides crop-specific potential yields under an intermediate level of technological inputs, which is characterized by improved varieties of seeds, partial mechanization and some use of chemicals. This technological level lies somewhere in between traditional and technologically advanced farming.

In this section we test the robustness of our results to an alternative definition of technical change that uses potential yields under intermediate inputs instead of low inputs to capture the level of agricultural technology before the introduction of GE seeds.

In table 1.22 we replicate the results presented in tables 1.8, 1.9 and 1.10 of the paper using a set of regressions of the following form:

$$\Delta y_j = \Delta\alpha + \beta\Delta A_j^{soy(h-m)} + \gamma\Delta A_j^{maize(h-m)} + \delta\ Rural_{j,1991} + \theta X_{j,1991} + \Delta\varepsilon_j \quad (1.14)$$

where Δy_j is change in agricultural or manufacturing outcomes, $\Delta A_j^{soy(h-m)}$ and $\Delta A_j^{maize(h-m)}$ are differences in potential yields in soy and maize between the high and the intermediate level of technological inputs, $Rural_{j,1991}$ is the share of rural population in 1991 and $X_{j,1991}$ is a set of municipality characteristics observed in 1991 including: average income per capita (in logs), population density (in logs) and literacy rate.

The outcome variables are the same as in tables 1.8, 1.9 and 1.10. The coefficient estimates presented in table 1.22 show that our main results are robust to this alternative definition of technical change in agriculture. It is possible to 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 1.5.2.2 and 1.5.2.3. The elasticities we get under this alternative specification tend to be smaller than those under our main specification. In particular, we obtain an elasticity of agricultural employment share to agricultural labor productivity of -0.115 (versus the -0.158 under our main specification). As for manufacturing, we obtain an elasticity of employment share to agricultural labor productivity of 0.087 (versus the 0.161 under our main specification).

In this section we showed that we obtain similar estimates in all our main regressions when using an alternative definition of technical change. We think the difference between high and low level of inputs is a more precise measure of technical change in agriculture than this alternative definition. This is because high and low level of technical inputs can be clearly mapped into actual agricultural technologies, while the loose definition of intermediate inputs could span different levels of actual agricultural technology between traditional farming and technologically advanced agriculture. In addition, if the improved seed varieties that are described as part of the bundle of intermediate inputs also capture part of the effect of GE seeds, then using this definition might miss part of the variation that we are trying to capture.

1.7 Final Remarks

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

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

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

We find that in municipalities where the new technology had a larger potential impact on soy yields, there was faster GE soy adoption, a reduction of labor intensity in agriculture and an expansion of manufacturing employment. In contrast, in municipalities where the new technology had a larger potential impact on maize yields, there was an increase of labor intensity in agriculture and a contraction of manufacturing employment. These different effects documented for soy and maize indicate that the factor bias of technical change is a key factor in the relationship between agricultural productivity and industrial growth in open economies.

References

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

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

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

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

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

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

Corden, W. M. and Neary, J. P. (1982). Booming sector and de-industrialisation in a small open economy. *Economic Journal*, pages 825–848.

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

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

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

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

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

Findlay, R. and Grubert, H. (1959). "Factor Intensities, Technological Progress, and the Terms of Trade". *Oxford Economic Papers*, pages 111–121.

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

Foster, A. D. and Rosenzweig, M. R. (2008). "Economic Development and the Decline of Agricultural Employment". *Handbook of Development Economics*, 4:3051–3083.

Gollin, D., Lagakos, D., and Waugh, M. E. (2014). The agricultural productivity gap. *Quarterly Journal of Economics*, 129(2):939–993.

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

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

Hornbeck, R. and Keskin, P. (2015). "Does Agriculture Generate Local Economic Spillovers? Short-run and Long-run Evidence from the Ogallala Aquifer". *American Economic Journal: Economic Policy*.

Hornbeck, R. and Naidu, S. (2014). "When the Levee Breaks: Black Migration and Economic Development in the American South". *American Economic Review*, 104(3):963–990.

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

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

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

IBGE (2008). *"Matriz de Insumo-Produto"*, volume 23. Rio de Janeiro, Brazil: Instituto Brasileiro de Geografia e Estatística (IBGE).

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

Krugman, P. (1987). The narrow moving band, the dutch disease, and the competitive consequences of mrs. thatcher: Notes on trade in the presence of dynamic scale economies. *Journal of Development Economics*, 27(1):41–55.

Kuznets, S. (1957). "Quantitative Aspects of the Economic Growth of Nations: II. Industrial Distribution of National Product and Labor Force". *Economic Development and Cultural Change*, 5:1–111.

Lagakos, D. and Waugh, M. E. (2013). Selection, agriculture, and cross-country productivity differences. *American Economic Review*, 103(2):948–80.

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

Lucas, R. J. E. (1988). On the mechanics of economic development. *Journal of Monetary Economics*, 22:3–42.

Lucas, R. J. E. (1993). Making a miracle. *Econometrica*, 61(2):251–272.

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

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

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

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

Neary, J. P. (1981). On the short-run effects of technological progress. *Oxford Economic Papers*, pages 224–233.

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

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

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

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

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

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

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

Syrquin, M. (1988). "Patterns of Structural Change". In Chenery, H. and Srinivasan, T., editors, *Handbook of Development Economics*, volume 1, chapter 7, pages 203–273. North Holland, Amsterdam and New York.

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

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

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

Figures and tables

Figure 1.1
Employment in agriculture, industry, services (1992-2009)

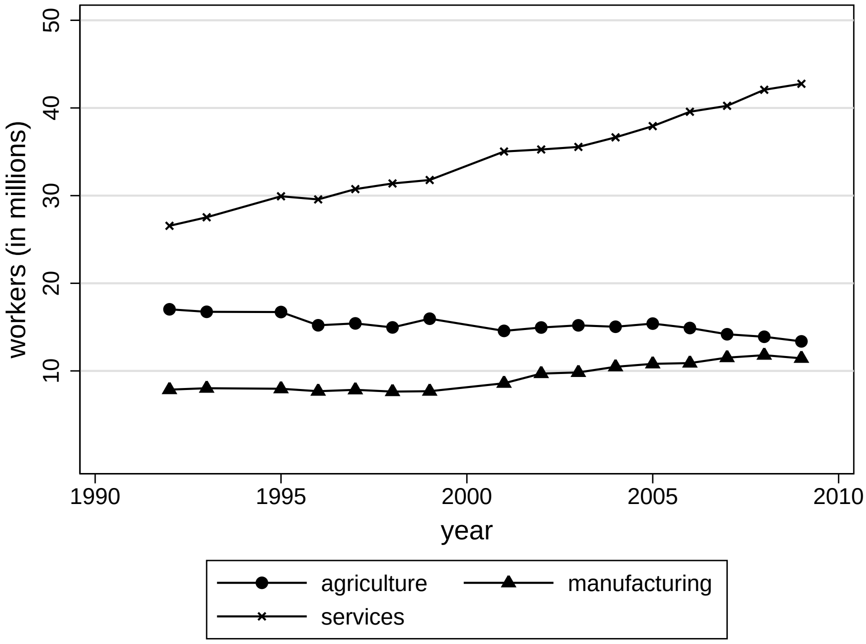

Notes: The figure shows the evolution between 1992 and 2009 of the total number of workers (expressed in million) employed by sector in Brazil. The sectors are: agriculture, industry and services. Services include: commerce, lodging and restaurants, transport, finance, housing services, other personal services, domestic workers and construction. Data source is PNAD, a national household survey representative at state level carried out yearly by the IBGE (the survey was not carried out in 1994 and in the census years: 1991, 2000 and 2010). Since the PNAD coverage changed over time, to harmonize the sample across years we exclude: (i) workers located in the states of: Rondonia, Acre, Amazonas, Roraima, Pará and Amapá (North macro-region) because only urban areas (and not rural areas) of these states were covered until 2004; (ii) workers located in the states of: Tocantins, Mato Grosso do Sul, Goias and the Distrito Federal because the sample of households in these states is not complete in the years from 1992 to 1997.

Figure 1.2
Area planted with soy (1980-2011)

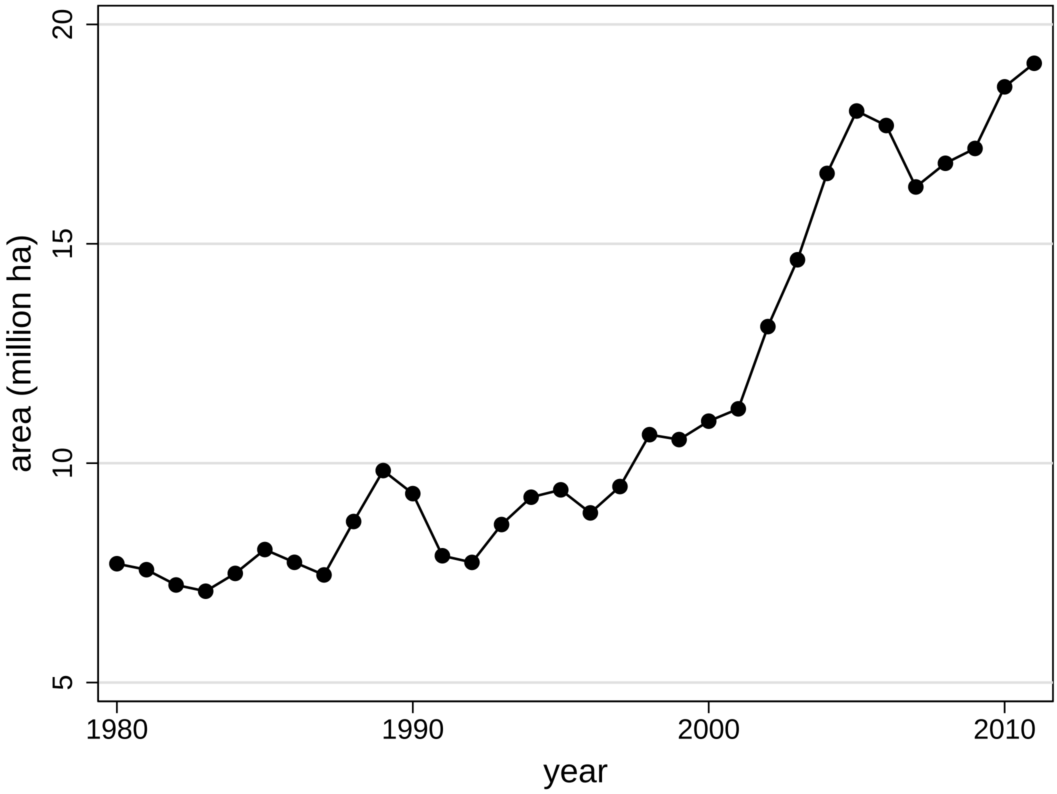

Notes: The figure depicts the evolution between 1980 and 2011 of the total area planted with soy in Brazil (expressed in million hectares). Data source is CONAB, Companhia Nacional de Abastecimento, which is an agency within the Brazilian Ministry of Agriculture. CONAB carries out monthly surveys to monitor the evolution of the harvest of all major crops in Brazil: the surveys are representative at state level and are constructed by interviewing on the ground farmers, agronomists and financial agents in the main cities of the country. All data can be downloaded at: `http://www.conab.gov.br/conteudos.php?a=1252\&t=`. In order to draw this graph we exclude land planted with soy in the states of: Rondonia, Acre, Amazonas, Roraima, Pará and Amapá (North macro-region) and Tocantins, Mato Grosso do Sul, Goias and the Distrito Federal. We do this in order to make this data comparable to the data drawn in figure 1.3.

Figure 1.3
Area planted per worker in soy production (1980-2011)

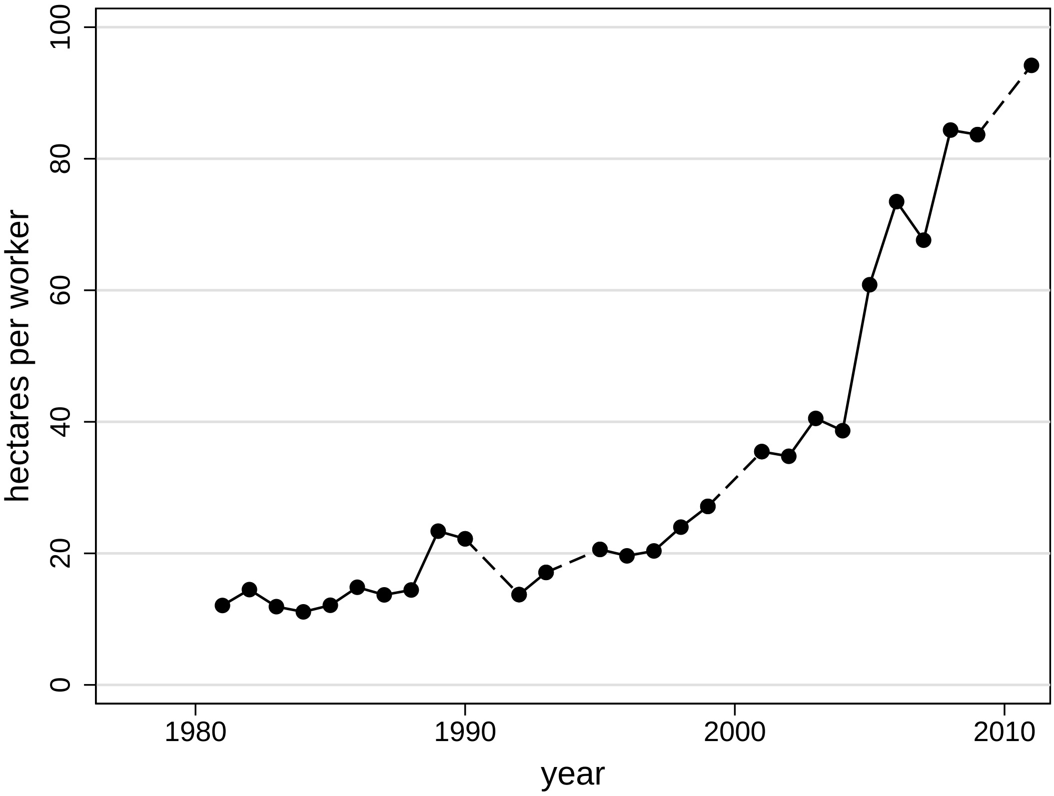

Notes: The figure depicts the evolution between 1980 and 2011 of the area planted with soy divided by the total number of workers employed in soy production in Brazil. Data sources are CONAB for area planted with soy and PNAD for the total number of workers in soy production. CONAB, Companhia Nacional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

Figure 1.4
Employment in soy production (1980-2011)

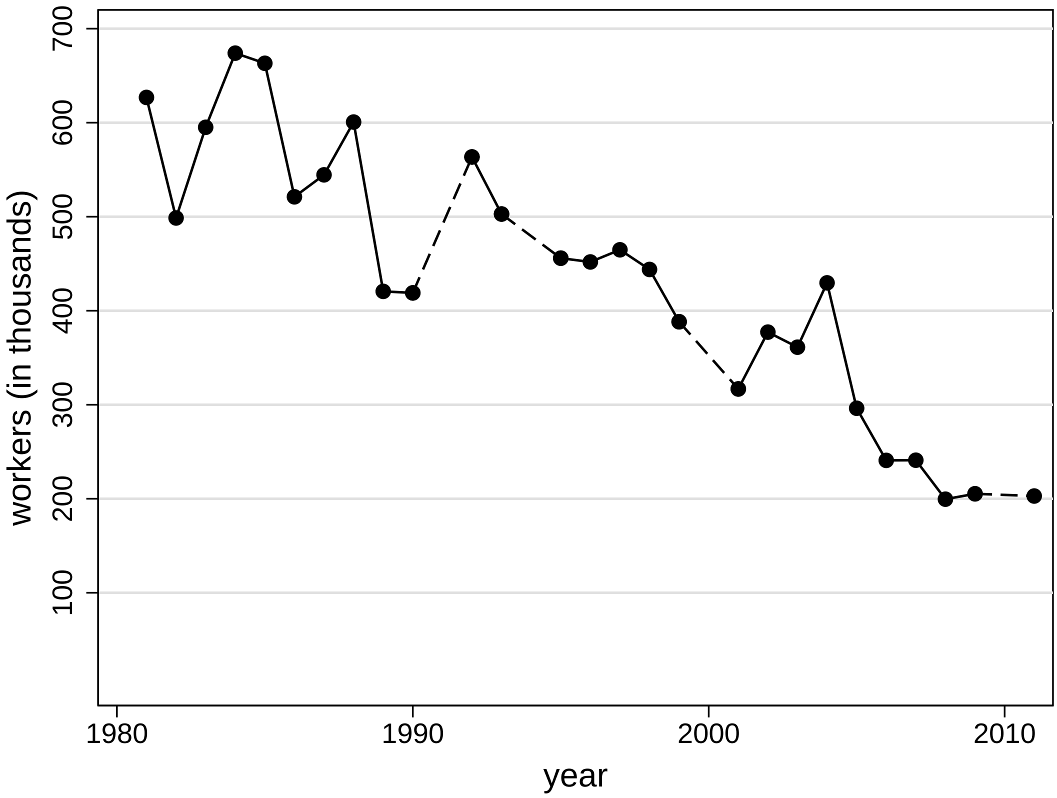

Notes: The figure depicts the evolution between 1980 and 2011 of the total number of workers employed in soy production (expressed in thousands) in Brazil. Data source is PNAD, a national household survey representative at state level carried out yearly by the IBGE (the survey was not carried out in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

Figure 1.5
Soy production per worker (1980-2011)

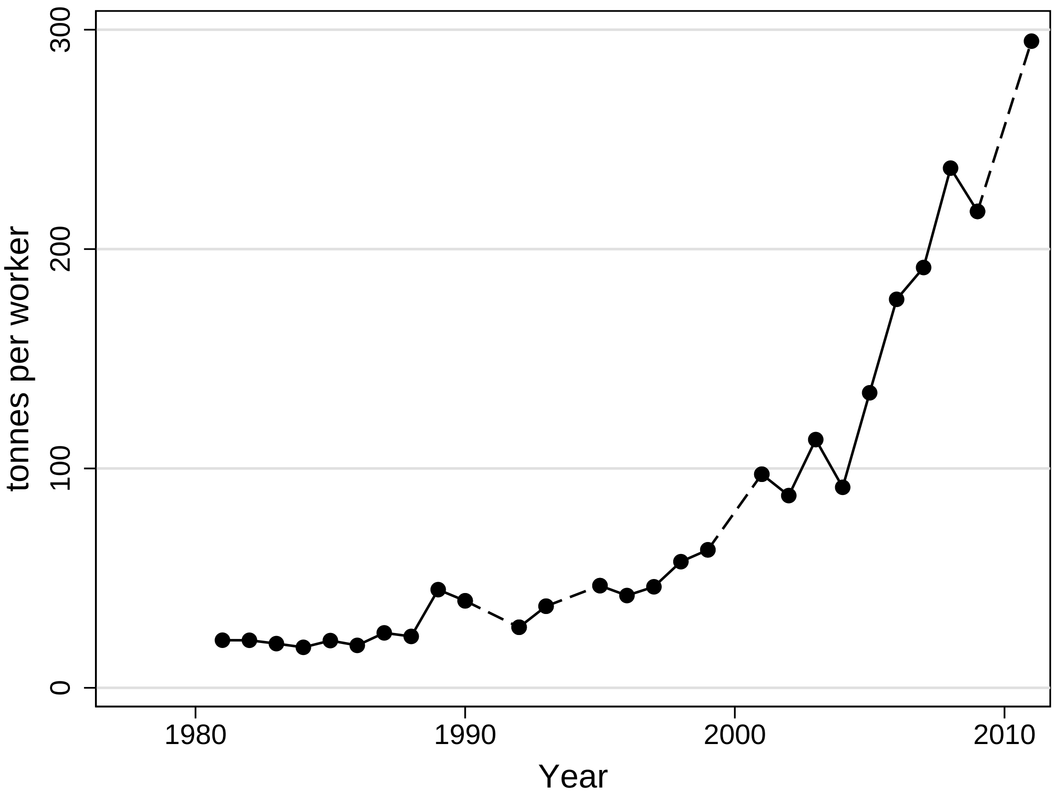

Notes: The figure depicts the evolution between 1980 and 2011 of soy production divided by the total number of workers employed in soy production in Brazil. Data sources are CONAB for soy production and PNAD for the total number of workers in soy production. CONAB, Companhia Nacional de Abastecimento,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

Figure 1.6
Area planted with maize (1980-2010)

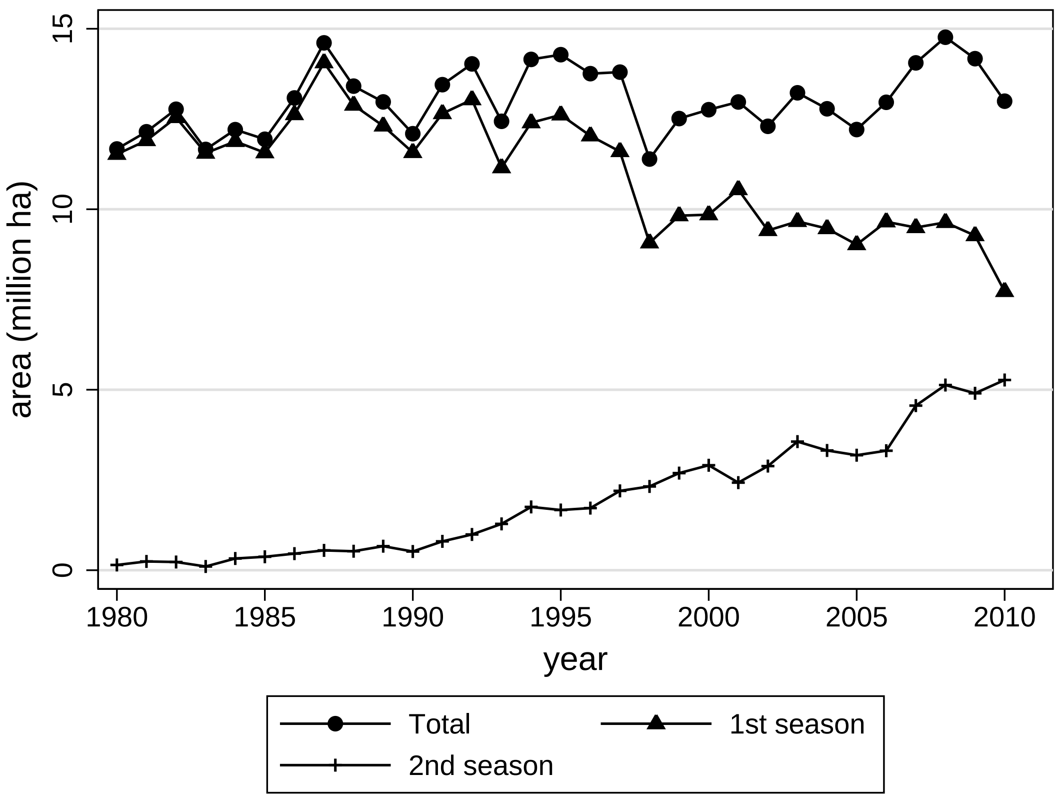

Notes: The figure depicts the evolution between 1980 and 2010 of the area planted with maize in Brazil (expressed in million hectares). The series show the total area planted with maize as well as the breakdown by the season of harvest (1st or 2nd season). Data source is CONAB, Companhia Nacional de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

Figure 1.7
Potential soy yield under low agricultural technology

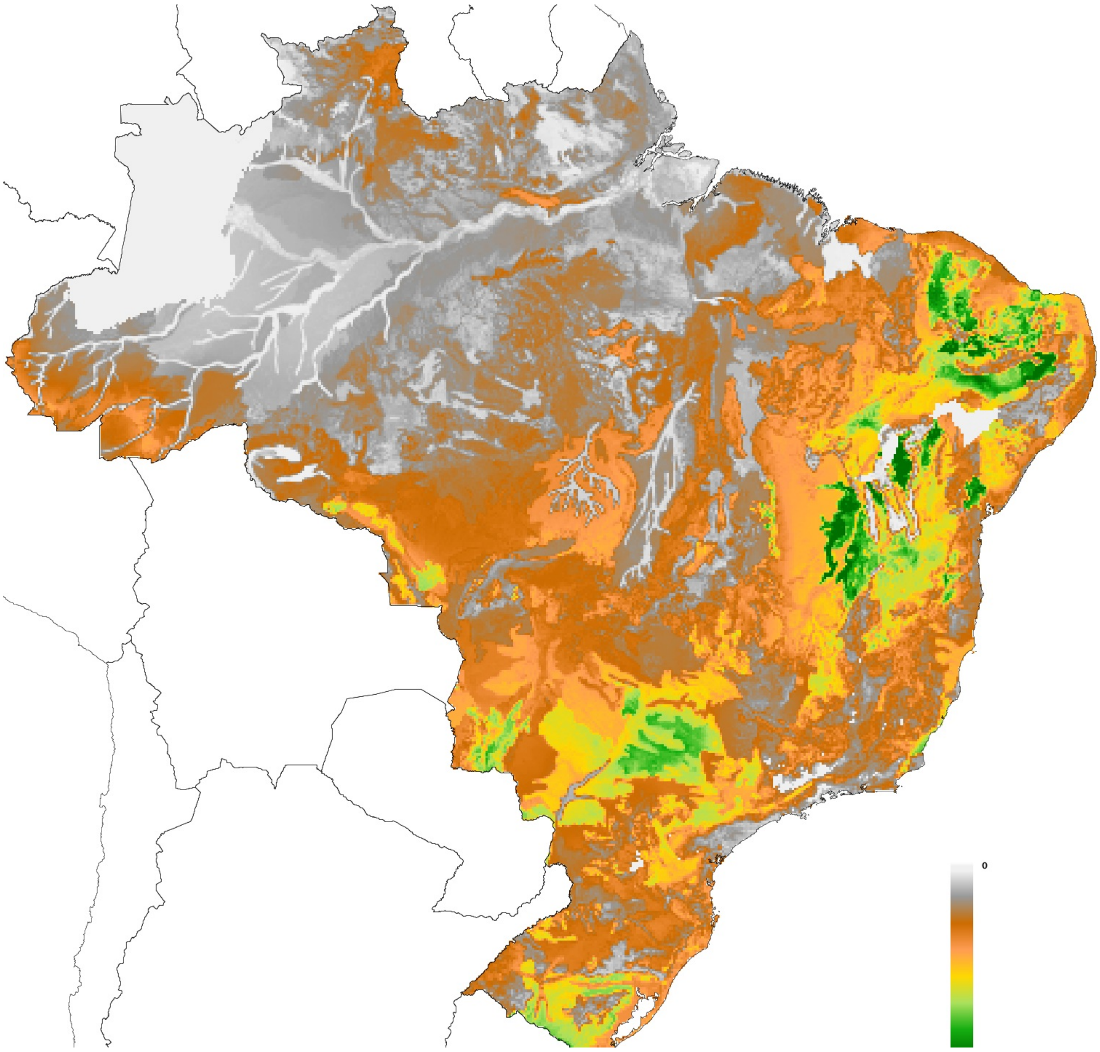

Notes: Data source is FAO-GAEZ. See section 4 for details.

Figure 1.8
Potential soy yield under high agricultural technology

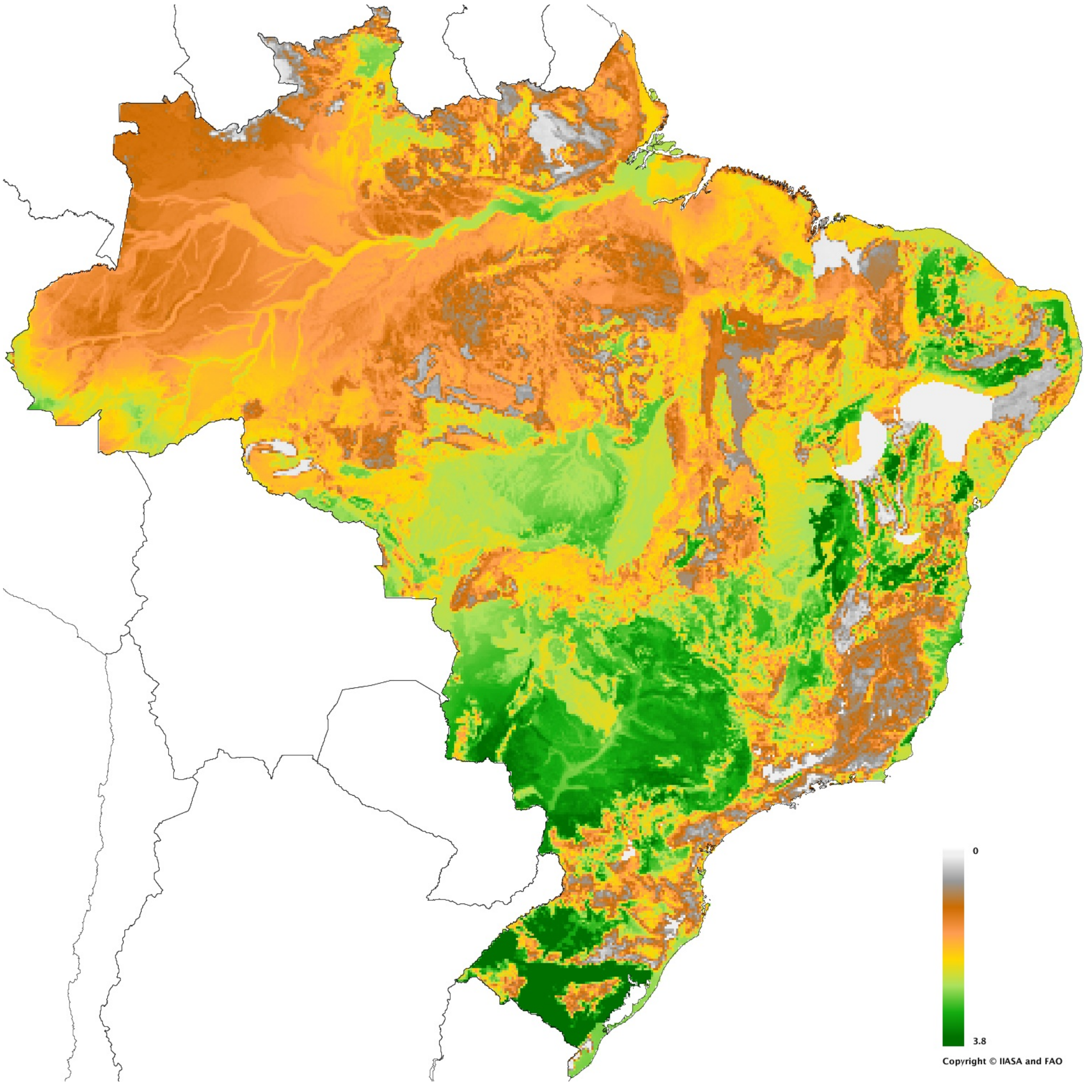

Notes: Data source is FAO-GAEZ. See section 4 for details.

Figure 1.9
Potential maize yield under low agricultural technology

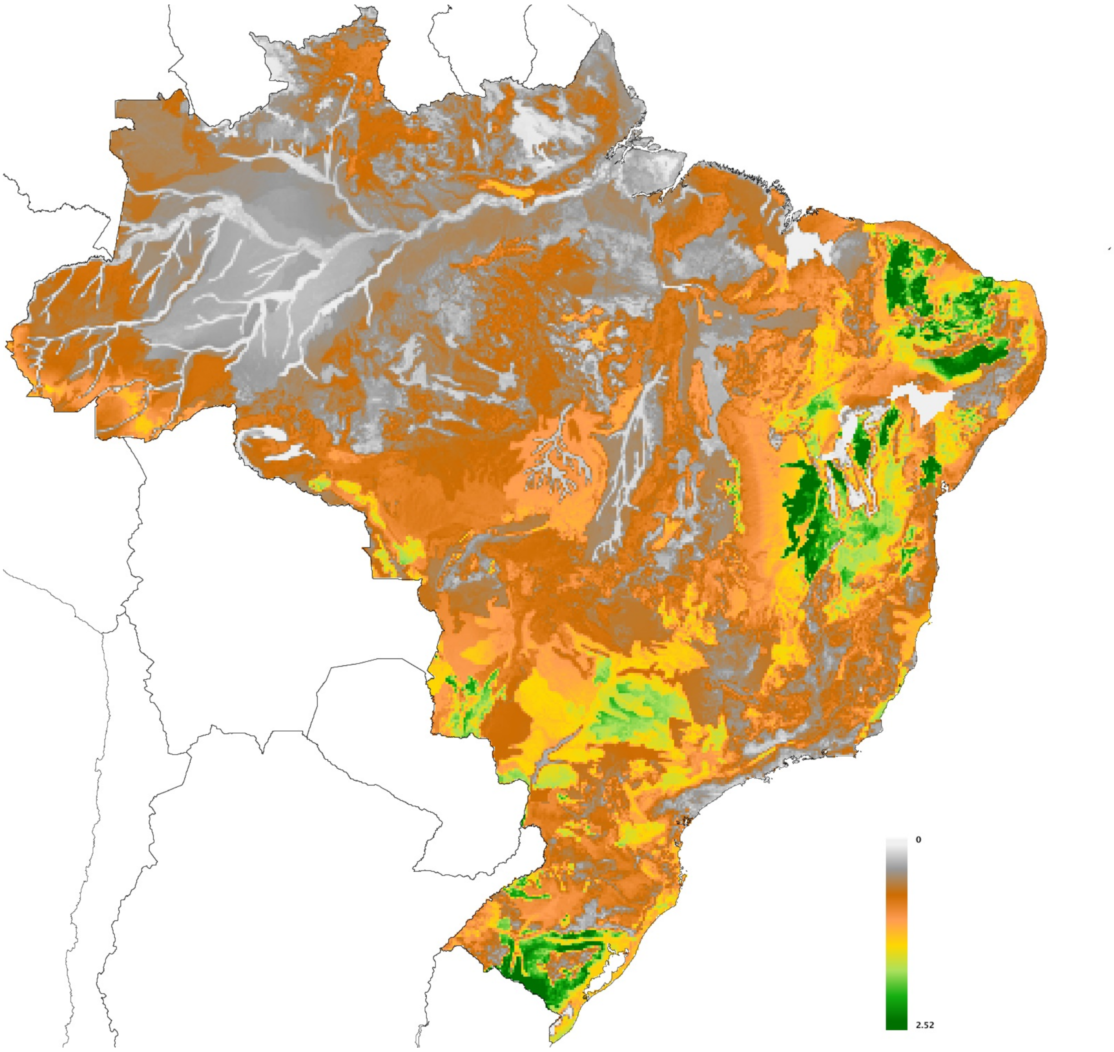

Notes: Data source is FAO-GAEZ. See section 4 for details.

Figure 1.10
Potential maize yield under high agricultural technology

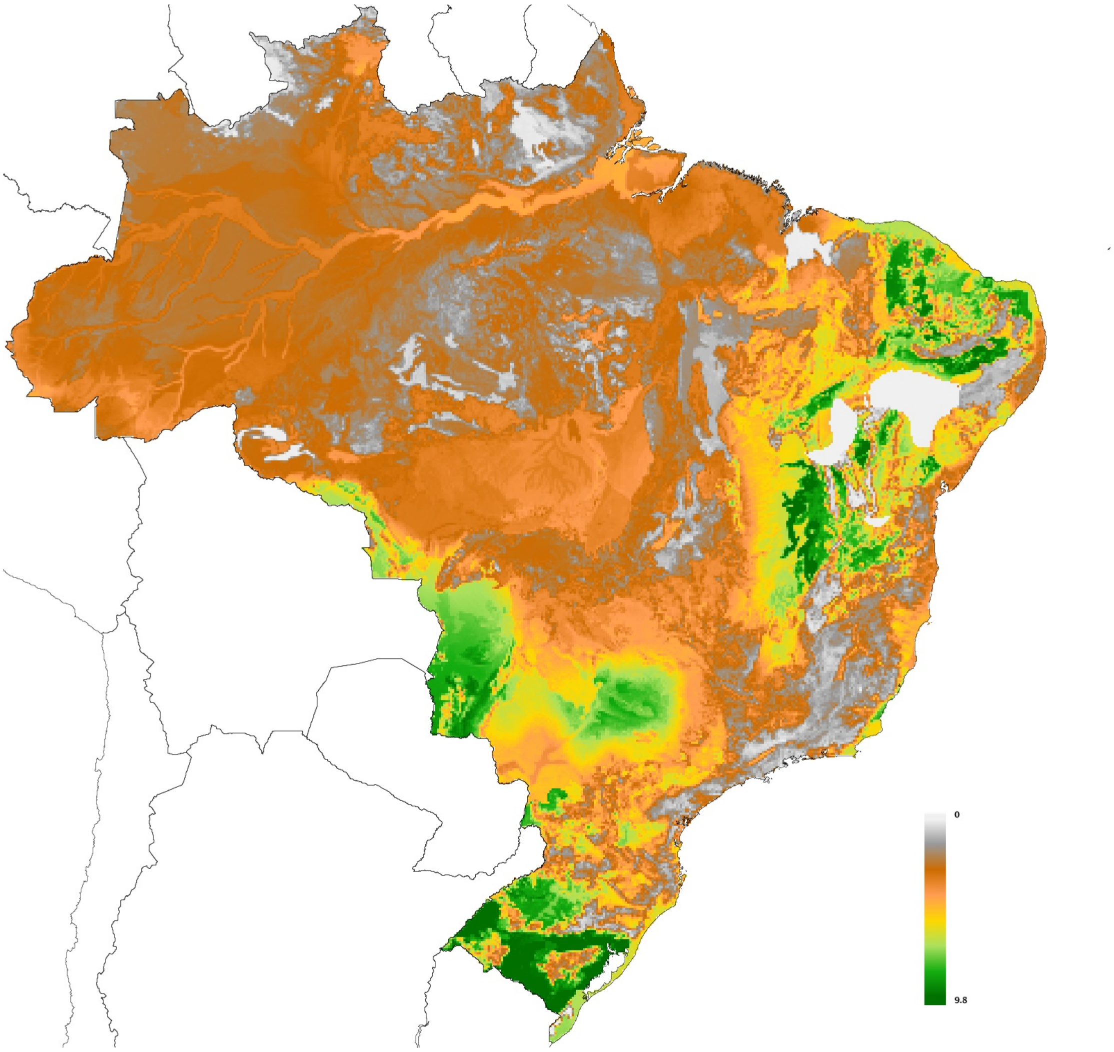

Notes: Data source is FAO-GAEZ. See section 4 for details.

Figure 1.11
Technological change in soy

Potential yield under high technology minus potential yield under low technology

Municipalities

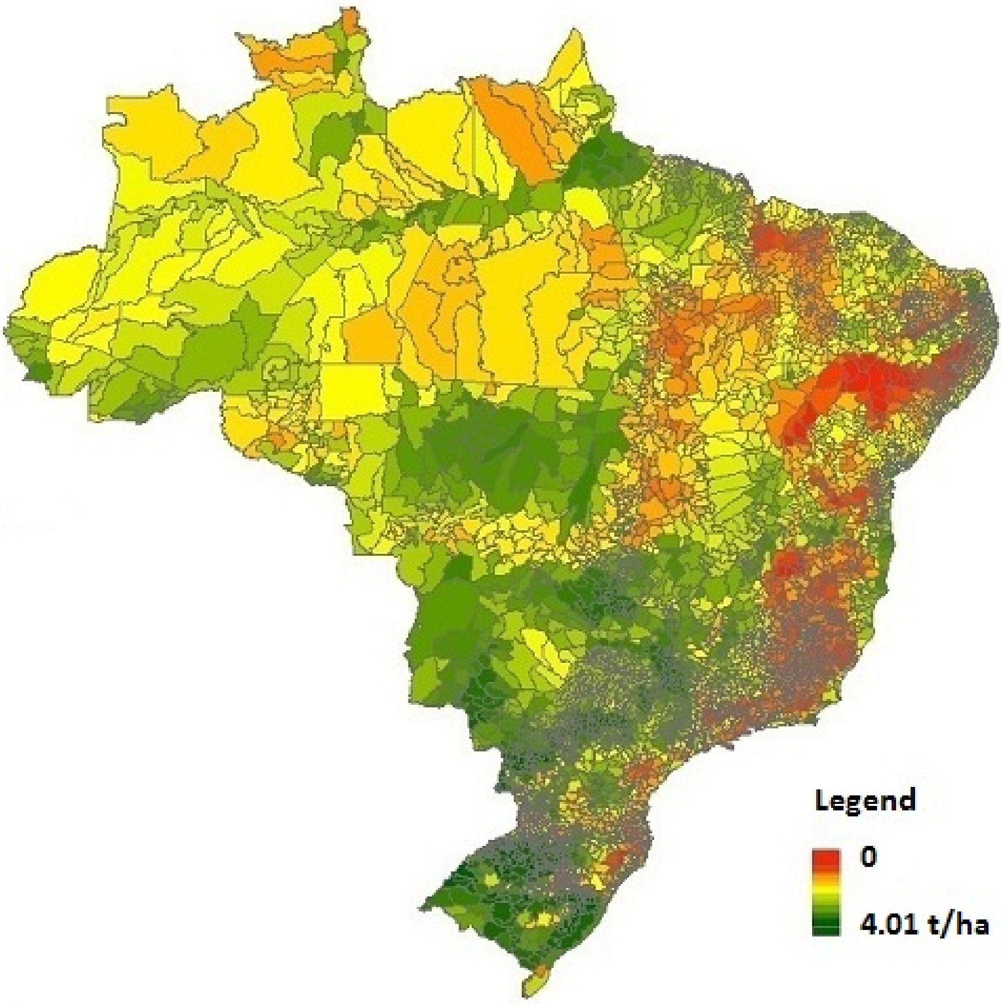

Notes: Authors' calculations from FAO-GAEZ data. See section 4 for details.

Figure 1.12
Technological change in soy

Potential yield under high technology minus potential yield under low technology

Micro-regions

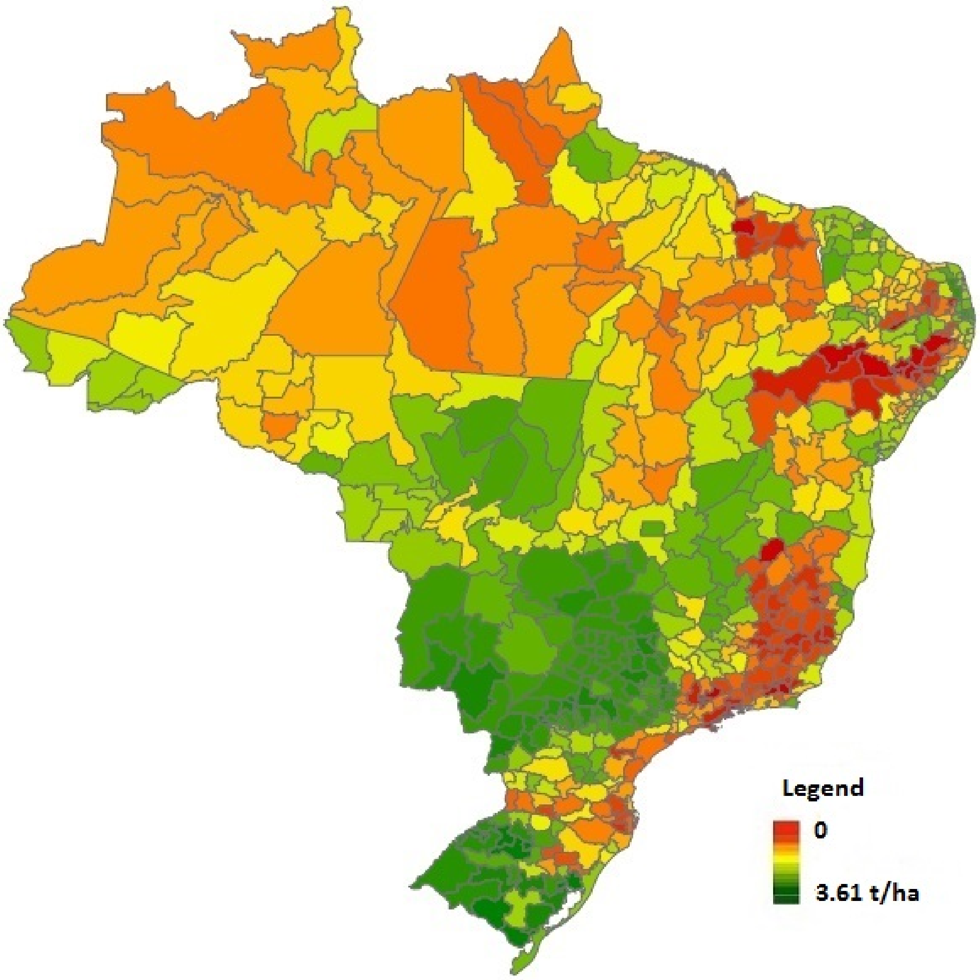

Notes: Authors' calculations from FAO-GAEZ data. See section 4 for details.

Figure 1.13
Agricultural frontier

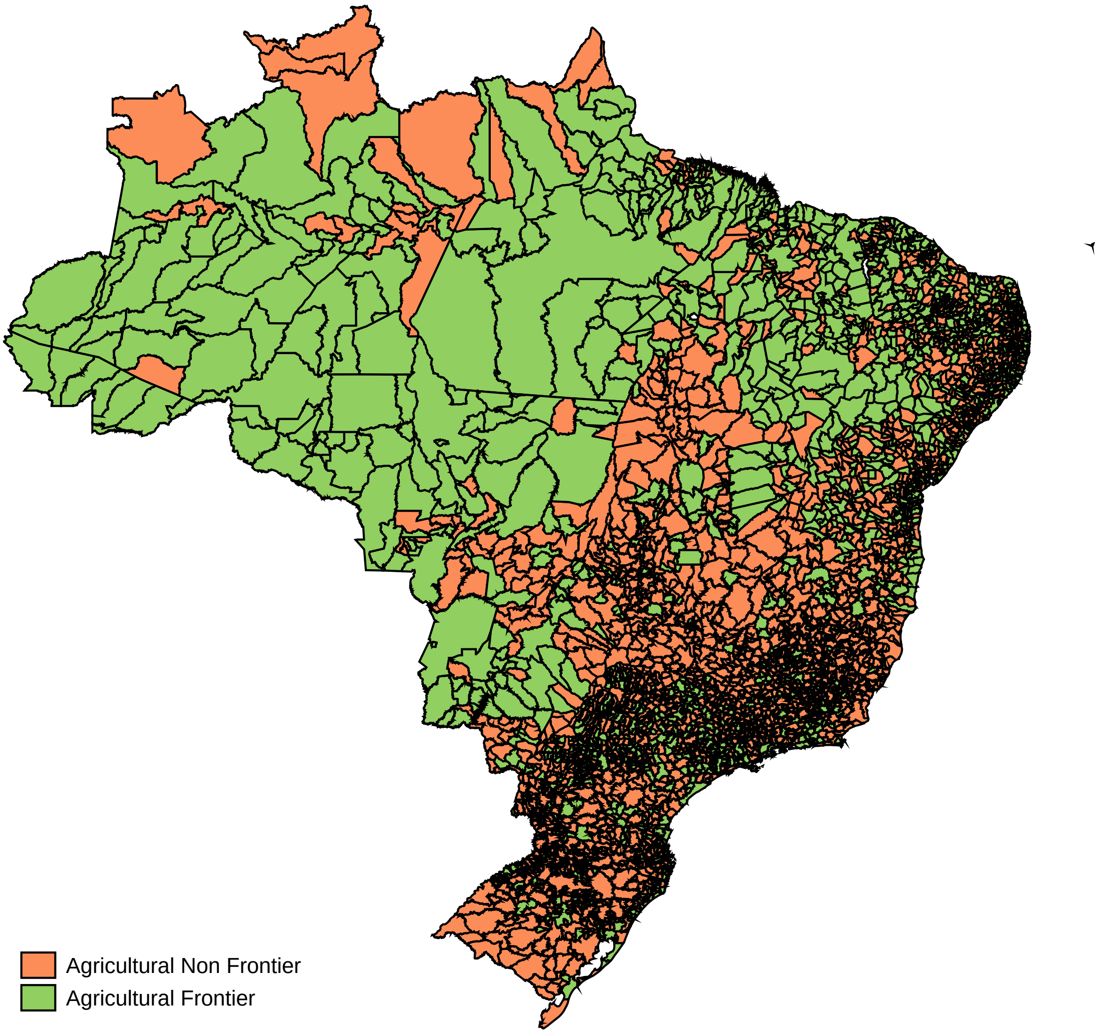

Notes: The figure shows our definition of agricultural frontier and agricultural non frontier in Brazil. We define municipalities that are part of the agricultural frontier as those that, between 1996 and 2006, experienced an increase in the use of agricultural land for the cultivation of permanent crops,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

Figure 1.14
Evolution of soy price (1990-2013)

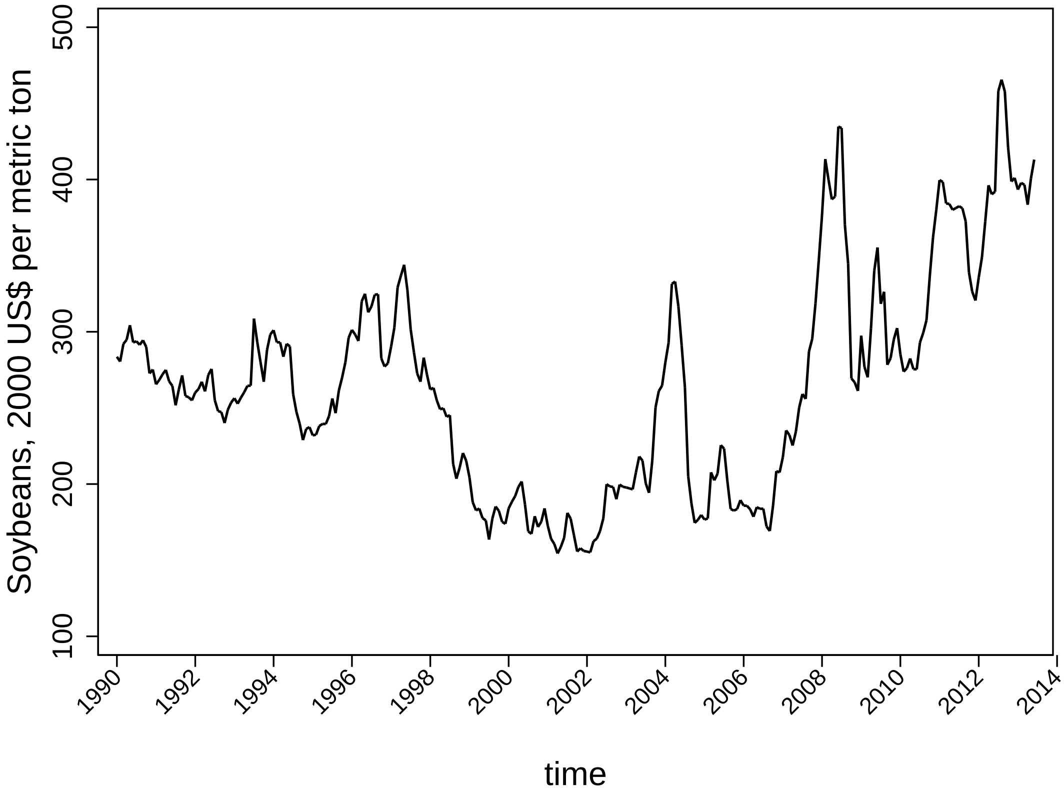

Notes: The figure shows the monthly evolution of soy real price between 1990 and 2013. Data are from the IMF Primary Commodity Prices database, series code: *PSOYB_USD*, expressed in nominal US$ per metric ton. We deflate the series using the US *Consumer Price Index for All Urban Consumers: All Items*,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

Figure 1.15
Evolution of maize price (1990-2013)

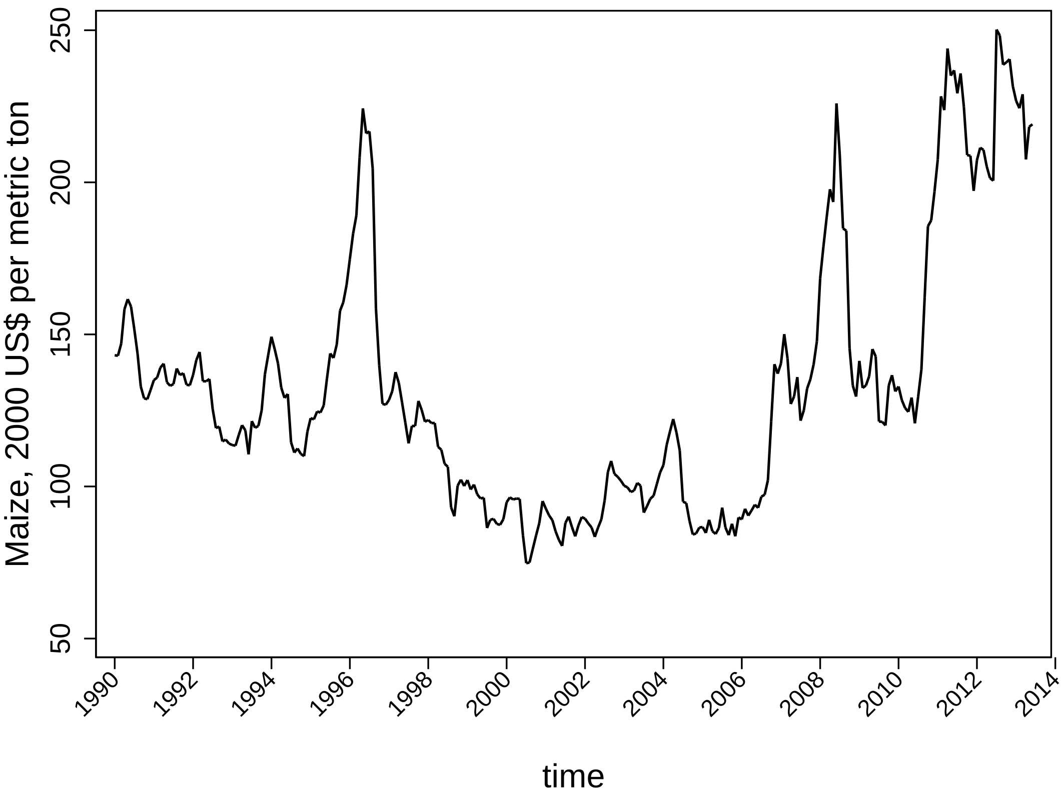

Notes: The figure shows the monthly evolution of maize real price between 1990 and 2013. Data are from the IMF Primary Commodity Prices database, series code: *PMAIZMT_US*, expressed in nominal US$ per metric ton. We deflate the series using the US *Consumer Price Index for All Urban Consumers: All Items*,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

Table 1.1
Land use (million ha)

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

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

Table 1.2
Labor intensity in Brazilian agriculture (1996-2006)

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

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

Table 1.3
Summary statistics of main variables at municipality level

PANEL A: Agricultural Census		1996 mean	1996 st.dev.	1996-2006 Change mean	1996-2006 Change st.dev.	obs.
Log labor productivity in agriculture		7.690	1.192	0.561	0.762	4,149
Log labor intensity in agriculture		-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
PANEL B: Population Census		**2000 mean**	**2000 st.dev.**	**2000-2010 Change mean**	**2000-2010 Change st.dev.**	**obs.**
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
PANEL C: Migration	**1991-2000 mean**	**1991-2000 st.dev.**	**1991-2000 obs**	**2000-2010 mean**	**2000-2010 st.dev.**	**2000-2010 obs.**
Net migration	-0.036	0.181	3992	-0.024	0.124	4149
PANEL D: FAO GAEZ	**Low inputs mean**	**Low inputs st.dev.**	**High inputs mean**	**High inputs st.dev.**	**Difference mean**	**Difference st.dev.**	**obs.**
Potential yield in soy	0.302	0.154	2.113	0.938	1.811	0.851	4,149
Potential yield in maize	0.992	0.494	4.066	2.197	3.073	1.811	4,149

Note: The table reports summary statistics of the main outcomes and independent variables used in the empirical section. In Panel A we report the main variables extracted from the Agricultural Census. In Panel B we report the main variables extracted from the Population Census. For each variable in panel A and B we report mean and standard deviation of its level in 1996 (2000) and of its change between 1996 (2000) and 2006 (2010). In Panel C we report summary statistics of the migration variable we estimate. In the first three columns of panel C we report mean, standard deviation and number of municipalities for net migration between 1991 and 2000; in the last three columns we report the same statistics for net migration between 2000 and 2010. In Panel D we report summary statistics of potential yields of soy and maize under low inputs, under high inputs, and of the difference in yields between high and low inputs. Data is from the FAO GAEZ dataset. For detailed variable definitions see Appendix B.

Table 1.4
Basic correlations in the data: agriculture

Productivity, labor intensity and employment share

	(1)	(2)	(3)
VARIABLES	Δ Log value per worker $_{2006-1996}$	Δ Log labor intensity $_{2006-1996}$	Δ Employment share $_{2010-2000}$
Δ Soy area share $_{2006-1996}$	0.583**	-0.479***	-0.090***
	(0.232)	(0.154)	(0.027)
Δ Maize area share $_{2006-1996}$	1.597***	0.737***	-0.014
	(0.184)	(0.119)	(0.019)
Observations	3,765	3,765	3,765
R-squared	0.023	0.008	0.003

Note: The table reports the OLS estimated coefficients of equation 7 in the text. The independent variables are defined as the share of farm land reaped with soy and maize. The dependent variables are reported on top of the respective columns. Value per worker is defined as total value of output of all agricultural activities divided by the total number of workers employed in agriculture. Labor intensity is the total number of workers employed in agriculture divided by total area in farms. Share of workers employed in agriculture is defined as total number of workers in agriculture divided by total number of workers in all sectors. The source of data for the independent variables and the dependent variables reported in columns 1 and 2 are the Agricultural Censuses of 1996 and 2006. The source for the employment share reported in column 3 are the Population Censuses of 2000 and 2010. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 1.5
Basic correlations in the data: manufacturing
Employment share, employment and wages

VARIABLES	(1) Δ Employment share $_{2010-2000}$	(2) Δ Log employment $_{2010-2000}$	(3) Δ Log wage $_{2010-2000}$
Δ Soy area share $_{2006-1996}$	0.106***	1.053***	0.150
	(0.022)	(0.226)	(0.113)
Δ Maize area share $_{2006-1996}$	0.001	0.018	-0.039
	(0.013)	(0.147)	(0.080)
Observations	3,765	3,765	3,765
R-squared	0.007	0.006	0.000

Note: The table reports the OLS estimates of the coefficients in equation 7 in the text. The independent variables are defined as the share of farm land reaped with soy and maize. The dependent variables are reported on top of the respective columns. Employment share in manufacturing is defined as number of people employed in the manufacturing sector (CNAE codes between 15 and 37) divided by total number of people employed in all sectors. Employment in manufacturing is the natural logarithm of people employed in the manufacturing sector. Wage is calculated as the logarithm of the average wage of manufacturing workers in 2000 Reais. The source of data for the independent variables are the Agricultural Censuses of 1996 and 2006. The source for the dependent variables are the Population Censuses of 2000 and 2010. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 1.6
Comparing municipalities below/above median increase in potential soy yield

	(1) Below ΔA^{soy} median	(2) Above ΔA^{soy} median	(3) Difference
Agricultural Employment Share $_{1991}$	0.500	0.443	-0.057***
			[0.007]
Manufacturing Employment Share $_{1991}$	0.080	0.097	0.017***
			[0.003]
Share Rural Pop $_{1991}$	0.516	0.404	-0.112***
			[0.007]
Log Income per Capita $_{1991}$	4.389	4.656	0.267***
			[0.018]
Log Pop Density $_{1991}$	3.155	3.219	0.064
			[0.041]
Literacy rate $_{1991}$	0.688	0.745	0.057***
			[0.005]
Observations	2075	2074	

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

Table 1.7
The effect of technological change on agriculture
GE soy adoption

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

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

Table 1.8
The effect of technological change on agriculture

Soy and maize expansion

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

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

Table 1.9

The effect of technological change on agriculture

Productivity, labor intensity and employment share

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Log value per worker $_{2006-1996}$		Δ Log labor intensity $_{2006-1996}$		Δ Employment share $_{2010-2000}$	
ΔA^{soy}	0.115***	0.131***	-0.057***	-0.064***	-0.018***	-0.021***
	(0.024)	(0.026)	(0.018)	(0.021)	(0.002)	(0.002)
ΔA^{maize}	-0.025**	-0.033***	0.031***	0.033***	0.005***	0.006***
	(0.011)	(0.011)	(0.008)	(0.009)	(0.001)	(0.001)
Share Rural Pop $_{1991}$	0.258***	0.125*	-0.136***	-0.177***	-0.091***	-0.076***
	(0.057)	(0.070)	(0.048)	(0.051)	(0.005)	(0.007)
Log Income per Capita $_{1991}$		-0.010		0.029		0.014***
		(0.045)		(0.039)		(0.004)
Log Pop Density $_{1991}$		-0.016		-0.017		-0.000
		(0.011)		(0.011)		(0.001)
Literacy Rate $_{1991}$		-0.270*		-0.124		-0.012
		(0.139)		(0.116)		(0.014)
Observations	4,149	4,149	4,149	4,149	4,149	4,149
R-squared	0.009	0.012	0.005	0.007	0.068	0.073

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

Table 1.10
The effect of agricultural technological change on manufacturing
Employment share, employment and wages

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Δ Employment share $_{2010-2000}$		Δ Log employment $_{2010-2000}$		Δ Log wage $_{2010-2000}$	
ΔA^{soy}	0.023***	0.021***	0.218***	0.186***	-0.032***	-0.024*
	(0.002)	(0.002)	(0.018)	(0.020)	(0.012)	(0.012)
ΔA^{maize}	-0.005***	-0.004***	-0.057***	-0.043***	0.018***	0.014**
	(0.001)	(0.001)	(0.009)	(0.009)	(0.005)	(0.005)
Share Rural Pop $_{1991}$	-0.006	0.011**	-0.186***	0.051	0.197***	-0.014
	(0.004)	(0.005)	(0.044)	(0.056)	(0.026)	(0.035)
Log Income per Capita $_{1991}$		0.002		0.093**		-0.107***
		(0.003)		(0.037)		(0.026)
Log Pop Density $_{1991}$		0.002**		0.020**		-0.035***
		(0.001)		(0.008)		(0.005)
Literacy Rate $_{1991}$		0.034***		0.197*		0.093
		(0.010)		(0.117)		(0.075)
Observations	4,149	4,149	4,149	4,149	4,149	4,149
R-squared	0.063	0.073	0.056	0.068	0.022	0.045

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

Table 1.11
The effect of agricultural technological change on sectoral employment shares
Employment shares

VARIABLES	(1) Δ Agriculture Employment share $_{2010-2000}$	(2) Δ Manufacturing Employment share $_{2010-2000}$	(3) Δ Services Employment share $_{2010-2000}$	(4) Δ Other Sectors Employment share $_{2010-2000}$
ΔA^{soy}	-0.021***	0.021***	-0.002	0.001
	(0.002)	(0.002)	(0.002)	(0.001)
ΔA^{maize}	0.006***	-0.004***	-0.000	-0.001***
	(0.001)	(0.001)	(0.001)	(0.001)
Share Rural Pop $_{1991}$	-0.076***	0.011**	0.043***	0.023***
	(0.007)	(0.005)	(0.005)	(0.004)
Log Income per Capita $_{1991}$	0.014***	0.002	-0.015***	-0.001
	(0.004)	(0.003)	(0.003)	(0.002)
Log Pop Density $_{1991}$	-0.000	0.002**	0.000	-0.002***
	(0.001)	(0.001)	(0.001)	(0.001)
Literacy Rate $_{1991}$	-0.012	0.034***	-0.009	-0.013*
	(0.014)	(0.010)	(0.010)	(0.007)
Observations	4,149	4,149	4,149	4,149
R-squared	0.073	0.073	0.103	0.045

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

Table 1.12
Variable factor endowment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Migration rate $_{2010-2000}$			Δ Agri Empl. Share $_{2010-2000}$		Δ Manuf Empl. Share $_{2010-2000}$	
VARIABLES	All	Non-Frontier	Frontier	Non-Frontier	Frontier	Non-Frontier	Frontier
ΔA^{soy}	-0.013***	-0.015***	-0.012**	-0.023***	-0.020***	0.023***	0.019***
	(0.004)	(0.005)	(0.006)	(0.003)	(0.004)	(0.002)	(0.004)
ΔA^{maize}	0.006***	0.007***	0.003	0.008***	0.003	-0.005***	-0.003*
	(0.002)	(0.002)	(0.003)	(0.001)	(0.002)	(0.001)	(0.002)
Share Rural Pop $_{1991}$	-0.078***	-0.095***	-0.035*	-0.081***	-0.061***	0.019***	-0.004
	(0.011)	(0.014)	(0.020)	(0.008)	(0.012)	(0.006)	(0.009)
Log Income per Capita $_{1991}$	0.051***	0.050***	0.047***	0.017***	0.008	0.006	-0.003
	(0.008)	(0.009)	(0.013)	(0.005)	(0.007)	(0.004)	(0.005)
Log Pop Density $_{1991}$	-0.006***	-0.002	-0.009***	-0.001	0.001	0.001	0.001
	(0.002)	(0.003)	(0.003)	(0.001)	(0.002)	(0.001)	(0.001)
Literacy Rate $_{1991}$	0.009	0.018	0.079**	-0.026	0.032	0.018	0.038**
	(0.023)	(0.027)	(0.038)	(0.017)	(0.024)	(0.012)	(0.016)
Observations	4,149	2,617	1,532	2,617	1,532	2,617	1,532
R-squared	0.104	0.119	0.113	0.080	0.076	0.076	0.066

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

Table 1.13
The effect of technological change on agriculture
Robustness of results displayed in tables 1.8 and 1.9 to controlling for additional initial municipality characteristics

VARIABLES	(1) Δ Soy area share $_{2006-1996}$	(2) Δ Maize area share $_{2006-1996}$	(3) Δ Log value per worker $_{2006-1996}$	(4) Δ Log labor intensity $_{2006-1996}$	(5) Δ Employment share $_{2010-2000}$
ΔA^{soy}	0.013***	0.001	0.164***	-0.088***	-0.008***
	(0.002)	(0.003)	(0.028)	(0.023)	(0.003)
ΔA^{maize}	-0.001	0.002**	-0.045***	0.039***	0.002
	(0.001)	(0.001)	(0.012)	(0.009)	(0.001)
Share Rural Pop $_{1991}$	0.022***	0.003	-0.030	-0.185**	0.024**
	(0.005)	(0.007)	(0.111)	(0.083)	(0.010)
Log Income per Capita $_{1991}$	0.003	0.003	0.087	0.027	0.004
	(0.003)	(0.004)	(0.054)	(0.044)	(0.005)
Log Pop Density $_{1991}$	-0.002***	0.003***	-0.018	-0.024**	-0.005***
	(0.001)	(0.001)	(0.012)	(0.011)	(0.001)
Literacy Rate $_{1991}$	0.069***	-0.004	-0.169	-0.214*	-0.015
	(0.007)	(0.013)	(0.144)	(0.122)	(0.014)
Log Agri Labor Prod $_{t-2}$	-0.001	-0.002	-0.081***	0.041**	-0.005***
	(0.001)	(0.002)	(0.021)	(0.019)	(0.002)
Log Avg Agri Wage $_{1991}$	-0.004**	-0.009***	-0.014	-0.037	-0.001
	(0.002)	(0.003)	(0.033)	(0.028)	(0.003)
Agri Employment Share $_{1991}$	0.024***	0.014	0.231*	0.031	-0.163***
	(0.006)	(0.009)	(0.124)	(0.096)	(0.011)
Observations	3,683	3,872	3,985	3,985	3,985
R-squared	0.121	0.015	0.017	0.011	0.134

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

Table 1.14

The effect of technological change on manufacturing

Robustness of results displayed in table 1.10 to controlling for additional initial municipality characteristics

VARIABLES	(1) Δ Employment share $_{2010-2000}$	(2) Δ Log employment $_{2010-2000}$	(3) Δ Log wage $_{2010-2000}$
ΔA^{soy}	0.009***	0.087***	-0.028**
	(0.002)	(0.022)	(0.014)
ΔA^{maize}	-0.001	-0.014	0.016***
	(0.001)	(0.009)	(0.006)
Share Rural Pop $_{1991}$	-0.052***	-0.185**	-0.128**
	(0.008)	(0.079)	(0.050)
Log Income per Capita $_{1991}$	0.007**	0.055	-0.076**
	(0.003)	(0.042)	(0.032)
Log Pop Density $_{1991}$	0.005***	0.037***	-0.031***
	(0.001)	(0.008)	(0.005)
Literacy Rate $_{1991}$	0.021**	0.005	0.086
	(0.010)	(0.117)	(0.078)
Log Agri Labor Prod $_{t-2}$	0.011***	0.119***	-0.000
	(0.001)	(0.013)	(0.008)
Log Avg Agri Wage $_{1991}$	-0.010***	-0.027	-0.031**
	(0.002)	(0.023)	(0.016)
Agri Employment Share $_{1991}$	0.107***	0.436***	0.174***
	(0.009)	(0.088)	(0.055)
Observations	3,985	3,985	3,985
R-squared	0.146	0.101	0.048

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

Table 1.15

The effect of agricultural technological change on manufacturing and migration

Robustness check of results reported in tables 1.10 and 1.12: controlling for pre-existing trends

VARIABLES	(1) Δ Log employment $_t$	(2) Δ Log wage $_t$	(3) Migration rate $_t$
ΔA^{soy} $\times$ (After 2003) $_t$	0.234***	-0.117***	-0.015**
	(0.025)	(0.018)	(0.006)
ΔA^{maize} $\times$ (After 2003) $_t$	-0.060***	0.053***	0.010***
	(0.012)	(0.008)	(0.003)
ΔA^{soy}	0.007	0.062***	-0.004
	(0.020)	(0.014)	(0.005)
ΔA^{maize}	-0.004	-0.027***	-0.003
	(0.010)	(0.006)	(0.002)
Share Rural Pop $_{t-2}$	0.255***	0.010	-0.128***
	(0.042)	(0.027)	(0.014)
Log Income per Capita $_{t-2}$	0.011	-0.058***	0.045***
	(0.027)	(0.020)	(0.006)
Log Pop Density $_{t-2}$	-0.016***	-0.003	-0.006**
	(0.006)	(0.004)	(0.002)
Literacy Rate $_{t-2}$	0.245***	0.190***	0.017
	(0.081)	(0.052)	(0.019)
(After 2003) $_t$	-0.233***	0.138***	-0.020***
	(0.032)	(0.022)	(0.006)
Observations	7,984	7,984	7,984
R-squared	0.031	0.018	0.096

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

Table 1.16
The effect of agricultural technological change on manufacturing
Robustness of results reported in table 1.10 to a larger unit of observation: micro-regions

VARIABLES	(1) Δ Employment share $_{2010-2000}$	(2) Δ Log employment $_{2010-2000}$	(3) Δ Log wage $_{2010-2000}$
ΔA^{soy}	0.017***	0.139***	-0.022
	(0.004)	(0.029)	(0.016)
ΔA^{maize}	-0.003	-0.037***	0.016**
	(0.002)	(0.013)	(0.007)
Share Rural Pop $_{1991}$	0.014	0.017	-0.103
	(0.012)	(0.121)	(0.089)
Log Income per Capita $_{1991}$	-0.002	0.058	-0.168**
	(0.007)	(0.088)	(0.073)
Log Pop Density $_{1991}$	0.004***	0.030***	-0.032***
	(0.001)	(0.011)	(0.007)
Literacy Rate $_{1991}$	0.016	0.007	0.128
	(0.021)	(0.261)	(0.180)
Observations	557	557	557
R-squared	0.101	0.107	0.239

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

Table 1.17
The effect of agricultural technological change on manufacturing
Robustness of results reported in table 1.10 to excluding sectors directly linked to soy and maize

VARIABLES	(1) Δ Employment share $_{2010-2000}$	(2) Δ Log employment $_{2010-2000}$	(3) Δ Log wage $_{2010-2000}$
ΔA^{soy}	0.013***	0.167***	-0.011
	(0.002)	(0.021)	(0.016)
ΔA^{maize}	-0.004***	-0.057***	0.010
	(0.001)	(0.009)	(0.007)
Share Rural Pop $_{t-2}$	0.012***	0.042	-0.014
	(0.004)	(0.058)	(0.044)
Log Income per Capita $_{t-2}$	-0.002	0.075*	-0.117***
	(0.002)	(0.038)	(0.027)
Log Pop Density $_{t-2}$	0.003***	0.034***	-0.040***
	(0.000)	(0.008)	(0.006)
Literacy Rate $_{t-2}$	0.025***	0.086	0.144*
	(0.007)	(0.124)	(0.084)
Observations	4,149	4,134	4,059
R-squared	0.037	0.042	0.030

Note: Employment in manufacturing sectors not directly linked to soy and maize is defined as number of people employed in all manufacturing sectors (CNAE codes between 15 to 37) except those employed in: food and beverages (code 15), manufacturing of other chemicals (code 24090) and manufacturing of goods from refined oil (code 23020). Employment share is defined as employment in these sectors divided by total number of people employed in all sectors. Employment in manufacturing is the natural logarithm of people employed in these sectors. Wage is calculated as the logarithm of the average wage of the workers employed in these sectors in 2000 Reais. ΔA^{soy} is defined as potential soy yield under high inputs minus potential soy yield under low inputs. ΔA^{maize} is defined as potential maize yield under high inputs minus potential maize yield under low inputs. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The source for the dependent variables are the Population Censuses of 2000 and 2010. In this case, changes in the dependent variable are calculated over the years 2000 and 2010. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1. The number of observations in columns 2 and 3 is smaller because in some municipalities sectors directly linked to soy and maize account for the whole manufacturing sector.

Table 1.18
The effect of agricultural technological change on manufacturing
Robustness of results reported in table 1.10 to controlling for commodity prices

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

Note: The table reports the OLS estimates of the coefficients in equation 13 in the text. The dependent variables are reported on top of the respective columns. Total employment is the natural logarithm of the total number of workers employed in manufacturing plants (CNAE codes 15 to 37) owned by firms that employ at least 30 employees within an AMC. The average wage is computed from manufacturing plants (CNAE codes 15 to 37) owned by firms that employ at least 30 employees. Wage is defined as the aggregate wage bill (in real terms) across firm within an AMC divided by total number of workers across the same firms within the same AMC. A^{soy} is defined as potential soy yield under high inputs for the years between 2003 and 2007, and the potential soy yield under low inputs for the years between 1996 and 2002. A^{maize} is defined as potential maize yield under high inputs for the years between 2003 and 2007, and potential maize yield under low inputs for the years between 1996 and 2002. $P^z A^z\ controls$ stand for the interaction of the potential yield of soy and maize under low inputs interacted with price levels of these crops between 1996 and 2007. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The source for the dependent variables is the plant-level supplement of the yearly industrial survey (PIA) for the years 1996 to 2007. The unit of observation is the AMC (Área Mínima Comparável). Standard errors clustered at AMC level are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 1.19
The effect of technological change on agriculture
Robustness of results reported in table 1.8 to correcting standard errors for spatial correlation

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

Note: The table reports the OLS coefficients of ΔA^{soy} and ΔA^{maize} from table 1.8. It reports below each coefficient three standard errors. The first are robust standard errors, and are identical to those reported in table 1.8. The second are standard errors clustered at microregion level. The third one are standard errors clustered at mesoregion level. The unit of observation is the AMC (Área Mínima Comparável). See table 1.8 for details on variable construction and sources.Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 1.20
The effect of technological change on agriculture
Robustness of results reported in table 1.9 to correcting standard errors for spatial correlation

VARIABLES	(1) Δ Log value per worker $_{2006}$	(2) Δ Log labor intensity $_{2006}$	(3) Δ Employment share $_{2010}$
ΔA^{soy}	0.131	-0.064	-0.021
Robust standard errors	(0.026)***	(0.021)***	(0.002)***
Microregion-clustered standard errors	(0.032)***	(0.026)**	(0.004)***
Mesoregion-clustered standard errors	(0.033)***	(0.030)**	(0.006)***
ΔA^{maize}	-0.033	0.033	0.006
Robust standard errors	(0.011)***	(0.009)***	(0.001)***
Microregion-clustered standard errors	(0.015)**	(0.012)***	(0.002)***
Mesoregion-clustered standard errors	(0.016)**	(0.015)**	(0.003)*
Controls	Y	Y	Y
Observations	4,149	4,149	4,149
R-squared	0.012	0.007	0.073

Note: The table reports the OLS coefficients of ΔA^{soy} and ΔA^{maize} from table 1.9. It reports below each coefficient three standard errors. The first are robust standard errors, and are identical to those reported in table 1.9. The second are standard errors clustered at microregion level. The third one are standard errors clustered at mesoregion level. The unit of observation is the AMC (Área Mínima Comparável). See table 1.9 for details on variable construction and sources.Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 1.21
The effect of technological change on manufacturing
Robustness of results reported in table 1.10 to correcting standard errors for spatial correlation

VARIABLES	(1) Δ Employment share $_{2010}$	(2) Δ Log employment $_{2010}$	(3) Δ Log wage $_{2010}$
ΔA^{soy}	0.021	0.186	-0.024
Robust standard errors	(0.002)***	(0.020)***	(0.012)*
Microregion-clustered standard errors	(0.004)***	(0.030)***	(0.014)*
Mesoregion-clustered standard errors	(0.006)***	(0.045)***	(0.019)
ΔA^{maize}	-0.004	-0.043	0.014
Robust standard errors	(0.001)***	(0.009)***	(0.005)**
Microregion-clustered standard errors	(0.002)**	(0.014)***	(0.006)**
Mesoregion-clustered standard errors	(0.003)	(0.027)	(0.008)*
Controls	Y	Y	Y
Observations	4,149	4,149	4,149
R-squared	0.073	0.068	0.045

Note: The table reports the OLS coefficients of ΔA^{soy} and ΔA^{maize} from table 1.10. It reports below each coefficient three standard errors. The first are robust standard errors, and are identical to those reported in table 1.10. The second are standard errors clustered at microregion level. The third one are standard errors clustered at mesoregion level. The unit of observation is the AMC (Área Mínima Comparável). See table 1.10 for details on variable construction and sources.Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 1.22
The effect of technological change on agriculture and manufacturing
Robustness of results reported on tables 1.8, 1.9 and 1.10 to alternative definition of technical change

	Agricultural outcomes					Manufacturing outcomes		
	Δ Soy area share	Δ Maize area share	Δ Log value per worker	Δ Log labor intensity	Δ Employment share	Δ Employment share	Δ Log Employment	Δ Log wage
Δ A $^{soy\ (h-m)}$	0.020***	0.002	0.163***	-0.145***	-0.019***	0.014***	0.155***	-0.040**
	(0.003)	(0.004)	(0.036)	(0.029)	(0.003)	(0.003)	(0.028)	(0.016)
Δ A $^{maize\ (h-m)}$	-0.001	0.004**	-0.043**	0.063***	0.005***	-0.001	-0.030**	0.023***
	(0.001)	(0.002)	(0.017)	(0.013)	(0.002)	(0.001)	(0.014)	(0.008)
Share Rural Pop $_{1991}$	0.038***	0.010	0.090	-0.179***	-0.069***	0.002	-0.017	-0.010
	(0.005)	(0.007)	(0.069)	(0.050)	(0.007)	(0.005)	(0.056)	(0.034)
Log Income p.c. $_{1991}$	-0.000	-0.005	-0.007	0.040	0.012***	0.004	0.111***	-0.106***
	(0.002)	(0.004)	(0.045)	(0.038)	(0.004)	(0.003)	(0.037)	(0.026)
Log Pop Density $_{1991}$	-0.001***	0.004***	-0.016	-0.019*	0.000	0.001	0.015*	-0.035***
	(0.001)	(0.001)	(0.011)	(0.010)	(0.001)	(0.001)	(0.008)	(0.005)
Literacy Rate $_{1991}$	0.056***	-0.009	-0.321**	-0.088	-0.007	0.030***	0.150	0.100
	(0.006)	(0.012)	(0.140)	(0.116)	(0.014)	(0.010)	(0.119)	(0.075)
Observations	3,652	3,652	4,149	4,149	4,149	4,149	4,149	4,149
R-squared	0.150	0.015	0.011	0.011	0.064	0.050	0.056	0.045

Note: The table reports the OLS estimates of the coefficients in equation 11 in the text. The dependent variables are reported on top of the respective columns. Soy and maize area share are defined as the share of farm land reaped with soy and maize. Value per worker is defined as total value of output of all agricultural activities divided by the total number of workers employed in agriculture. Labor intensity is the total number of workers employed in agriculture divided by total area in farms. Share of workers employed in agriculture is defined as total number of workers in agriculture divided by total number of workers in all sectors. Employment share in manufacturing is defined as number of people employed in the manufacturing sector (CNAE codes between 15 and 37) divided by total number of people employed in all sectors. Employment in manufacturing is the natural logarithm of people employed in the manufacturing sector. Wage is calculated as the logarithm of the average wage of manufacturing workers in 2000 Reais. $\Delta A^{soy\ (h-m)}$ is defined as potential soy yield under high inputs minus potential soy yield under intermediate level of inputs. $\Delta A^{maize\ (h-m)}$ is defined as potential maize yield under high inputs minus potential maize yield under intermediate level of inputs. The source of data for the independent variables is the FAO-GAEZ v3.0 dataset. The source for the dependent variables in columns 1-4 are the Agricultural Censuses of 1996 and 2006. The source for the dependent variables in columns 5-8 are the Population Censuses of 2000 and 2010. The unit of observation is the AMC (Área Mínima Comparável). Robust standard errors are reported in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Appendix A

Theory

This appendix contains proofs of the results stated in the Theory Section 3. As most results depend on the effect of different types of technical change on the marginal product of labor, we start by obtaining $MPL_a = \frac{\partial Q_a}{\partial L_a}$ by differentiating the agricultural production function described by equation (1.2) w.r.t. labor:

$$MPL_a = A_a \Theta^{\frac{\sigma}{\sigma-1}-1} \gamma L_a^{\frac{\sigma-1}{\sigma}-1} A_L^{\frac{\sigma-1}{\sigma}}$$

where, to save space, we define $\Theta \equiv \gamma \left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}} + (1-\gamma)\left(A_T T\right)^{\frac{\sigma-1}{\sigma}}$.

Proof that $\frac{\partial MPL_a}{\partial A_L} < 0$ iff the elasticity of substitution is smaller than the land share of output.

$$\frac{\partial MPL_a}{\partial A_L} \quad = \quad A_a \Theta^{\frac{1}{\sigma-1}} \gamma L_a^{\frac{-1}{\sigma}} A_L^{\frac{-1}{\sigma}} \frac{\sigma - 1}{\sigma} \left[1 + \frac{1}{\sigma - 1} \Theta^{-1} \gamma \left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}}\right]$$

Then, if $\sigma < 1$, $\frac{\sigma-1}{\sigma} < 0$, then $\frac{\partial MPL_a}{\partial A_L} < 0$ iff $1 + \frac{1}{\sigma-1}\Theta^{-1}\gamma\left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}} > 0$ which is true as long as σ satisfies the following condition:

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

Proof that $\frac{\partial MPL_a}{\partial L_a} < 0$:

$$\frac{\partial MPL_a}{\partial L_a} = A_a \Theta^{\frac{\sigma}{\sigma-1}-1} \gamma \frac{1}{\sigma} \left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}-2} A_L^2 \left[\Theta^{-1} \gamma \left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}} - 1\right] < 0$$

because $\gamma \in (0,1)$, and σ, A_a, A_L, A_T, L_a, T, Θ are positive and $\Theta^{-1}\gamma \left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}} < 1$.

Proof that $\frac{\partial MPL_a}{\partial A_a} > 0$:

$$\frac{\partial MPL_a}{\partial A_a} = \left[\gamma \left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}} + (1-\gamma)\left(A_T T\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}-1} \gamma \left(A_L L_a\right)^{\frac{\sigma-1}{\sigma}-1} A_L > 0$$

because $\gamma \in (0,1)$, and A_L, A_T, L_a and T are positive.

Proof that $\frac{\partial MPL_a}{\partial A_T} > 0$:

$$\frac{\partial MPL_a}{\partial A_T} = \frac{1}{\sigma} A_a \Theta^{\frac{\sigma}{\sigma-1}-2} \gamma (1-\gamma) L_a^{-\frac{1}{\sigma}} A_L^{\frac{\sigma-1}{\sigma}} A_T^{-\frac{1}{\sigma}} T^{\frac{\sigma-1}{\sigma}} > 0$$

because $\gamma \in (0,1)$, and σ, A_a, A_L, A_T, L_a, T, Θ are positive.

Appendix B

Data

This appendix contains a detailed description of the main variables used in the empirical section.

Value of output per worker in agriculture. Data on value of agricultural production per worker is from the Brazilian Agricultural Census (IBGE (1996, 2006)) and it is sourced from the IBGE online data repository SIDRA.[1]. The variable ”value of output per worker in agriculture” is defined as the logarithm of total value of agricultural production divided by the total number of workers employed in agriculture. Total value of production in agriculture in 1996 comes from tables: 500 (value of production by seasonal crop), 513 (value of production by permanent crop), 527 (value of production by horticulture product), 534 (value of production by forestry product), 551 (value of production by vegetable extraction product) and 338 (value of animals by type). Total value of production in agriculture in 2006 comes from tables: 1823 (value of production by seasonal crop), 1177/1178 (value of production by permanent crop), 818 (value of production by horticulture product), 815 (value of production by forestry product), 816 (value of production by vegetable extraction product), 782 (value of bovines), 937 (value of swines), 943 (value of poultry). Number of workers in agriculture comes from table 321 in 1996 and table 956 in 2006. In tables 1.13 and 1.14 we also use as a control for the initial level of agricultural productivity the value of output per worker in agriculture from the 1985 Agricultural Census. Data for the 1985 Census was manually extracted from the original documentation available from the IBGE online library.[2] To-

[1]At `www.sidra.ibge.gov.br`.
[2]At `www.biblioteca.ibge.gov.br`

tal value of production in agriculture in 1985 comes from table 96 (expenditures, value of production and revenues), while number of workers in agriculture comes from table 87.

The definition of agricultural workers in the Agricultural Censuses of 1996 and 2006 deserves a more detailed discussion. This variable includes: employees, family members employed in farm activities, sharecroppers and people who reside in the farm and perform agricultural activities without a formal contract. Receiving a wage is not a necessary condition to be considered an agricultural worker, so subsistence farmers are also included in this definition. In addition, workers that perform non-agricultural activities within the farm, such as truck drivers, accountants, mechanics or any other worker performing tasks in support of the main activity of the farm are counted as agricultural workers. Instead, family members or farm residents who were not employed in farm activities, domestic workers that are employed within the household of the owner but are not assigned agricultural tasks, and workers that are hired by a service provider company that has a contract with the farm are not counted as agricultural workers. There are two potential issues with this definition of agricultural workers in the Agricultural Census. The first issue is about how seasonal workers are accounted for. In particular, there is a potential double counting problem for seasonal workers that work in more than one farm. This is because the variable is collected at farm level, and not at individual level, and the definition of agricultural workers includes both permanent (employed for at least 180 days during the Census year) and seasonal workers, even when the latter are not employed on the reference date of the Census (p. 53, IBGE, 2006). Therefore, if seasonal workers are employed by more than one farm during a year, and they are recorded by each of them in the Census, the variable "number of workers" may overestimate the actual number of workers. Notice that seasonal workers account for 10.35% of total workers in 1996 (p. 43, IBGE, 1996) and 14% in 2006 (SIDRA). The second issue is that the variable "number of workers" in the Agricultural Census does not include employees hired by service provider companies that are contracted by the farm to perform agricultural activities. The Agricultural Census of 2006 was the first to report the number of employees of service providers hired by farms (separately from number of workers), while their number was not reported in the 1996 Census. This means that the only way to construct a definition of agricultural workers that is consistent across the Agricultural Censuses of 1996 and 2006 is to exclude workers contracted by service providers in 2006 (IBGE, 2006, p. 34). Notice that in 2006, 238.825 agricultural establishments contracted service providers to perform tasks on their land: this accounted for 4.61% of all Brazilian establishments in that year (SIDRA). Notice that these two issues do not affect the agricultural employment variables calculated using the Population Census.

All variables come at the municipality level and have been aggregated at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE.[3] All values are reported at current values. In order to convert 1996 and 2006 values to 2000 Reais we used the Índice Geral de Preços - Disponibilidade Interna (IGP-DI) prepared by the Fundaçao Getulio Vargas. We set the index equal to 1 in 2000 and divided the value of production in 1996 and 2006 by the value of the index in those years. First differences are defined between 1996 and 2006 and in all regressions with this variable we winsorize top and bottom 1% of observations of the variable to reduce the importance of extreme values in our estimates.

Labor intensity in agriculture. Data on labor intensity in agriculture come from the Brazilian Agricultural Census (IBGE (1996, 2006)). The variable is defined as the logarithm of total number of workers employed in farms divided by the total area in farms (in hectares). Number of workers in agriculture in 1996 comes from table 321 in 1996 (people employed in farms) and table 956 in 2006 (people employed in farms on 12/31). A detailed discussion of the definition of agricultural workers in the Agricultural Census is reported in the description of the variable ”value of output per worker in agriculture”. Total land in farms comes from table 314 in 1996 and table 787 in 2006 (both reporting the area in farms). All variables come at the municipality level and have been aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE. First differences are defined between 1996 and 2006.

Employment share in agriculture. Data on employment share in agriculture come from Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 1991, 2000 and 2010). The Supplement reports detailed information on employment for a sample of the Brazilian population that is representative of the whole population at the municipal level. The variable is calculated as total number of people

[3]At `http://www.ipeadata.gov.br/iframe\_dicionario.aspx?width=1074\&height=480`

who reported working in the agricultural sector divided by the total number of people that reported being employed in any sector of the economy.[4] Sector of employment is reported according to the first version of the CNAE-Domiciliar in both 2000 and 2010. We define the agricultural sector as any sector with code from 01000 through 09999 (section "A": *agricultura, pecuária, silvicultura e exploração florestal* and section "B": *pesca*). Notice that this definition includes also workers that are employed by firms that provide services related to agriculture (code 01401 in the CNAE Domiciliar). In the year 1991 sector of employment is reported according to the old classification. For these years, we define the agricultural sector as any sector with code from 010 through 049: these include both agriculture and fishery. In every year, the sector of activity refers to the sector of the firm where the worker is employed, even if the worker performs a task that is not directly related to this activity. Employed people are defined as anyone who reported being employed during the reference week, both in permanent and seasonal jobs. People who during the reference week worked for no compensation helping someone else in the household and people who worked in agriculture or fishery for subsistence are also counted as employed. People who suspended work during the reference week for holidays, strike or leave are also considered as employed. Original data come at the individual level: we aggregate these data at the municipality level using the individual weights provided by the IBGE. Data at the municipality level have been in turn aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE. First differences are defined between 2000 and 2010. We do not define first differences between 1991 and 2000 because between these two years the IBGE changed its definition of employment in two ways. First, it started to count zero-income workers as employed. 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. Since these changes made the definition of total employment not homogeneous between 1991 and 2000, we do not compute changes in the employment shares. See text for further details.

[4]Both in 2000 and in 2010 some workers did not declare a valid sector of employment and were classified as working in a specific category: "*setor maldefinido*"(badly defined sector). These were around 1.2% in 2000 and around 6.2% in 2010. We do not count these workers in the denominator of the employment shares. This is the correct treatment if these workers are drawn randomly from the other sectors of the economy.

Employment share in manufacturing. Data on employment share in manufacturing come from Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 2000 and 2010). The Supplement reports detailed information on employment for a sample of the Brazilian population that is representative of the whole population at the municipal level. The variable is calculated as total number of people who reported working in the manufacturing sector divided by the total number of people that reported being employed in any sector of the economy.[5] Sector of employment is reported according to the first version of the CNAE-Domiciliar in both 2000 and 2010. We define the manufacturing sector as any sector with code from 15000 through 37999 (section "D": *indústrias de transformação*). In every year, the sector of activity refers to the sector of the firm where the worker is employed, even if the worker performs a task that is not directly related to this sector. Since 2000, employed people are defined as anyone who reported being employed during the reference week, both in permanent and seasonal jobs. People who during the reference week worked for no compensation helping someone else in the household and people who worked in agriculture or fishery for subsistence are also counted as employed people. People who suspended work during the reference week for holidays, strike or leave are also considered as employed. Original data come at the individual level: we aggregate these data at the municipality level using the individual weights provided by the IBGE. Data at the municipality level have been in turn aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE. First differences are defined between 2000 and 2010.

Employment in manufacturing. Data on employment in manufacturing come from Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 1991, 2000 and 2010). The Supplement reports detailed information on employment for a sample of the Brazilian population that is representative of the whole population at the municipal level. The variable is defined as the logarithm of the total number of people who reported working in the manufacturing sector. Sector of employment is reported according to the first version of the CNAE-Domiciliar in both 2000 and

[5] Both in 2000 and in 2010 some workers did not declare a valid sector of employment and were classified as working in a specific category: "*setor maldefinido*"(badly defined sector). These were around 1.2% in 2000 and around 6.2% in 2010. We do not count these workers in the denominator of the employment shares.

2010. For these years, we define the manufacturing sector as any sector with code from 15000 through 37999 (section "D": *indústrias de transformação*). In the years 1991 and 2000 sector of employment is reported according to the old classfication. For these years, we define the manufacturing sector as any sector with code from 100 through 300. In every year, the sector of activity refers to the sector of the firm where the worker is employed, even if the worker performs an activity that is not directly related to this sector. Employed people are defined as anyone who reported being employed during the reference week. People who suspended work during the reference week for holidays, strike or leave are also considered as employed. Original data come at the individual level: we aggregate these data at the municipality level using the individual weights provided by the IBGE. Data at the municipality level have been in turn aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE. First differences are defined between 2000 and 2010 and between 1991 and 2000 for the pre-trends regression discussed in section 6.2 and presented in table 13.

Wage in manufacturing. Data on the wage in manufacturing come from Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 1991, 2000 and 2010). The Supplement reports detailed information on income from employment for a sample of the Brazilian population that is representative of the whole population at the municipal level. The variable is calculated as the logarithm of the average wage of manufacturing workers 10 years or older in 2000 Reais. See the discussion of employment in manufacturing for details on the definition of the manufacturing sector. The Sample Supplement reports the income from the principal job of every worker. We estimate the average wage in every municipality by looking only to the income from the principal job of workers who report being employees (*empregado com carteira de trabalho assinada* or *empregado sem carteira de trabalho assinada*): this excludes entrepreneurs and self-employed workers, whose income from the principal job may include profits or rents but it includes informal workers.[6] Average wage of manufacturing workers at the municipality level is calculated as the weighted average of the wage of all employees who report working in the manufacturing sector as principal job, using the individual weights provided by the IBGE. Average wage at the level of AMC

[6]Income of entrepreneurs and self-employed workers is reported to the Census as business revenues minus business expenses.

(micro-region) is calculated as a weighted average of the income in all municipalities belonging to an AMC (micro-region) using the number of manufacturing employees in every municipality as weight. All wages are reported in current values of the currency in circulation in the year of the Census: these are thousand Cruzeiros in 1991 and Reais in 2000 and 2010. In order to convert 1991 values to Reais we divided wages in thousand Cruzeiros by 2.75 million. In order to deflate all current values to July 2000 Reais we used the monthly *Índice Nacional de Preços ao Consumidor* (INPC). We set the index equal to 1 in July 2000 (the reference month for the 2000 Census) and divided the 1991 values in Reais by the INPC index in August 1991 (the reference month for the 1991 Census) and the 2010 values by the INPC index in July 2010 (the reference month for the 2010 Census). First differences are defined between 2000 and 2010 and between 1991 and 2000 for the pre-trends regression discussed in section 6.2 and presented in table 13.

Share of agricultural land cultivated with soy. Data on the share of agricultural land cultivated with soy come from the Brazilian Agricultural Census (IBGE (1996, 2006)). The variable is defined as area reaped with soy divided by total land in farms. Area reaped with soy comes from table 501 in 1996 (*área colhida por produtos das lavouras temporárias e condição do produtor*); and table 1823 in 2006 (*produção, venda, valor da produção e área colhida da lavoura temporária por produtos da lavoura temporária e grupos e classes de atividade*). Total land in farms comes from table 314 in 1996 (*área dos estabelecimentos por grupo de atividade econômica e condição legal das terras*) and table 787 in 2006 (*número de estabelecimentos e área dos estabelecimentos agropecuários, por condição legal do produtor em relação ás terras, sexo do produtor, grupos de atividade econômica e grupos de área total*). All variables come at the municipality level and have been aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE. First differences are defined between 1996 and 2006 and in all regressions with this variable we windsorize top and bottom 1% of observations of the variable to reduce the importance of extreme values in our estimates.

Share of agricultural land cultivated with maize. Data on the share of agricultural land cultivated with maize come from the Brazilian Agricultural Census (IBGE (1996, 2006)). The variable is defined as area reaped with maize divided by total land in farms. Area reaped with maize comes from table 501 in 1996 (*área colhida por produtos das lavouras temporárias e condição do produtor*); and table 1823 in 2006 (*produção, venda, valor da produção e área colhida da lavoura temporária por produtos da lavoura temporária e grupos e classes de atividade*). Total land in farms comes from table 314 in 1996 (*área dos estabelecimentos por grupo de atividade econômica e condição legal das terras*) and table 787 in 2006 (*número de estabelecimentos e área dos estabelecimentos agropecuários, por condição legal do produtor em relação ás terras, sexo do produtor, grupos de atividade econômica e grupos de área total*). All variables come at the municipality level and have been aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE. First differences are defined between 1996 and 2006 and in all regressions with this variable we windsorize top and bottom 1% of observations of the variable to reduce the importance of extreme values in our estimates.

Share of agricultural land cultivated with genetically engineered soy. Data on the share of agricultural land cultivated with genetically engineered (GE) soy come from the Brazilian Agricultural Census (IBGE (1996, 2006)). The variable is defined as area reaped with GE soy divided by total land in farms. Area reaped with GE soy comes from table 824 in 2006 (*produção, venda, valor da produção e área colhida da lavoura temporária por produtos da lavoura temporária, tipo de semente, tipo de colheita, tipo de cultivo e destino da produção*) and it is assumed to be 0 in 1996. Total land in farms comes from table 314 in 1996 (*área dos estabelecimentos por grupo de atividade econômica e condição legal das terras*) and table 787 in 2006 (*número de estabelecimentos e área dos estabelecimentos agropecuários, por condição legal do produtor em relação ás terras, sexo do produtor, grupos de atividade econômica e grupos de área total*). All variables come at the municipality level and have been aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE. First differences are defined between 1996 and 2006 and in all regressions with this variable we windsorize top and bottom 1% of observations of the variable to reduce the importance of extreme values in our estimates.

Share of agricultural land cultivated with traditional soy. Data on the share of agricultural land cultivated with traditional (non-GE) soy come from the Brazilian Agricultural Census (IBGE (1996, 2006)). The variable is defined as area reaped with soy minus the area reaped with GE soy divided by total land in farms. Area reaped with soy comes from table 501 in 1996 (*área colhida por produtos das lavouras temporárias e condição do produtor*); and table 1823 in 2006 (*produção, venda, valor da produção e área colhida da lavoura temporária por produtos da lavoura temporária e grupos e classes de atividade*). Area reaped with GE soy comes from table 824 in 2006 (*produção, venda, valor da produção e área colhida da lavoura temporária por produtos da lavoura temporária, tipo de semente, tipo de colheita, tipo de cultivo e destino da produção*) and it is assumed to be 0 in 1996. Total land in farms comes from table 314 in 1996 (área *dos estabelecimentos por grupo de atividade econômica e condição legal das terras*) and table 787 in 2006 (*número de estabelecimentos e área dos estabelecimentos agropecuários, por condição legal do produtor em relação ás terras, sexo do produtor, grupos de atividade econômica e grupos de área total*). All variables come at the municipality level and have been aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE. First differences are defined between 1996 and 2006 and in all regressions with this variable we windsorize top and bottom 1% of observations of the variable to reduce the importance of extreme values in our estimates.

Change in potential soy yield. Data on the change in potential soy yield come from the FAO GAEZ v3.0 database.[7] To construct the change in potential soy yield we use two variables from the *Suitability and Potential Yield* Series: *Total production capacity for low input level rain-fed soybean* and *Total production capacity for high input level rain-fed soybean*. These series are aimed at capturing the *potential production capacity in terms of output density which equals total grid cell production potential divided by grid cell area*. Variables are expressed in tons per hectare. The time is set to the baseline 1961-1990, which means that the FAO-GAEZ agricultural model has been applied considering the average climate of the period 1961-1990. To construct the change in potential soy yield we subtract the low input level variable from the high input level variable in each municipality. Data has then been aggregated at the at AMC-level using the correspondence proposed by IPEA and IBGE.

Change in potential maize yield. Data on the change in potential maize yield come

[7] All data can be downloaded from www.gaez.fao.org.

from the FAO GAEZ v3.0 database.[8] To construct the change in potential maize yield we use two variables from the *Suitability and Potential Yield* Series: *Total production capacity for low input level rain-fed maize* and *Total production capacity for high input level rain-fed maize*. These series are aimed at capturing the *potential production capacity in terms of output density which equals total grid cell production potential divided by grid cell area*. Variables are expressed in tons per hectare. The time is set to the baseline 1961-1990, which means that the FAO-GAEZ agricultural model has been applied considering the average climate of the period 1961-1990. To construct the change in potential maize yield we subtract the low input level variable from the high input level variable in each municipality. Data has then been aggregated at the at AMC-level using the correspondence proposed by IPEA and IBGE.

Share of rural population. Data on the share of rural population come from Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 1980 and 1991). The Supplement reports information on the area where people live (rural or urban) for a sample of the Brazilian population that is representative of the whole population at the municipal level. The variable is calculated considering only people 10 years or older, as the total number of people living in rural areas divided by the total number of people. The type of area where a person lives is classified under four categories in 1980 (two categories for cities and villages, and two for rural areas, depending on density) and under eight categories in 1991 (three categories for cities and villages, and four categories for rural areas). In both years IBGE defines the boundaries of the cities using the most recent municipal law on the matter. We define rural population as everybody who does not live within the boundary of a city or village. Original data come at the individual level: we aggregate these data at the municipality level using the individual weights provided by the IBGE. Data at the municipality level have been in turn aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE.

Income per capita. Data on income per capita come from Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 1980, 1991, 2000 and 2010). The Supplement reports information on income received from any source for a sample of the Brazilian population that is representative of the whole population

[8]All data can be downloaded from www.gaez.fao.org.

at the municipal level. The variable is calculated as the natural logarithm of the average real income of people between 10 and 60 years old earning strictly positive income. The IBGE collects data on income coming from both labor and other sources (including pensions, social programs, rents, capital income etc). We defined income for every person as the sum of income coming from all sources. In order to avoid our results being affected by pensions, we restrict our sample to positive income earners between 10 and 60 years old. Original data come at the individual level: we compute average income at the municipality level using the individual weights provided by the IBGE. Average income at the level of AMC is calculated as a weighted average of the income in all municipalities belonging to an AMC using the number of income earners in every municipality as weight. We matched municipalities to AMC and micro-regions using the correspondence proposed by IPEA and IBGE. All incomes are reported in current values of the currency in circulation in the year of the Census: these are Cruzeiros in 1980, thousand Cruzeiros in 1991, Reais in 2000 and 2010. In order to convert values to Reais, we divided income in Cruzeiros by 2.75 billion in 1980; and we divided income in thousand Cruzeiros by 2.75 million in 1991. In order to deflate all current values to July 2000 Reais we used the monthly *Índice Nacional de Preços ao Consumidor* (INPC). We set the index equal to 1 in July 2000 (the reference month for the 2000 Census) and divided the 1980 and 1991 values in Reais by the INPC index in August 1980 and August 1991 respectively (these were the reference months for the 1980 and 1991 Censuses).

Population density. Data on population density are constructed from the Brazilian Population Censuses (IBGE, 1980 and 1991) and geo-referenced maps of Brazil prepared by GADM (`http://www.gadm.org/`) using data from IBGE. The Brazilian Censuses report information on the total number of people of any age living in a municipality. Geo-referenced on Brazilian municipalities has information on the area of each municipality. The variable is calculated as the natural logarithm of the total number of people living in a municipality divided by the area of the municipality in square kilometers. Data at the municipality level have been in turn aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE.

Literacy rate. Data on the literacy rate come from Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 1980 and 1991). The

Supplement reports information on literacy for a sample of the Brazilian population that is representative of the whole population at the municipal level. The variable is calculated considering only people 10 years or older, as the total number of people who is able to read and write divided by the total number of people. In 1980 IBGE classified people who used to be able to read and write but were currently unable to do so in a separate category: we do not consider these people as literate. Original data come at the individual level: we aggregate these data at the municipality level using the individual weights provided by the IBGE. Data at the municipality level have been in turn aggregated at the at the level of AMC and micro-region using the correspondence proposed by IPEA and IBGE.

Wage in agriculture. Data on the wage in agriculture in 1991 come from Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 1991). The Supplement reports detailed information on income from employment for a sample of the Brazilian population that is representative of the whole population at the municipal level. The variable is calculated as the logarithm of the average wage of agricultural workers 10 years or older in 2000 Reais. Sector of employment is reported according to the old classfication. We define the agricultural sector as any sector with code from 010 through 049. See the discussion of employment in agriculture for details on the definition of the agricultural sector in 1991. The Sample Supplement reports the income from the principal job of every worker. We estimate the average wage in every municipality by looking only to the income from the principal job of workers who report being employees (*empregado com carteira de trabalho assinada* or *empregado sem carteira de trabalho assinada*). This definition of workers excludes entrepreneurs and self-employed workers, whose income from the principal job may include profits or rents.[9] The definition however includes all informal workers. Average wage of agricultural workers at the municipality level is calculated as the weighted average of the wage of all employees who report working in the agricultural sector as principal job, using the individual weights provided by the IBGE. Average wage at the level of AMC (micro-region) is calculated as a weighted average of the income in all municipalities belonging to an AMC (micro-region) using the number of agricultural employees in

[9] Income of entrepreneurs and self-employed workers is reported to the Census as business revenues minus business expenses.

every municipality as weight. All wages are reported in current values of the currency in circulation in the year of the Census: these are thousand Cruzeiros in 1991: see the discussion of income per capita for details on the way we deflate 1991 wages to 2000 Reais.

Employment in manufacturing firms with more than 30 employees. Data on employment in manufacturing firms with more than 30 employees come from PIA, *Pesquisa Industrial Anual*. The data is collected yearly by the IBGE and made available for research purposes in the facilities of the IBGE - Rio de Janeiro upon approval of a research proposal. We restrict our analysis to the manufacturing sector as defined by CNAE 1.0 (code 15 to 37) and use micro-data from 1996 to 2006. We construct our measure of employment starting from variable $V0194$, which is defined in the original documentation as: *Total pessoal ocupado em 31/12* or end-of-year number of workers. We sum this variable across all plants located in the same municipality in each year. In order to have a representative sample at municipality level we use only plants that are part of firms with 30 or more employees (which are sampled with probability one).

Average wage in manufacturing firms with more than 30 employees. Data on wages in manufacturing firms with more than 30 employees come from PIA, *Pesquisa Industrial Anual*. The data is collected yearly by the IBGE and made available for research purposes in the facilities of the IBGE - Rio de Janeiro upon approval of a research proposal. We restrict our analysis to the manufacturing sector as defined by CNAE 1.0 (code 15 to 37) and use micro-data from 1996 to 2006. We construct our measure of average wage starting from variable $V0195$, which is defined in the original documentation as: *Salários, retiradas e outras remuneraçãoes* and includes total wage remuneration as well as pension contributions and other worker benefits. The average wage is calculated summing this variable across all plants located in the same municipality in each year and dividing it by the total number of workers in the same plants in the same year. In order to have a representative sample at municipality level we use only plants that are part of firms with 30 or more employees (which are sampled with probability one).

Migration rate. The migration rate is estimated with data from the Sample Supplement of the Brazilian Population Censuses (*questionário da amostra*: IBGE, 1991, 2000 and 2010). To construct net migration in a municipality between two Census years we follow the standard cohort average method (Shryock et al., 1980, pp. 630-635). Letting the first Census year be year 0, and the second Census year be year t, this implies computing, for every age $a + t$, the total number of people living in each municipality m in year t, and subtract from this number the total number of people of that age that one would expect given the population of age a living in that municipality in year 0 and given the average survival rate for that age group in Brazil.

Formally, this implies computing the following formula:

$$M_{m,a+t} = \frac{M'_{m,a+t} + M''_{m,a+t}}{2}$$

where $M'_{m,a+t}$ and $M''_{m,a+t}$ are net migrants in municipality m computed with the "forward survival rate method" and with the "reverse survival rate method", respectively:

$$\begin{aligned} M'_{m,a+t} &= P^t_{m,a+t} - s^t_a P^0_{m,a} \\ M''_{m,a+t} &= \frac{P^t_{m,a+t}}{s^t_a} - P^0_{m,a} \end{aligned}$$

In these equations, $P^t_{m,a+t}$ is the total number of people of age $a+t$ living in municipality m in year t and $P^0_{m,a}$ is the total number of people of age a observed in the same municipality in year 0. s^t_a is the probability that a person of age a in year 0 is still alive after t years. Both $P^0_{m,a}$ and $P^t_{m,a+t}$ are observed in the Brazilian Censuses of 1991, 2000 and 2010. Survival probabilities are estimated using data on the whole Brazilian population with the formula:

$$s^t_a = \frac{P^t_{a+t}}{P^0_a}$$

where $P^0_a = \sum_m P^0_{m,a}$ and $P^t_{a+t} = \sum_m P^t_{m,a+t}$. Notice that with this formula net migration for Brazil as a whole is zero by construction.

Let M^t_{a+t} be the estimate of net number of migrants in a municipality with age $a+t$ in Census year t: total net migration M^t is the sum of net migrants with age between 10 and 60 years in the initial period[10]. We compute migration rates MR^t_m as net migrants

[10]Since the last three Censuses were taken in 1991, 2000 and 2010, we can estimate migration between 1991 and 2000 and between 2000 and 2010. In the first case migrants are aged between 19 and 69 years in 2000, while for the migration rate between 2000 and 2010 they are aged between 20 and 70 in 2010.

in a municipality divided by the total number of people living in that municipality in year 0:

$$MR_m^t = \frac{M_m^t}{P_m^0}$$

The cohort average method implicitly imposes two assumptions that are worth discussing briefly. First, the method assumes that Brazil is a closed country (i.e. nobody emigrate or immigrate from abroad between two Census years). Second, the method assumes that in a given year everybody has the the same survival probability conditional on age. We discuss these assumptions in turn.

We deal with the closed country assumption by computing net migration only of people born in Brazil.[11] In other words, in every Census year, we compute $P_{m,a}^0$, $P_{m,a+t}^t$ and s_a^t using only people that report being born in Brazil, which means that our migration rates refer to Brazilian-born only. We do this because people born outside of Brazil are on average more likely to enter or leave the country between two Census years. Note however, that this decision creates three additional issues. First, it requires to interpret the survival probability as the probability of re-observing a person somewhere in Brazil after t years, rather than the probability of this person of simply being alive after t years. Second, it implicitly assumes that nobody born in Brazil emigrated outside of Brazil before year 0 and then returned between the year 0 and the year t. Finally, it entirely abstracts from migration of foreigners (both within the country and from abroad). The first of these three issues is inconsequential, but the other two may bias our estimates if Brazil as a whole experienced large migration flows between 1991 and 2010, or if foreign-born people accounted for a large share of the population. In practice this does not seem to be the case. In 2010, only 0.19% of the Brazilian-born population aged 15 to 70 reported having moved to Brazil from abroad during the previous 10 years. Moreover, people born outside of Brazil always account for less than 1% of the total population living in Brazil between 1991 and 2010 (the exact figures are: 0.52% in 1991, 0.4% in 2000 and 0.31% in 2010).

In order to allow different survival probability across different groups of the population we compute net migration and survival probabilities separately for men and women

[11]This is what Shryock et al. (1980) recommend. The standard alternative is to compute the survival probabilities only of Brazilian-born and then apply these probabilities to both nationals and foreigners.

and then add them to compute total migration.

Chapter 2

Turnover taxes and productivity. Evidence from a Brazilian tax reform. *

2.1 Introduction

Most production processes are sequential and involve a chain of firms that exchange inputs. In such production processes, distortions to trade between firms along the production chain can have large effects on aggregate productivity because production losses may cumulate along the production chain. Models with production chains have therefore recently been used to reexamine the aggregate implications of firm-level distortions in growth and development (e.g. Jones (2011)), business cycle economics (e.g. Acemoglu et al. (2012)) and international trade (e.g Antràs and Chor (2013)). In this paper we use a model with production chains for a quantitative study of the productivity and welfare losses of distortive taxation. In particular we are interested in quantifying the productivity and welfare gains of a tax reform that eliminates distortive turnover taxes, which are still used in many developing countries (according to the World Bank, 45 countries were still using turnover taxes in 2012).

Turnover taxes are taxes levied on the full value of firms' revenues (or turnover). Because turnover taxes do not allow firms to deduct the cost of intermediate inputs from the tax base, they raise the cost of purchasing intermediate inputs from other firms relative to producing inputs in-house. This is well-known to reduce the aggregate efficiency

*Joint with Antonio Ciccone.

of production (Diamond and Mirrlees (1971a)) which is why international economic institutions have been recommending that turnover taxes should be abolished in favor of less distortive taxes (e.g. Keen (2009)). Our objective is to quantify the productivity and welfare effects of substituting a turnover tax with a non-distortionary tax when there are production chains. To do so we develop a model with production chains and calibrate it using information on input and output flows reported in input-output tables. Two other aspects of the model we develop are important. First, the model can accommodate production chains of any length. This allows us to quantify the productivity and welfare effects of turnover taxes for production chains of different length. Second, the model reduces to standard models of input-output linkages used by Jones (2011, 2013) and Antràs et al. (2012) when production chains become infinitely long. Thus, the model we develop is easy to relate to standard models of input-output linkages.

One important issue that arises when calibrating a model with production chains using input-output tables is that such tables do not contain enough information to calibrate both the intermediate input intensity and the number of stages of production. This point can be seen in a simple example. Suppose that we know that there are two firms in a sector and that firms in the sector have bought a total value of 10 million Euros of goods from other firms in the sector. This data is consistent with firm A buying 10 million Euro of goods from firm B. In this case goods are exchanged between firms once. But the data is also consistent with firm A buying 2.5 million Euro of a good produced by firm B, adding 5 million of value through processing and then selling 7.5 million Euro of the processed good to firm B. In this second case goods are exchanged between firms twice. Because the total value of transactions is the same in the two cases, knowing the value of goods exchanged across firms is not enough to know the number of stages of production. This is an important limitation of the data in input-output tables, as generally the productivity and welfare effects of turnover taxes will depend on the intermediate input intensity of production as well as the number of stages of production. We deal with this limitation of the data in input-output tables in the following way. We first set the number of stages of production in the model as an exogenous parameter. Next, we calibrate the remaining parameters to match the data in input-output tables. This allows us to derive the production gains of eliminating a tax on turnover as a function of the number of stages of production.

We calibrate the model to Brazilian input-output tables and use it to examine the productivity and welfare effects of the conversion of a turnover tax of 3.65% into a value added tax — a tax reform that Brazil actually undertook in 2002. A main conclusion of our analysis is that such gains are much larger in the case with long production chains than the case with short production chains. For example, we find that the reform results in productivity gains that are 4.1 times larger when the production process has more than 11 stages of production than when it has a single stage of production. However, for the intermediate input intensities observed in Brazil, the productivity gains of such a tax reform remain small even when production chains are infinitely long. For example, eliminating the tax on turnover increases national income by at most 0.05% in the baseline case. For the productivity effects of the Brazilian turnover tax to become large, intermediate input intensities would have to be substantially larger than what we observe from the input-output tables. For example, with infinitely long production chains, the productivity gain of the tax reform would be around 1% (0.95%) of national income if the overall intermediate input intensity of the Brazilian economy – the value of inputs over gross output – was 70% instead of the observed 43%. And the productivity gain of the tax reform would be around 1.33% of value added if the overall intermediate input intensity of the Brazilian economy was 95%.

Our baseline calibration for Brazil is based on an input-output matrix with 100 sectors. Because such a detailed input-output matrix will be unavailable for many developing countries, we also ask what productivity effect of the 2002 tax reform we would have found if our calibration had treated Brazil as a single-sector economy (with an overall intermediate inputs intensity of production of 43%). Somewhat surprisingly, we find almost the same effects as in the case of the input-output matrix with 100 sectors. Hence, the number of production stages appears to matter more for the quantitative results of the tax reform than the detail of sectoral disaggregation.

Our main results are obtained assuming an elasticity of substitution between intermediate inputs and labor equal to one . For robustness we examine how results depend on the elasticity of substitution. We find lower productivity gains brought by the reform when firms have a very low or very high elasticity of substitution between intermediate inputs and labor. The productivity effects of turnover taxes are low when the elasticity of substitution between intermediate inputs and labor is low because firms' technologies are such that they cannot alter their production decisions in response to the tax. In the opposite case when the elasticity of substitution between labor and inputs is very high, firms respond to the tax by changing significantly their input mix. However, because labor is a good substitute of intermediate inputs in this scenario, the productivity distortions of producing in-house rather than buying from other firms will be small. With numerical simulations we show that the elasticity that yields the highest welfare gain of the Brazilian tax reform depends on the number of stages of production. For example, we find that the largest welfare gains are obtained for an elasticity of substitution equal to 50 when the production process involves 3 stages of production. In this scenario the welfare gains of the reform would be slighly larger than 1% of national income. When the production process requires an infinitely long chain instead, the largest welfare gain of the reform are obtained for an elasticity of substitution equal to 37, when the welfare gains of the reform are equal to 0.87% of national income.

The results of the paper have implications for tax design, especially in developing countries. Taxes on turnover have been long known to reduce productivity (Diamond and Mirrlees (1971a,b)). However, it has also been argued that this type of taxes can be more desirable than other tax schemes. Taxes on turnover do not allow to deduct the cost of intermediate inputs from the tax base. Thus, firms cannot reduce their taxes by inflating the cost of inputs and for this reason they are harder to evade than value added taxes or taxes on profits. Best et al. (2013) present evidence consistent with this argument, by showing a large positive effect of turnover taxes on tax compliance in Pakistan, where firms are taxed on either turnover or profits, depending on which tax liability is larger.[2,3] The large informal sectors observed in developing countries, combined with the lower evasion guaranteed by turnover taxes may explain why these taxes are still very popular among countries with limited tax capacity (Besley and Persson (2013)). As of 2012, 45 low and middle income countries had some form of tax on turnover in place (World

[2]The presence of an informal sector also creates theoretical reasons for the optimality of input taxation. When intermediate inputs are produced by informal firms, taxing inputs of production is optimal because it allows to raise revenue on at least part of the product of the informal sector (Newbery (1986); Emran and Stiglitz (2005)).

[3]There are also reasons to believe that a well enforced value added tax increases tax compliance. The general consensus is that the credit method used to collect value added taxes create "chains of firms" that exchange inputs of production and that are either all formal or all informal (Keen (2009); De Paula and Scheinkman (2010)). The creation of such production chains may increase the effectiveness of tax audits, as they create incentives for audited firms to demand receipts from their suppliers (Pomeranz (2013)). However, this implies that value added taxes may raise the returns of a given level of tax enforcement, but there is nothing in the literature that suggests that VAT increase tax compliance *per se*.

Bank (2006, 2012)), despite the fact that the conversion of taxes on turnover into modern value added taxes is part of the standard reform package promoted by the International Monetary Fund and the World Bank (Keen (2009)).

Related literature

The paper contributes to several literatures. In public finance, the evaluation of the social costs of taxes was pioneered by Harberger (1962). Following his tradition, several authors have evaluated the efficiency costs of taxes that cumulate along the production chain (Bovenberg (1987); Gottfried and Wiegard (1991); Piggott and Whalley (2001)). These works evaluate the welfare loss caused by the inefficient design of value added taxes, and they do not analyze the welfare cost of input taxation that affects all the sectors of the economy. The closest precedent to our result on the welfare gain of eliminating a turnover taxes is Keen (2013), who analyzes production inefficiencies generated by turnover taxes. He derives an approximation of the deadweight loss of a turnover tax in two special cases. First, when there is a single stage of production and arbitrary elasticity of substitution between inputs. Second, when there is an arbitrary number of stages but every stage before the last one produces output according to a Leontief production function. In both cases, the turnover tax distorts decisions of a single stage of production, when intermediate inputs can be substituted with labor. Relative to his work, we assume that firms produce with either a Cobb-Douglas, or a general constant elasticity of substitution technology that combines labor and intermediate inputs. These assumptions allow us to study processes in which turnover taxes distort production decisions several times along a production chain, rather that a single time at the end of it. The model is flexible enough to analyze formally the interaction between the length of the production chain and the elasticity of substitution between intermediate inputs and labor in determining the welfare cost of turnover taxes.

In the literature of input-output economics several authors have pointed out that input-output linkages can magnify distortions or changes in productivity that happen in few sectors (Leontief (1936); Hirschman (1958); Hulten (1978); Long and Plosser (1983); Basu (1995) and more recently Ciccone (2002); Gabaix (2011); Jones (2011); Oberfield (2011)).[4] Similar ideas have been developed in the international trade liter-

[4]Bartelme and Gorodnichenko (2014) are to our knowledge the only authors who explored these ideas empirically. They use a panel of input-output tables from several countries to show that the process of economic growth is associated with an increase in the importance of input-output linkages.

ature, that has shown how the effect of tariffs cumulates along the production chain in the same way as a tax on inputs does (Corden (1966); Yi (2003); Fally (2012); Antràs et al. (2012); Baldwin and Venables (2013)).

Finally, some of the results we present are relevant for the literature on misallocation. In particular, Jones (2011) shows that firm-level TFP may be amplified by input-output linkages and can have very large aggregate output effects. As Jones, we find that a tax on turnover distorts the way resources are allocated, and its effect is magnified by the input-output structure of the economy. However, we also find that the deadweight loss implied by the turnover tax is relatively small. The source of the difference between our estimates and those of Jones, is the consideration of tax revenues. In his paper Jones considers wedges that reduce the productivity of different sectors and are entirely wasteful.[5] When the distortion to relative prices comes from a tax however, part of the increased cost of production is collected by the fiscal authority in the form of tax revenue, and should be accounted in welfare calculations.

The paper proceeds as follows. Section 2.2 develops the main ideas with a simple example. Section 2.3 presents background on the 2002-3 Brazilian tax reform. Section 2.4 derives a general model that delivers explicit expressions for the welfare gain of the tax reform. Section 2.5 presents an extension of the model where we relax the Cobb-Douglas assumption and assume that production takes place with a generic constant elasticity of substitution function. Section 2.6 concludes. Explicit derivation of analytical results are in the appendix.

2.2 An example

We now illustrate the main idea of the paper. Consider an economy where firms produce goods using labor and goods that are produced by other firms (intermediate inputs). Suppose that a tax on the full value of the good is applied every time a firm buys or sells a good to another firm or to final consumers. To estimate the deadweight loss of such

[5]In this sense the way sectoral distortions show up in aggregate TFP in Jones (2011) differs to the way distortions to the prices faced by individual firms reduce aggregate production in models of heterogeneous firms such as Banerjee and Duflo (2005), Restuccia and Rogerson (2008) and Hsieh and Klenow (2009). The easiest way to see the difference is to assume that symmetric (i.e. identical) positive wedges throughout the economy. In models of misallocations of factors across firms symmetric wedges create no output loss. In Jones (2011) a positive wedge create output loss, even when it is the same for every sector.

a tax one would have to observe a detailed input-output table for each good and firm. This table would have to report the value of production for each good and firm as well as the value of each input, the firms from which inputs are bought and the good for which inputs are used. With such data, plus the price of each good, one could keep track of every transaction and evaluate the deadweight loss of the tax. In practice, such detailed information is never available. What is usually available however is a sectoral input-output table. In this paper we therefore investigate whether such input-output tables can be used to evaluate the deadweight loss of the tax.

The main limitation of input-output tables when evaluating the deadweight loss of turnover taxes is that they specify the value of goods exchanged within and across sectors, but contain no direct information on the number of times goods are exchanged. Moreover, it is impossible to use the information in input-output tables to calibrate both the intermediate input intensities of production and the number of stages of production. This can be seen by considering the situation of the two firms exchanging 10 million Euro of goods discussed in the introduction. In that case knowing the value of the goods exchanged was not enough to determine the number of transaction between them. Although this information can never be recovered from input-output tables, it matters for the welfare gains of a reform that eliminates a tax on turnover, because these are generally higher when firms exchange inputs more often. We now illustrate this in a simple (calibrated) example.

The example is a one sector model. Hence, the corresponding input-output table would only tell us the total value of goods produced by firms in the sector and the total value of goods that firms in the sector bought from other firms in the sector. Suppose that within the sector firms produce different varieties that are either used as intermediate inputs or consumed. We assume the simplest possible structure and consider two versions of the example that differ in the production structure. In the first version, the economy produces two varieties and in the second version the economy produces three varieties. In both cases production is organized sequentially. To emphasize this point, we refer to the number of different varieties as to “stages of production”. The objective of the example is to illustrate that for a given observed input-output table, the production parameters and the deadweight loss of a turnover tax depend on the number of stages of production.

Two production stages. In this version the economy produces two varieties of the same good, variety 1 and variety 2. Variety 1 is produced by firm 1 with labor only and it is entirely sold to firm 2, which uses it to produce variety 2. Variety 2 is then sold to households as consumption good.

Three production stages. In the version with three production stages the economy produces three varieties, variety 1, variety 2, and variety 3. Variety 1 is produced by firm 1 with labor only and sold as an intermediate input to firm 2 to produce variety 2. Variety 2 is then sold to firm 3 which uses it to produce variety 3. Variety 3 is sold to households as consumption good.

We assume that firms operate under perfect competition and that all production that combines labor and intermediate inputs takes place with Cobb-Douglas production functions. The two different production structures and the production functions are described in figure 2.1. It can be seen that the example with three stages of production simply adds a stage of production to the example with two stages of production. An important aspect of this setup is that the only parameter that requires calibration in both cases is the intermediate input intensity of production (which is called ϕ in the first case and ψ in the second).

Preferences. In both versions of the example, a representative household suppies labor inelastically and consumes a privately produced good C and a public service supplied by the government. Her utility is:

$$U = C \cdot G^{\alpha}$$

where G is a public service supplied by the government and C is the quantity of the privately produced consumption good.

Government. The government collects taxes and uses them to hire workers and produce G units of the public service, which is produced using labor only.

The government can choose to raise taxes in two ways. It can tax the total revenues of every firm using a turnover tax of rate t. This tax makes labor cheaper than intermediate inputs for all firms that use both factors in production and therefore it generally distorts production. Alternatively, the government can impose a tax of rate v on the consumption of the household. This tax does not affect the price of labor relative to intermediate inputs and therefore it does not distort production (in our model this tax does not distort any economic decision and is therefore equivalent to a lump-sum tax).[6]

Input-output tables. The information in input-output tables does not allow us to see whether the production structure corresponds to the model with two or three stages of production. However, it does allow us to calibrate the intermediate inputs intensity of production in the two cases. To see this notice that the total value of production in the economy (V) and the value of intermediate inputs used (X) are in the two cases respectively:

$$\begin{aligned} X &= P(1)Q(1) \\ V &= P(1)Q(1) + P(2)Q(2) \end{aligned} \quad \text{and} \quad \begin{aligned} X &= P(1)Q(1) + P(2)Q(2) \\ V &= P(1)Q(1) + P(2)Q(2) + P(3)Q(3) \end{aligned}$$

The assumption of Cobb-Douglas production allows to solve for ϕ in the two stages case, and for ψ in the three stages case. These are the solutions of the following two equations:

$$\frac{X}{V} = \frac{\phi}{1+\phi} \qquad \text{and} \qquad \frac{X}{V} = \frac{\psi + \psi^2}{1 + \psi + \psi^2}$$

Tax reform. Once we have used the input-output tables to calibrate the intermediate input intensity of production, we can evaluate the welfare gains of a reform that substitutes a turnover tax with a consumption tax. Here as in the rest of the paper we consider a tax reform that does not affect the level of tax revenues. Thus, the tax rate of the consumption tax is set to the level that guarantees the level of revenues needed to finance the level of public service provided before the reform.

Taking the wage to be the numeraire ($W = 1$), the inelastic supply of labor implies that welfare can always be written as $U = \left(\overline{L}/P_c\right) \cdot G^{\alpha}$, where $\overline{L}$ is the total amount of labor the household supplies, P_c is the price paid by households for one unit of the consumption good. Because we maintain the provision of public service fixed, in the example the welfare change is equal to the log change of the price of the consumption

[6]From the point of view of production, a consumption tax is equivalent to a tax on value added. We study the welfare gains of a reform that converts a tax on turnover into a consumption tax because the latter is easier to work with than a value added tax.

good. Thus, when the production process involves S separate stages of production, the welfare gains of the reform is:

$$\Delta \log U = \log \left(\frac{P^{tt}(S)}{P^{v}(S)} \right)$$

where $P^{tt}(S)$ and $P^{v}(S)$ are respectively the prices paid by the household for one unit of consumption good when there is a turnover tax and the price of the same good when there is a consumption tax.

As stated before, here as in the rest of the paper we keep tax revenues fixed to a given value in every tax scheme. Thus, we obtain the tax rate of the turnover tax t (consumption tax v) as the tax rate that raises the tax revenue we observe. In the example we assume that the government needs to raise an amount of taxes equal to 4.6% of GDP. This is what turnover taxes raised as a share of Brazilian GDP in the two years before the tax reform. We then set the "effective" turnover tax in our example equal to the value that ensures these tax revenues. Formally this means that the tax rates in the example with two production stages and in the example with three production stages solve:

$$0.046 = \frac{t(2)}{1+t(2)} \left(1 + \frac{\phi}{1+t(2)} \right)$$
$$0.046 = \frac{t(3)}{1+t(3)} \left[1 + \frac{\psi}{1+t(3)} + \left(\frac{\psi}{1+t(3)} \right)^2 \right]$$

Where $t(2)$ is the effective tax rate in the economy with two stages of production and $t(3)$ is the effective tax rate in the economy with three stages. Because $\phi = \psi + \psi^2 \equiv x$, the two equations can be written as:

$$\begin{aligned} 0.046 &= \frac{t(2)}{1+t(2)} \left(1 + \frac{x}{1+t(2)} \right) \\ 0.046 &= \frac{t(3)}{1+t(3)} \left(1 + \frac{x}{1+t(3)} \right) - \frac{1}{1+t(3)} \left(\frac{t(3)\psi}{1+t(3)} \right)^2 \end{aligned}$$

Since the last term of the second equation is negative, it follows that $t(3) > t(2)$ which means that the effective turnover tax must be higher in the economy with 3 stages than in the economy with 2 stages. As a result, tax reform yields a larger welfare gains in the economy with 3 stages of production, because in this case the effective turnover tax that the reform abolishes is higher. This is illustrated in figure 2.2, which plots the welfare gain of the reform for different observed input share X/V in the version of the example that involves three stages of production (in green) and in the version of the example that involves two stages of production (in blue). Comparing the two gains for $X/V = 43\%$ (the observed input share for Brazil in 2005, 2 years after the reform) one can see that the calibrated welfare gain of the reform is 3.72 times larger in the production process with three stages than in the production process with only two stages.

Infinite number of production stages. It is instructive to look at the case in which the final good is produced after an infinitely long production chain. There are two reasons to discuss this special case. First, because the welfare gain of the tax reform grows with the number of stages of production, this case provides an upper bound for the gains in a model with a single final good. Second, when the number of stages of production grows to infinity, the model with a single final good becomes identical to a model in which a single sector produces the final good using some of its own output as input of production. This is a type of model that is close to standard models used in the growth literature (e.g. Jones (2011)), and it is interesting to see how the model we develop here relates to these standard models.

When there are an infinite number of stages of production, the contribution to total output of the initial stage of production becomes negligible. Because every stage after the initial one is assumed to produce with an identical Cobb-Douglas function, in this special case the observed input share of the economy equals the input elasticity parameter of the production function. We call this input elasticity $X/V = o$. Solving the model with an infinite number of stages yields the following equations for the welfare gains and the effective tax rate as a function of the input elasticity o, tax revenues T and national income $W\overline{L} = \overline{L}$:

$$\begin{aligned}
\lim_{S\to\infty} \Delta \log U(S) &= \log\left[(1+t(S))^{\frac{1}{1-o}}\left(1-\frac{T}{\overline{L}}\right)\right] \\
\lim_{S\to\infty} t(S) &= \frac{T}{\overline{L}-T}(1-o)
\end{aligned}$$

Figure 2.3 shows the welfare gain of the reform as a function of the observed input share when tax revenues over national income are kept constant at 4.6%. The blue and the green lines show the welfare gains of the reform in the two and three stages economies. The red line represents the welfare gain of the reform when the number of stages of production goes to infinity. The plot shows that the welfare gains are always larger in the economy with an infinitely long production chain. However, even as the observed input share approaches 100%, the welfare gains do not become infinitely large. This result is in stark contrast with what happens when the tax rate is fixed (and equal to the statutory tax rate, say) and tax revenues are allowed to adjust as input share changes. In this case, tax revenues collected on an infinitely long production chain grow without bounds as the input share approaches 100%. For this reason, in this case also the welfare gain of eliminating a turnover tax grows to infinity when the observed input share approaches unity. When tax revenues are kept constant instead, the effective tax rate needed to raise a given amount of tax revenues becomes smaller as the observed input share grows. In the limit, this tax rate approaches zero as the observed input share approaches unity. A lower effective tax rate partly compensate for the larger distortions created by a turnover tax in an economy that produces with a larger input elasticity. As a result, the welfare gain of eliminating such tax does not become indefinitely large as input share tends to one.

Finally, this simple example allows to ask what would be the welfare gain of the reform if the turnover tax was raising a higher share of revenues relative to total income than what we observe for Brazil in 2002. The interest of this exercise lays in the possibility that, although the Brazilian reform brought only modest gains, it was nonetheless a necessary step towards a more efficient tax system that would allow the government to raise more revenues. In figure 2.4 we plot the welfare gains of the reform as a function of the observed input share when the tax revenue raised by the turnover tax before the reform is equal to 10%. This is on the upper end of what developed countries usually raise by means of value added taxes: according to OECD data cited by Keen and Lockwood (2006), in 2005 only Hungary and Iceland were raising more than 10% of GDP with this type of tax. In the graph, we plot the welfare gain for an economy that produces with two stages (in blue), with three stages (in green) and with an infinite number of stages (in red). For the input share observed in Brazil in 2005 (43%), the welfare gains of the reform are around 0.24% of total income even when there is an infinite number of stages. Figure 2.5 repeats the same exercise assuming that the turnover tax that is eliminated was raising revenues equal to 15% of total income, a level that is higher than the level observed anywhere today for VATs. In this case, for the observed input share of Brazil the welfare gain of the reform would be at most 0.56% of national income, when production requires an infinite number of stages.

2.3 The 2002-3 Brazilian tax reform

In this section, we provide background information on the 2002 Brazilian tax reform and few relevant aspects of the Brazilian fiscal system. In Brazil, every business has to pay several different taxes and social contributions.[7] This paper focuses on two of the social contributions paid by firms: the *Programa de Integração Social* (PIS: contribution for the social integration programme) and the *Contribuição para Financiamento da Seguridade Social* (COFINS: contribution for the funding of social security).[8] Until 2002 these contributions were levied on the total value of turnover, with no deductions for the cost of intermediate inputs. The tax rate of PIS was 0.65%, while the tax rate of COFINS was 2% between 1991 and 1998, and 3% since 1998. Except for the different tax rate, until 2002 the two taxes were very similar in terms of taxpayers, tax base and exceptions.

Between December 2002 and December 2003 the Brazilian Federal Congress passed two separate laws that modified the regime of both PIS and COFINS. The first reform, in December 2002, modified the regime of PIS and it allowed to deduct the cost of intermediate inputs from the the tax base, effectively converting PIS into a VAT. The

[7]Brazilian companies have to pay a corporate income tax on their profits (the *Imposto de Renda de Pessoa Jurídica*, IRPJ), and two different VATs on the value of their shipments (the *Imposto sobre Produtos Industrializados*, IPI, paid to the Federal Government, and the *Imposto sobre a Circulação de Mercadorias e Serviços*, ICMS, paid to the State where the firm operates). They are also responsible to pay a contribution for a social integration programme (*Programa de Integração Social*, PIS), a contribution for public servant assets (*Programa de Formação do Patrimônio do Servidor Público*, PASEP), a contribution to finance social security (*Contribuição para Financiamento da Seguridade Social*, COFINS), an union contribution, a retirement contribution for each employee, a contribution to finance education and municipal taxes according to local laws.

[8]PIS was introduced with the complementary law no. 7 on the 7th of September 1970. COFINS was created by the complementary law no. 70 on the 31st of December 1991, and it replaced the existing *Fundo de Investimento Social* (FINSOCIAL).

law also increased the tax rate of PIS by 1 percentage point, to 1.65%, in order to avoid a fall in fiscal revenues. The second reform, in December 2003, modified the regime of COFINS, and it allowed to deduct the cost of intermediate inputs from the tax base of this tax too. The new tax rate of COFINS was set to 7.6%, a number that was chosen by applying to the existing tax rate the same multiplier adopted to adjust the rate of PIS in the VAT regime (roughly 2.53; see Werneck (2006) for a discussion of this choice).[9] In general, the two reforms aimed at maintaining PIS and COFINS similar taxes: the law that modified PIS contained the express provision for the Government to convert COFINS into a VAT within a year, and the COFINS reform is almost identical to the PIS reform in content, structure and wording. For this reasons in this paper we treat the two reforms as a single one.

The stated objective of these reforms was to increase competitiveness of the Brazilian manufacturing sector, and the policies were strongly encouraged by the International Monetary Fund (IMF) as part of the agreements to obtain recovery loans in the aftermath of the capital account crisis of 1998 (IMF (2002) and IMF (2003)).

In the next two sections we develop a formal framework to evaluate the welfare gain of this reform under different assumptions regarding the structure of production.

2.4 Cobb-Douglas production

The example discussed in section 2.2 is intuitive and simple to analyze. However it is special in that it assumes that production can be organized in only two ways. Moreover, it abstracts from input-output linkages between sectors producing different goods, and imposes a special functional form on production. In this section and in the next section, we develop two extensions of the example. In both models, we consider production processes that involve an arbitrary number of stages. In section 2.4.1 we consider an economy with a realistic input-output structure in which several sectors exchange intermediate inputs. In section 2.5 we consider a single-sector economy, but relax the Cobb-Douglas assumption.

[9] For the PIS reform, see law no. 10637 passed on the 30th of December 2002, especially clause 3. For the COFINS reform, see law no. 10833 passed on the 29th of December 2003, especially clause 3.

2.4.1 The model

Setup. In this model, a representative household values the consumption of two types of goods. The first good is a public service that the government finances with taxes and provides to the household free of charge. The second good is an aggregate good made of a bundle of N differentiated products.

Production of public service only requires labor, and the government hires workers with the revenues collected with taxes. Production of the N differentiated products that compose the second good is carried out by private firms working in N separated sectors of the economy. Within each of the N sectors of the economy there are many firms that produce gross output using labor and intermediate goods produced by firms operating in all other sectors of the economy. In the input-output table of this economy we observe how much the firms of each sector buy from other firms, and how much they sell to other firms in the economy. As in the simple example developed in section 2.2 however, these aggregate flows may be generated by firms exchanging goods an arbitrary number of times, and this number will not be observable. As in the single-sector example, we will think of the inputs sold by the same sector at different stages as different varieties of the same good.

The household supplies labor inelastically. In the model, prices in the private sectors are determined by technology and perfect competition among many firms that produce and sell the same varieties. Given homothetic preferences, consumption shares are then uniquely determined by prices. The level of consumption of private goods depends on the total labor supply which determines the income of the household. Given the quantity of goods that the household wants to consume, the quantity produced by every sector is then pinned down by the structure of production.

Households. The representative consumer has the following Cobb-Douglas utility:

$$U = C \cdot G^{\alpha}$$

The arguments of the utility function are G, the amount of public service received from the government, and C, a composite of the N private goods purchased by households. The composite private good is a Cobb-Douglas aggregate of the goods produced in the

N sectors of the economy:

$$C = \prod_{i=1}^{N} C_i^{\beta_i} \tag{2.1}$$

In (2.1), C_i is consumption of the i-th good and $\sum_i \beta_i = 1$.

The household supplies inelastically $\overline{L}$ hours of labor. At the given hourly wage W, she earns a total income of $W\overline{L}$ that she can spend on the consumption of the private goods C_i. Given Cobb-Douglas preferences, consumption of each good i is equal to $\beta_i W\overline{L}$.

Firms. There are N separate sectors where private firms produce differentiated final goods. For each of the N final goods produced and sold to consumers there are S different varieties produced. For every final good i, the first variety is produced using only labor:

$$Q_i(1) = A_i L_i(1)$$

In the production function of firms producing variety "1" in sector i, A_i is a sector-specific Hicks-neutral productivity parameter and $L_i(1)$ is the amount of labor used. The other $S-1$ varieties in the same sector are produced with labor and intermediate inputs according to a sector-specific production function:

$$Q_i(s) = A_i L_i(s)^{1-\omega_i} \left[\prod_{j=1}^{N} M_{ij}(s)^{\omega_{ij}}\right] \qquad s = 2, 3, \ldots, S$$

where $\omega_i = \sum_j \omega_{ij}$. We assume that in every sector i, Hicks-neutral productivity parameter and input elasticities are sector-specific. However we do not allow these parameters to vary across varieties of the same sector. In other words, we assume that $\omega_{ij}(s) = \omega_{ij}(s') = \omega_{ij}$ and $A_i(s) = A_i(s') = A_i$ for all $s \neq s'$.

In order to derive a tractable solution to the model we make the following assumption:

Assumption 1. *The output of variety s of every sector is entirely absorbed by the production of the varieties $s+1$ of the economy. Only the output of varieties "S" is directly consumed. Formally, for every sector i, the following two conditions are assumed to be* *always true:*

$$\begin{aligned} Q_i(s) &= \sum_{j=1}^{N} M_{ji}(s+1) \qquad\qquad s = 1, 2, \ldots, S-1 \\ Q_i(S) &= C_i \end{aligned}$$

Although extreme, this assumption allows to solve the model and to derive a formula for the aggregate tax distortion. Moreover, it allows to order varieties according their position along the production chain, and it justifies calling the different varieties "stages of production".

Firms are owned by the household, but perfect competition and constant return to scale drive their profits to zero. Zero profit imply that prices equal unit costs. Thus prices are:

$$\begin{aligned} P_i(1) &= \frac{W}{A_i} && s = 1 \\ P_i(s) &= \frac{1}{A_i}\left(\frac{W}{1-\omega_i}\right)^{1-\omega_i}\left[\prod_{j=1}^{N}\left(\frac{P_j(s-1)}{\omega_{ij}}\right)^{\omega_{ij}}\right] && s = 2, 3, \ldots, S \end{aligned} \tag{2.2}$$

Given the wage W, (2.2) is a system of $N \times (S+1)$ equations in $N \times (S+1)$ unknown prices. Prices at stage "1" are pinned down by wage and technology alone. Then, all prices from stage "2" through "S" are determined by wage, technology and the prices of the varieties produced by the stage immediately upstream. In appendix A we derive explicitly the solution of the price vector at every stage of production.

Government. The government produces a single service with a constant return to scale technology that requires only labor:

$$G = L_G$$

The government provides to the household an exogenous level of this public service $\overline{G}$. The public good is provided free of charge, and it is financed with tax revenues.

We assume that the government can tax the activities of private firms. It can do so in one of two ways. First, it can impose a turnover tax equal to $t\%$ of the value of revenues earned by every firm in the economy. Because in this economy some firms earn revenues by selling intermediate inputs to other firms, a turnover tax is distortionary, because it raises the price of intermediate inputs relative to the price of labor. If the government chooses to raise taxes in this way it will collect tax revenues equal to:

$$T^{tt} = t \sum_{i=1}^{N} \sum_{s=1}^{S} MC_i^{tt}(s) Q_i^{tt}(s)$$

where $Q_i^{tt}(s)$ is the quantity produced of variety s in sector i when the turnover tax is in place and $MC_i^{tt}(s)$ is the marginal cost of producing this variety.

Second, the government can also decide to collect taxes by imposing a consumption tax equal to $v\%$ of the value of purchases of final consumers. Because such a consumption tax does not affect firms producing intermediate inputs, it does not distort production decisions. As a result, the allocation of resources within the private sector will be optimal. If the government chooses to raise revenues with the consumption tax it will collect:

$$T^{v} = v \sum_{i=1}^{N} MC_i^{v}(S) Q_i^{v}(S)$$

where $Q_i^v(S)$ is the quantity of the variety S produced by sector i and sold to final consumers when the private sector is taxed with the consumption tax and $MC_i^v(s)$ is the marginal cost of producing this variety.

Equilibrium. We solve the model as follows. First, we take the wage to be the numeraire and fix it to 1. For a given tax scheme, before-tax prices are given by equations (2.2), which depend on the wage, technological parameters and the tax scheme. Because labor supply is fixed, the household can spend an income equal to $\overline{L}$. Total income, along with the price of the final goods fixed by equations (2.2) determines the quantity of each good i the household demands.

All product markets must clear. Thus, demand of consumption by the households determines the quantity produced at the last stage of production. Production of the $N \times S$ varieties produced at stages $S-1, S-2, \ldots, 1$ is then determined recursively by the structure of production implied by assumption 1. Recall that assumption 1 states that the whole production of stage s, $Q_i(s)$ is absorbed in the production of stage $s+1$.

Given Cobb-Douglas production, this implies that $V_i(s)$, the value of output of sector i at stage s, satisfies the following relation: $V_i(s) = \omega_{1i}V_1(s+1) + \ldots \omega_{Ni}V_N(s+1)$. Collecting the $V_i(s)$ into the vector $\boldsymbol{V}(s)$, and the ω into the input-output matrix Ω we can write:

$$\boldsymbol{V}(s) = \Omega'\boldsymbol{V}(s+1) \tag{2.3}$$

I start from stage S, where final demand determines production: $V_i(S) = \beta_i(S)\overline{L}$. From there, the system of equations (2.3) can be solved backwards for all stages of production $S-1, S-2, \ldots, 1$.

The government can only choose how much of the public service it wants to provide to the household. We do not model this decision explicitly, and simply assume that the government wants to provide a fixed amount of services, equal to $\overline{G}$. Given this decision, the government needs to pay for L_G hours of work, which require tax revenues equal to $T = L_G$. Given the decisions of the firms, this total level of revenues uniquely determines the tax rate. This is t if it chooses to rely on a turnover tax or v if it chooses to rely on the consumption tax.

In this way, the model uniquely determines quantities G, T, L_G, $\{C_i,\ Q_i(s), L_i(s), M_{ij}(s)\}$, prices W, $\{P_i(s)\}$ and the tax rate (either t or v).

Welfare gain of the reform. In this section we derive the welfare gain of a tax reform that converts a turnover tax t into a consumption tax with rate v. We assume that the government always wants to raise the same amount of revenues when it switches from the turnover to the consumption tax. Thus, when the turnover tax is converted into a consumption tax, the new tax rate v is fixed so that the total value of revenues collected is not affected:

$$t\sum_{i=1}^{N}\sum_{s=1}^{S} MC_i^{tt}(s)Q_i^{tt}(s) = v\sum_{i=1}^{N} MC_i^{v}(S)Q_i^{v}(S) \tag{2.4}$$

When the wage is the numeraire, and the government provides a level of public service equal to $\overline{G}$ the indirect utility of the household is equal to:

$$U(P_C) = \left(\frac{\overline{L}}{P_C}\right)\overline{G}^{\alpha}$$

where P_C is the price index of the Cobb-Douglas aggregate (2.1):

$$P_C = \prod_{i=1}^{N} \left[\frac{P_i(S)}{\beta_i} \right]^{\beta_i} \tag{2.5}$$

Thus, the welfare gain of the reform can be calculated as the percentage change in consumption of the final goods aggregate. This is summarized by the percentage change in the price index of the aggregate final good P_C:

$$u^v - u^{tt} = p_C^{tt} - p_C^v$$

where a lowercase letter denotes the logarithm of a variable: $x \equiv \log X$ and we use tt and v to label equilibrium variables when taxes are collected on turnover or on consumption respectively. In the appendix we show that given the price equations (2.2) and the revenue neutrality of the reform (2.4) the welfare gain of the reform can be written as:

$$\begin{aligned} u^v - u^{tt} = \sum_{i=1}^{N} \beta_i \cdot & \left[\log(1+t) \sum_{j=1}^{N} \lambda_{ij} \right] \\ & + \log \left[1 - \frac{t}{1+t} \sum_{i=1}^{N} \beta_i \left(\sum_{j=1}^{N} \lambda_{ji}^{tt} \right) \right] \end{aligned} \tag{2.6}$$

where β_i are the consumption share parameters of the Cobb-Douglas aggregator (2.1) and λ_{ij} and λ_{ij}^{tt} are the ij-th elements of the following two matrices:

$$\Lambda = \left[\sum_{s=0}^{S-1} \Omega^s \right] \qquad \text{and} \qquad \Lambda^{tt} = \left[\sum_{s=0}^{S-1} \left(\frac{1}{1+t} \Omega' \right)^s \right]$$

The first matrix is a measure of how much input was used along the whole chain of production. In equation (2.6), this measure controls the extent to which the tax increases the cost of all the inputs exchanged along the production line for a given final good i, and through this channel how much it increases its final price relative to its undistorted price. The second matrix summarizes the revenues that can be raised along the production chain.

Discussion. This model allows to extend some of the intuitions introduced in section 2.2. First, the underlying parameters of production are not directly observable in standard input-output tables. In particular, the input elasticities of the different sectors, the ω_{ij} in the production functions, depend on the input-output flows reported in input-output tables as well as on the total number of stages of production S, which is not observed. Intuitively, for a given observed input share, longer production processes imply that the same inputs must be used across more stages. Thus, for a given observed input share, the input share underlying the production of longer processes must be smaller. In section 2.4.3 we explain how we calibrate the input elasticities to match the input shares observed in the input-output tables.

Second, it is interesting to note that as the number of stages S approaches infinity, the formula of the matrix that governs the price distortion Λ becomes $[I - \Omega]^{-1}$. The term $[I - \Omega]^{-1}$, appears also in Fally (2012) and Antràs et al. (2012). Fally (2012) in particular, calls it the "length of the production chain" and relates it to the total cumulative transport costs that are associated to trade when there is a constant transport cost on every transaction. The similarity should not surprise, because the turnover tax acts as a constant transport cost in our model.

2.4.2 Special case: a single final good

Before explaining how we calibrate to the Brazilian economy the model described in the previous section it is useful to study the special case in which $N = 1$ and the economy produces a single final good. This special case allows to gain some intuition over the mechanism that makes the distortion of a turnover tax grow as the production process becomes longer. The special case with a single final good is also a natural extension of the examples developed in section 2.2, where the production process is not restricted to have either two, three or an infinite number of production stages.

When private firms produce a single final good $Q(S) = C$, household's utility equals:

$$U = \frac{\overline{L}}{P(S)} \cdot G^{\alpha}$$

where $P(S)$ is the price of good $Q(S)$. Production structure in this case is a simplified version of the structure described in section 2.4.1, and variety $s = 1, 2, \ldots, S$ are produced with the following technologies:

$$
\begin{aligned}
Q(1) &= AL(1) && s = 1 \\
Q(s) &= AL(s)^{1-\omega} M(s)^{\omega} && s = 2, 3, \ldots, S
\end{aligned}
$$

Also in this case we simplify the structure of the model and assume the following:

Assumption 2. *The output of variety s is entirely absorbed by the production of the variety $s + 1$. Only the output of variety "S" is directly consumed. Formally:*

$$
\begin{aligned}
Q(s) &= M(s+1) \qquad s = 1, 2, \ldots, S-1 \\
Q(S) &= C
\end{aligned}
$$

The welfare gain of a reform that converts a turnover tax of rate t into a consumption tax of rate v is in this case:

$$
u^{v} - u^{tt} = p^{tt}(S) - p^{v}(S)
$$

where $p^{tt}(S)$ and $p^{v}(S)$ are the logarithm of the price of the final good with the turnover tax and with the consumption tax respectively. Using the condition that the government fixes the new tax rate v in order to raise the same amount of revenues collected with the turnover tax, the formula for the welfare gain (2.6) in this case becomes:

$$
u^{v} - u^{tt} = \log(1+t) \cdot \lambda + \log\left[1 - \frac{t}{1+t}\lambda^{tt}\right] \tag{2.7}
$$

where λ and λ^{tt} are a function of the input elasticity of the only sector of the economy ω and the tax rate of the turnover tax t:

$$
\lambda = \left[\sum_{s=0}^{S-1} \omega^{s}\right] \qquad \text{and} \qquad \lambda^{tt} = \left[\sum_{s=0}^{S-1} \left(\frac{\omega}{1+t}\right)^{s}\right]
$$

Equation (2.7) allows to develop some intuition over how the length of the production chain S affects the welfare gain of a reform that eliminates a tax on turnover. The equation shows that the overall gain depends on two terms. The first term captures by how much the price of the final good grows as a result of having a tax that makes input of production more expensive. This term depends on λ. In the case of single final good it is easy to show that this parameter is uniquely determined by the input share in the whole economy observed when the turnover tax has been removed and it is not affected by the number of stages of production. In other words, observing the total value of inputs over the gross value of production after the reform has eliminated the turnover tax is sufficient to calibrate the parameter λ, without further information on the number of stages of production S. To see this, write:

$$\begin{aligned} o &= \frac{\sum_{s=2}^{S} X^v(s)}{\sum_{s=1}^{S} V^v(s)} \\ &= \frac{\sum_{s=1}^{S-1} \omega^s}{\sum_{s=0}^{S-1} \omega^s} \\ &= \frac{\lambda - 1}{\lambda} \end{aligned}$$

The second term of the welfare gain captures the revenues that are raised by the turnover tax along the production chain, and it depends on λ^{tt}. Contrary to what happens with λ, for a given observed input share o, the parameter λ^{tt} does depend on the number of stages of production S. In particular, numerical simulations show that in the case of a single final good for a given o, the parameter λ^{tt} is a decreasing function of the number of stages of production S.[10] Because total tax revenue raised with the turnover tax is equal to $T = t\left(\overline{L}\lambda^{tt}\right)/(1+t)$, this implies that the tax base of the turnover tax is a decreasing function of the number of stages of production. Thus, looking at an economy with an observed input share o, a given tax rate of t raises more revenues when the production process has fewer stages. Alternatively, the same level of tax revenue T can be raised with lower tax rates in an economy that has fewer production stages. Both exercises yield the same conclusion: namely, that for a constant input share o the welfare gain of a reform that converts a turnover tax into a consumption tax is greater in longer production processes.

The intuition for the result is the following. Longer production processes allow production to be more specialized across different units of production. This gives greater

[10] It seems to be possible to prove this result analytically. We did not have time to include this formal proof in the current version of the paper.

opportunity to substitute the taxed input with labor, and allows longer production processes to shift production towards the last stage of production to a greater extent. This in turn reduces the overall tax base, and it is the source of the greater welfare gain that we estimate for longer production processes in the next section.

2.4.3 Welfare gains

In this section we use the model developed in section 2.4.1 to estimate the welfare gains of the Brazilian 2002-3 tax reform. First, we present the source of the data we use. Next, we explain how we calibrate the parameters of the model to the data and present the estimates for the welfare gain.

Data. We calibrate the main parameters of the model presented in section 2.4.1 using the Brazilian input-output tables. The tables are compiled by the *Instituto Brasileiro de Geografia e Estatística* (IBGE) based on national accounts statistics (IBGE (2008))[11] and report total value of production, total value of inputs used and total value of goods sold to domestic and foreign consumers for 110 commodities and 55 industries.

We use the Brazilian input-output tables for the year 2005. The original tables are prepared following the international standards and report "rectangular" matrices that allow for the production of many goods by a single sector. The "make" matrix reports total production of each of 110 commodities carried out by 55 different industries, while the "use" matrix reports the consumption of the 110 commodities as inputs by the same 55 industries. We use the "make" matrix (table 1) and the domestic "use" matrix at base price (table 3) to compute a "square" matrix of direct requirement coefficients. We construct a 100×100 commodity-by-commodity matrix of direct requirement coefficients by first constructing a 110×110 input-output table and then aggregating the 11 agricultural commodities reported in the table.[12] We follow the "Industry Technology Assumption" (Guo et al. (2013)) and compute the typical entry of matrix O, the value

[11] See `http://www.ibge.gov.br/home/estatistica/economia/matrizinsumo_produto/default.shtm`.

[12] These are: rice, maize, wheat, sugarcane, soy, manioc, tobacco, cotton, citrus, coffee and other agricultural products.

of inputs j used to produce 1 Real of product i (o_{ij}) as:

$$o_{ij} = \sum_{k=1}^{N} \left(\frac{X_{kj}}{\sum_{j'=1}^{N} V_{kj'}} \cdot \frac{V_{ki}}{\sum_{k'=1}^{N} V_{k'i}} \right)$$

where V_{ki} is the value of output of commodity i produced by industry k from the "make" matrix, $\sum_{j'=1}^{N} V_{kj'}$ is the total value of output in industry k, $\sum_{k'=1}^{N} V_{k'i}$ is the total value of the production of commodity i by any industry and X_{kj} is the value of domestically-produced commodity j purchased by industry k as an input from the "use" matrix. In words, the "Industry Technology Assumption" construct the share of inputs j in the production of good i as the weighted share of inputs used by all industries k that produce good i, with weights equal to the share of production of commodity i that is accounted by each industry k. Consumption shares of good i, β_i, are computed using data on the value of goods sold to domestic and foreign consumers.

Data on aggregate revenues collected with COFINS and PIS as a percentage of Brazilian GDP come from the Brazilian Ministry of Finance.[13]

Calibration. With the data from the Brazilian input-output tables, we now use the model of section 2.4.1 to evaluate the productivity gains of the 2002-3 Brazilian tax reform.

The reform eliminated two turnover taxes with statutory tax rates equal to 3.65% in total. We estimate the welfare gains of this reform using equation (2.6). In this equation, the consumption shares of the N sectors, β_i, are directly observable from national accounts. However, neither the number of stages of production S, nor the input elasticities collected in the matrix Ω are directly reported in official statistics. Moreover, although in principle the statutory tax rate t is observable, in practice the average tax rate effectively enforced may be lower due to tax evasion. Because the actual welfare gain of the reform depends on the tax rate that is effectively enforced before the policy change, it is important to estimate the gains while allowing the effective tax rate to be different from the statutory one.

In what follows we use the information contained in standard input-output tables to calibrate the input elasticities in Ω. We also use the input-output tables and the observed

[13]See `http://www.receita.fazenda.gov.br/Historico/EstTributarios/Estatisticas/default.htm`.

value of tax revenues collected to calibrate the effective tax rate for the same production processes. Because input-output tables and tax revenues are not enough to calibrate both the input elasticities, the tax rate and the number of stages of production, we proceed in the following way. First, we set the number of stages of production in our model as an exogenous parameter. Next, we calibrate the remaining parameters (Ω and t) to match the data contained in the input-output tables and in the tax revenue record. This allows us to derive the production gains of eliminating a tax on turnover as a function of the number of stages of production. We now illustrate the calibration of Ω and t.

The Brazilian input-output tables report information about "O": the matrix collecting the "direct coefficients" of the N sectors in the economy. For any sector i in the economy, the direct coefficient o_{ij} contains information about the total value of inputs purchased by sector i from sector j per each dollar of production of sector i. Calling V_i the gross value of production of sector i and X_{ij} the total value of inputs that sector i purchases from sector j, o_{ij} is:

$$o_{ij} = \frac{X_{ij}}{V_i} \tag{2.8}$$

The procedure described here shows how, once the number of stages S is fixed, the production parameters in Ω can be calculated from the matrix of direct coefficients O.

In the model of section 2.4.1, inputs are purchased at every stage of production after the first one. At the same time, output is sold at every stage. Thus, equation (2.8) is equal to:

$$o_{ij} = \frac{\sum_{s=2}^{S} X_{ij}(s)}{\sum_{s=1}^{S} V_i(s)}$$

where $X_{ij}(s)$ is the value of inputs purchased by firms producing at stage s in sector i from firms producing at stage $s-1$ in sector j and $V_i(s)$ is the gross output of firms producing at stage s in sector i. Because ω_{ij} are assumed to be identical across stages, $X_{ij}(s)/V_i(s) = \omega_{ij}$ for all stages after the first one. Thus, the last equation can be written as:

$$\begin{aligned} o_{ij} &= \omega_{ij} \left(\sum_{s=2}^{S} \frac{V_i(s)}{V_i} \right) \\ &= \omega_{ij} \left(1 - \frac{V_i(1)}{V_i} \right) \end{aligned}$$

Equation (2.9) makes clear that the observed direct coefficients are equal to the underlying input elasticities times the share of gross output produced by stages that use intermediate inputs in production. Intuitively, the direct coefficients reported in the input-output matrix are equal to a weighted average of the input elasticities of every stage of production, with weights equal to the share of production of every stage. Because stage 1 only uses labor, its input elasticities are all equal to 0. Moreover, because within a sector, input elasticities are identical across the following S stages of production, direct coefficients are equal to the underlying input elasticities times the share of output produced by the stages following the first one. For production processes that involve a large number of stages, the share of output contributed by stage 1 becomes negligible, and observed direct coefficients become identical to the true input elasticities.

Collecting equations (2.9) into matrices, the true parameters ω of the model must solve the following system of equations:

$$\begin{bmatrix} o_{11} & o_{12} & \dots & o_{1N} \\ o_{21} & o_{22} & \dots & o_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ o_{N1} & o_{N2} & \dots & o_{NN} \end{bmatrix} =$$

$$\begin{bmatrix} 1-\theta_1(1) & 0 & \dots & 0 \\ 0 & 1-\theta_2(1) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1-\theta_N(1) \end{bmatrix} \cdot \begin{bmatrix} \omega_{11} & \omega_{12} & \dots & \omega_{1N} \\ \omega_{21} & \omega_{22} & \dots & \omega_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \omega_{N1} & \omega_{N2} & \dots & \omega_{NN} \end{bmatrix} \tag{2.9}$$

where $\theta_i(1) \equiv V_i(1)/V_i$ is the share of production of the first stage in sector i.

There are two challenges to solve the system (2.9). First, we must express the shares of output at the first stage, $\theta_i(1)$, in terms of the parameters of the model. This is made possible by the sequential structure imposed to the model by assumption 1, that implies that the vector of output of the N sectors at stage s can be written as $\boldsymbol{V}(s) = \Omega' \boldsymbol{V}(s+1)$ for all stages from 1 through $S-1$. Solving the system recursively, the vector of output at stage 1 is: $\boldsymbol{V}(1) = (\Omega')^{S-1} \boldsymbol{V}(S)$. This in turn means that the vector $\boldsymbol{\theta}(1)$ of shares of output at stage 1 depends on the input elasticities and on the number of stages of production:

$$\boldsymbol{\theta}(1) = (\Omega')^{S-1}\, \boldsymbol{\iota} \tag{2.10}$$

where $\boldsymbol{\iota}$ is a $N \times 1$ vector of ones.

The second issue to address is that for large input-output matrices, the system (2.9) is hard to solve numerically. We use Brazilian input-output tables with 100 industries: this result in a system of 10000 non-linear equations in 10000 unknowns. We make the system tractable by exploiting the fact that the ratio of any two input elasticities of a given sector i depends only on observables direct coefficients. Cobb-Douglas production in fact implies:

$$\frac{\omega_{ij}}{\omega_{ip}} = \frac{o_{ij}}{o_{ip}} \tag{2.11}$$

Thus, if we select an input j that is used by every other sector in the economy (i.e. such that $o_{ij} \neq 0$ for all sectors i), we can re-write the matrix Ω as:

$$\begin{bmatrix} \omega_{11} & \omega_{12} & \dots & \omega_{1N} \\ \omega_{21} & \omega_{22} & \dots & \omega_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \omega_{N1} & \omega_{N2} & \dots & \omega_{NN} \end{bmatrix} =$$

$$\begin{bmatrix} \omega_{1j} & 0 & \dots & 0 \\ 0 & \omega_{2j} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \omega_{1j} \end{bmatrix} \cdot \begin{bmatrix} \frac{o_{11}}{o_{1j}} & \frac{o_{12}}{o_{1j}} & \dots & \frac{o_{1N}}{o_{1j}} \\ \frac{o_{21}}{o_{2j}} & \frac{o_{22}}{o_{2j}} & \dots & \frac{o_{2N}}{o_{2j}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{o_{N1}}{o_{Nj}} & \frac{o_{N2}}{o_{Nj}} & \dots & \frac{o_{NN}}{o_{Nj}} \end{bmatrix} \tag{2.12}$$

This way of expressing matrix Ω allows us to write system (2.9) as a system of only 100 non-linear equations in 100 unknowns. This system is solvable if there is at least one sector whose inputs are used by every other sector in the economy: this is necessary in order to avoid that any of the entries in the second matrix of equation (2.12) has a denominator equal to 0. We select sector “information services”, and estimate the input shares of this input for all 100 sectors of the economy. This sector is one of the 7 sectors whose output is used by every other sector in the economy as input. Next, we use these estimates to back out all the other 9900 parameters of the Ω matrix, using equations (2.11).

To summarize, we use the structure of the model in section 2.4.1 and calibrate the input elasticities of production so that they match the observed direct coefficient reported in the input-output matrix of Brazil. Thus, the ω_{ij} are the solutions to the system (2.9), where the $\theta_i(1)$ are defined by (2.10). We solve the system numerically on a 100×100 matrix by substituting the matrix Ω with the two matrices on the right-hand side of (2.12): this reduces the complexity of the model and allows us to solve it for only 100 unknowns. Finally, we use equations (2.11) to retrieve all the other unknowns of system (2.9).

I now explain how we calibrate the effective tax rate of the turnover tax. From the public records made available by the Brazilian Ministry of Finance we observe the value of revenues raised by COFINS and PIS as a share of the Brazilian GDP in the years before the reform. In the years 2001 and 2002 total revenues from these two turnover tax was equal to 4.63% and 4.62% respectively. Thus, we set $T/\overline{L}$ in the model to 0.046, and set the effective tax revenue of the turnover tax to hit this aggregate value. Because we still do not have enough information to calibrate the number of stages of production we proceed as for the calibration of Ω. First, we set the number of stages of production as an exogenous parameter. Next, we use the procedure outlined above to back out the implied input elasticities ω_{ij}. Finally, we use the definition of tax revenues in our model, and find the effective tax rate t that raises 4.6% of national income in turnover taxes, for the assumed number of stages. This requires solving the following equation in t:

$$
\begin{aligned}
0.046 &= \sum_{s=1}^{S} MC_i^{tt}(s) Q_i^{tt}(s) \\
&= \frac{t}{1+t}\left[\sum_{i=1}^{N} \beta_i \left(\sum_{j=1}^{N} \lambda_{ji}^{tt}\right)\right]
\end{aligned}
$$

where λ_{ij}^{tt} is the ij-th entry of the matrix Λ^{tt} defined in section 2.4.1 and depends on the number of stages of production S.

Results. Once we calibrate Ω and t, we use the formula (2.6) to estimate the welfare gains of the reform. We do this for several exogenous production lengths, in order to evaluate how the number of stages affects the overall welfare gains of the reform. Baseline results are shown in figure 2.6, where we plot on the horizontal axis the number of stages on which we calibrate Ω and t, and on the vertical axis the welfare gain of converting a turnover tax that raises revenues equal to 4.6% of national income into a consumption tax. The graph shows that the welfare gain of the reform is 4.1 times larger when the production process has an infinite number of stages compared to a production process in which there are only two stages of production. Most of this increase happens relatively quickly: the welfare gain of the reform is 3.95 times larger when the production process involves 8 stages compared to when the production process involves only two stages. Overall, for the input-output table observed in Brazil, the productivity gains of the tax reform remain small even when production chains are infinitely long. For example, eliminating the tax on turnover increases aggregate production by at most 0.05%, when the production chain is infinitely long.

This baseline estimate is obtained by calibrating an input-output with 100 sectors. An interesting question is what would have been the welfare gain of the reform had we treated Brazil as a single-sector economy with an overall intermediate inputs intensity ω equal to the input share observed in the aggregated input-output matrix. This is an interesting exercise because many of the developing countries that make use of turnover taxes do not produce detailed input-output matrices. To answer this question, we adapt equation (2.6) to the special case in which there is a single final good produced, and calculate the welfare gain of the reform using equation (2.7). As for the model with many sectors, we set the number of stages as an exogenous parameter, and calibrate the input elasticity ω so that the observed input share in the undistorted economy o is equal to 43%, the observed input share in Brazil in 2005. For the same number of stages, we also set the effective tax rate t equal to the tax rate that would raise revenues equal to 4.6% of national income.

The results of this exercise are shown in figure 2.7. The figure plots for any given number of stages of production (on the horizontal axis) the welfare gain of the reform for a single-sector economy with an overall intermediate inputs intensity of production of 43% (on the vertical axis). To ease comparison with the previous exercise, we plot in the same figure the welfare gain obtained calibrating the model with 100 sectors shown in figure 2.6. The graph shows that the welfare gain of the reform are several times larger in longer production processes: in the single-sector model, the welfare gains of the reform are 7.1 times larger when the number of stages converges to infinity than when inputs are exchanged a single time. Interestingly, for any given production length, the welfare gain of the reform are always larger when we calibrate the model using the information from the full input-output matrix. However the difference between the two estimates is small: when the number of stages converges to infinity, the welfare gain calculated using the information from the full input-output matrix is only 1.14 times larger than the welfare gains obtained assuming that Brazil is an economy that produces a single final good.

2.5 Constant elasticity of substitution production

In the model of this section we relax the assumption of Cobb-Douglas production. The objective is to evaluate how the elasticity of substitution between labor and intermediate inputs affects the welfare gains of the reform.

2.5.1 The model

The model developed in this section is identical to the model of section 2.4.2, except for the production function, which is assumed to have a more general structure.

Households. Preferences are identical to the ones derived in the special case with a single final good in section 2.4.2. The representative consumer has preferences defined over a single final good C supplied by private firms and a public good G provided by the government:

$$U = C \cdot G^{\alpha}$$

The household supplies inelastically $\overline{L}$ hours of work, receiving an hourly wage of W. Because we set the wage to be the numeraire, total income in the economy is $\overline{L}$.

Firms. The model has a single sector that produces S varieties.[14] Each variety is produced under perfect competition. Production of the first variety requires only labor. The following varieties are produced combining labor and intermediate inputs using a generalized constant elasticity of substitution (CES) function. Thus, technology is described by:

$$
\begin{aligned}
&Q(1) = BL(1) && s = 1 \\
&Q(s) = B\left[L(s)^{\frac{\eta-1}{\eta}} + \delta M(s)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} && s = 2, 3, \ldots, S
\end{aligned}
\tag{2.13}
$$

In (2.13) B is a Hicks-neutral productivity parameter, $L(s)$ and $M(s)$ are labor and intermediate inputs used by firms producing variety s, η is the elasticity of substitution between the two factors of production and δ is a distribution parameter that governs the importance of intermediate input relative to labor. We assume that the productivity parameter B is identical across all varieties and that both the distribution parameter δ and the elasticity of substitution η are the same across all varieties that use intermediate inputs.

In order to solve the model we make the following assumption:

Assumption 3. *The output of every variety s is entirely used as an intermediate input by the production of variety $s+1$. Only the output of variety S is consumed. Formally:*

$$
\begin{aligned}
Q(s) &= M(s+1) \qquad s = 1, 2, \ldots, S-1 \\
Q(S) &= C
\end{aligned}
$$

As assumption 1, assumption 3 gives to the production process a clear sequential structure. It also allows to pin down the price of every variety produced. The price of variety 1 is determined by the wage alone. The price of the following varieties is then determined recursively with the wage and the price of the variety produced immediately

[14] Extending the theory to allow for N sectors, each producing many varieties combining labor and inputs from all other sectors with arbitrary elasticity of substitution is relatively easy. However, the calibration of the model to the observed input-output matrix presents some challenges that we have still not been able to solve. In section 2.4.3 we have shown that considering more than one sector does not have a large effect on the estimated welfare gains of the reform. In this section we explore whether maintaining the single-sector assumption and allowing the elasticity of substitution to vary has larger effects on the welfare gains of the reform.

upstream. Formally this can be written:

$$\begin{aligned} &P(1) = B^{-1}W \\ &P(s) = B^{-1}\left[W^{1-\eta} + \frac{\gamma}{A}P(s-1)^{1-\eta}\right]^{\frac{1}{1-\eta}} \qquad\qquad s = 2, 3, \ldots, S \end{aligned} \tag{2.14}$$

where we define $A \equiv B^{1-\eta}$ and $\gamma \equiv \delta^{\eta}$.

Government. As in section 2.4.1, the government produces a single service G with the following constant return to scale technology:

$$G = L_G$$

Of this service, the government provides the household an exogenous level $\overline{G}$, free of charge. The government finances the production of this services with tax revenues.

As in the previous model, we assume that the government can chose to raise tax with either a turnover tax or a tax on consumption. A turnover tax with rate t would raise total revenues equal to:

$$T^{tt} = t\sum_{s=1}^{S} MC^{tt}(s)Q^{tt}(s)$$

A tax on consumption with a rate of v would raise total revenues equal to:

$$T^{v} = vMC^{v}(S)Q^{s}(S)$$

In the following we consider welfare gains of a reform that converts a turnover tax into a consumption tax, while leaving the total revenue raised by the tax unaffected. Before that, we briefly explain how the model solves.

Equilibrium. Income is equal to $W\overline{L} = \overline{L}$. The prices of the varieties produced at every stage are uniquely determined by the wage, the structure of production and the type of taxation. Thus, the system of prices (2.14) determines the price of the consumption good, which is the variety produced at the last stage of production. Given this price and the income in the economy, the household's preferences will determine the demand of the good produced at the last stage of production. Once demand of the last variety is determined, the market clearing conditions and the structure of production determine the quantity produced at every stage of production.

Welfare gain of the reform. We derive here the expressions to calculate the welfare gain of a reform that converts a turnover tax into a consumption tax. As in the exercises of section 2.4, we assume that the reform leaves the overall level of revenues unaffected:

$$t \sum_{s=1}^{S} MC^{tt}(s) Q^{tt}(s) = v MC^{v}(S) Q^{v}(S)$$

This also means that the government will be able to produce and provide the same amount of public good before and after the reform. Because income in units of labor is unaffected by the reform, the welfare gain of the reform can be written as a function of the change in price of the final good:

$$u^{v} - u^{tt} = p^{tt}(S) - p^{v}(S)$$

In the appendix B we show that the welfare gain in this model can be written as:

$$\begin{aligned} u^{v} - u^{tt} &= \log \left[\frac{(1+t) MC^{tt}(S)}{(1+v) MC^{v}(S)} \right] \\ &= \log \Bigg\{ (1+t) \cdot \frac{1 - \left[\frac{\gamma}{A}(1+t)^{1-\eta}\right]^{S}}{1 - \frac{\gamma}{A}(1+t)^{1-\eta}} \cdot \frac{1 - \frac{\gamma}{A}}{1 - \left(\frac{\gamma}{A}\right)^{S}} \times \\ & \quad \left[1 - \frac{t}{1+t} \left(1 + \sum_{s=1}^{S-1} \prod_{k=s+1}^{S} \frac{\omega^{tt}(k)}{1+t} \right) \right] \Bigg\} \end{aligned} \tag{2.15}$$

In (2.15) $\omega^{tt}(k)$ is the cost share of intermediate inputs at stage k:

$$\omega^{tt}(k) = \frac{\gamma \left[P^{tt}(k-1)\right]^{1-\eta}}{1 - \gamma \left[P^{tt}(k-1)\right]^{1-\eta}} \tag{2.16}$$

In the appendix B we also show that given the way prices are set in this economy, the sum of products in the second line of equation (2.15) can be written as a function of the parameters of the model:

$$\sum_{s=1}^{S-1} \prod_{k=s+1}^{S} \frac{\omega^T(k)}{1+t} = \frac{\left[\frac{\gamma}{A}(1+t)^{-\eta}\right]^S}{1-\left[\frac{\gamma}{A}(1+t)^{1-\eta}\right]^S} \times \tag{2.17}$$
$$\left\{ \frac{\left[\frac{\gamma}{A}(1+t)^{-\eta}\right]^{1-S} - 1}{1-\frac{\gamma}{A}(1+t)^{-\eta}} + \frac{1+t-(1+t)^S}{t} \right\}$$

In the next section we explain how we calibrate the parameters of this model to the Brazilian economy and use equation (2.15) to calculate the welfare gains of the 2002-3 reform.

2.5.2 Welfare gains

Calibration. In this section we explain how we estimate the welfare gain of the 2002-3 Brazilian tax reform using equation (2.15) and data from the Brazilian economy.

As in section 2.4.3, we use two observable moments to calibrate the model. The first one is the intermediate input share of the Brazilian economy, as observed after the 2002-3 reform. The second one is the revenues raised by turnover taxes before the reform as a share of Brazilian GDP. As it was the case for the model described in section 2.4, these variables are not sufficient to calibrate all the parameters of the model: the number of stages of production S, the elasticity of substitution between intermediate inputs and labor η, the distribution parameter δ, the productivity B and the effective tax rate t. For this reason, we extend the approach adopted in section 2.4.3 and proceed as follows. First, we set the number of stages of production and the elasticity of substitution as exogenous parameters. Next, we use information on the intermediate input share coming from the input-output tables along with the structure of the model to calibrate the parameters of the production function. This procedure does not allow to identify separately the distribution parameter δ and the technology parameter B. Instead, it allows to set the joint parameter $\gamma/A = \delta^\eta/B^{1-\eta}$. Because in equations (2.15) and (2.17) δ and B always enter as $\delta^\eta/B^{1-\eta}$, this joint parameter is sufficient to calculate the welfare gain of the reform. Finally, we use information about the tax revenue collected before the reform to calibrate the effective tax rate t. This approach is equivalent to the one we adopted in section 2.4.3, where we also set the tax rate in order to match the available data on tax revenues. The procedure builds on the idea that the information on the total value of tax revenues can be observed with greater precision than the tax rate effectively paid by Brazilian firms before the reform.

We start by describing the procedure to calibrate γ/A. Since the economy produces a single final good, we use the information on the aggregate share of gross output that is used as intermediate input by any sector of the economy after the tax reform. If V is gross output and X is gross value of intermediate inputs, this share is:

$$\begin{aligned} o &= \frac{X}{V} \\ &= \frac{\sum_{s=2}^{S} X(s)}{\sum_{s=1}^{S} V(s)} \end{aligned}$$

Multiplying and dividing every term in the numerator by the value of of all varieties produced in the following stages, we can write the observed input share as a function of the underlying intermediate input shares at all stages of production, $\omega(k)$:

$$\frac{o}{1-o} = \sum_{s=1}^{S-1} \prod_{k=s+1}^{S} \omega(k) \tag{2.18}$$

When the production function has constant elasticity of substitution, the terms $\omega(k)$ on the right-hand side of equation (2.18) are:

$$\omega(k) = \frac{\gamma}{A} \frac{1-\left(\frac{\gamma}{A}\right)^{k-1}}{1-\left(\frac{\gamma}{A}\right)^{k}}$$

Plugging this expression of $\omega(k)$ back into equation (2.18), it is possible to write the input share o observed after the reform as a function of the parameters of the model (see also the derivations in appendix B):

$$\frac{o}{1-o} = \frac{\left(\frac{\gamma}{A}\right)^{S}}{1-\left(\frac{\gamma}{A}\right)^{S}} \left[\frac{\left(\frac{\gamma}{A}\right)^{S-1}-1}{1-\frac{\gamma}{A}} - S + 1 \right] \tag{2.19}$$

Equation (2.19) is the key expression used to calibrate the parameters γ/A in the constant elasticity of substitution model. Because in the 2005 Brazilian input-output tables we observe an aggregate input share equal to 43%, we set γ/A equal to the value that solves (2.19) when $o = 0.43$. We do this for several different values for the elasticity of substitution and the number of stages.

In the second step of the calibration, we use the values of γ/A that we have found in the first step of the procedure along with the tax revenue observed before the reform to calibrate the effective tax rate t. We do so by writing down the value of tax revenue as a share of total income as a function of the cost shares of every stage of production:

$$\begin{aligned}\frac{T}{L} &= t\sum_{s=1}^{S} MC^{tt}(s)Q^{tt}(s) \\ &= \frac{t}{1+t}\left[1+\sum_{s=1}^{S-1}\prod_{k=s+1}^{S}\frac{\omega^{tt}(k)}{1+t}\right]\end{aligned}$$

Using the expression for $\omega^{tt}(k)$ reported in (2.16), the last equation can be written as (see also the derivations in appendix B):

$$\begin{aligned}\frac{T}{L} =& \frac{t}{1+t}\left\{1+\frac{\left[\frac{\gamma}{A}(1+t)^{-\eta}\right]^{S}}{1-\left[\frac{\gamma}{A}(1+t)^{1-\eta}\right]^{S}}\times\right. \\ & \left.\left[\frac{\left[\frac{\gamma}{A}(1+t)^{-\eta}\right]^{S-1}-1}{1-\frac{\gamma}{A}(1+t)^{-\eta}}+\frac{1+t-(1+t)^{S}}{t}\right]\right\}\end{aligned} \tag{2.20}$$

Equation (2.20) is the key expression used to calibrate the effective tax rate t in the constant elasticity of substitution model. Because Brazilian turnover taxes collected around 4.61% and 4.63% of GDP in the two years leading to the tax reform, we set the left hand side of equation (2.20) equal to 0.046, and for every number of stages, elasticity of substitution and corresponding γ/A, we set the effective tax rate t equal to the value that solves equation (2.20).

Results. Once we calibrate γ/A and the tax rate, we use equation (2.15) to evaluate the welfare gains of the 2002-3 tax reform. As in section 2.4.3, we perform this exercise for several different values of production length and elasticity of substitution, in order to analyze how these two parameters affect the welfare gains.

The baseline results of this exercise are shown in figure 2.8, where we plot on the horizontal axis the number of stages of production and on the vertical axis the welfare gains implied by equation (2.15). In the figure, we plot the welfare gains for three different values of the elasticity of substitution η: 0.5 (in blue), 1 (in green) and 5 (in red). An elasticity of substitution equal to 0.5 corresponds to a case in which labor and the taxed input are complement in production. The other two values of η imply that labor and the intermediate input are substitutes in production. The unit-elasticity case reproduces the single-sector Cobb-Douglas model discussed in section 2.4.2, while the case with $\eta = 5$ allows for greater substitutability than in the Cobb-Douglas model.

Figure 2.8 shows that the welfare gains are larger the higher the elasticity of substitution. When the number of stages goes to infinity, an elasticity of substitution equal to 0.5 implies welfare gains that are half of those obtained in the Cobb-Douglas model (50.11%). When the elasticity of substitution is equal to 5 instead, the welfare gains of the reform in an economy that produces with an infinitely long production process are 4.8 times larger than with the Cobb-Douglas technology, and equal to 0.23% of national income.

The welfare gain of the reform becomes larger as the elasticity of substitution grows because the easier it is to substitute intermediate inputs with labor, the more a tax on turnover distorts production decisions. In the extreme case in which production at all stages after the first one is carried out with a Leontief technology ($\eta = 0$) a tax on turnover does not distort decisions, and for this reason it creates no loss at all. On the other hand, the relationship is not increasing for all values of elasticity of substitution and also when the elasticity is very high, production losses of a turnover tax turn out to be small. This is because when intermediate inputs are easy to substitute with labor, the two inputs have a similar effect on production, and the change in the input mix caused by a tax on turnover does not affect overall production very much. This inverse-U shaped relationship between elasticity of substitution and welfare gains is shown in figure 2.9. In this figure we plot on the horizontal axis the elasticity of substitution, while on the vertical axis we show how the welfare gain changes as a function of it. The figure shows the relationship between these two variables for three different production structures: one in which production is carried out in 3 stages (in blue), one in which production requires 5 stages (in green) and one in which production requires an infinite number of stages (in red). The inverse-U shape pattern appears in the three production structures. Interestingly, the elasticity of substitution that maximizes the welfare gains depends on the production structure, and it is larger for shorter production processes. In particular, the welfare gains of the reform are highest in the 3-stages economy when the elasticity of substitution between labor and intermediate inputs is equal to 50. In this case, the welfare gain are much higher than in the Cobb-Douglas case, and equal to 1.02% of national income. The highest welfare gain is reached for lower elasticities in the economies that produce with 5 and an infinite number of stages stages: 39 and 37 respectively.

Overall, this exercise suggests that the elasticity of substitution plays an important role in determining the welfare gains of the reform. In particular, if labor and intermediate inputs are highly substitutable, the welfare gains calculated with the Cobb-Douglas model may significantly underestimate the actual effects of the reform.

2.6 Discussion

In this paper we use a tax reform in Brazil to study the welfare gain of eliminating distortionary turnover taxes. The main challenge when quantifying the welfare gain of such a reform is that when the production process is sequential, the welfare gain of removing turnover taxes depends on the number of stages of production. Our approach is to develop a model of sequential production and calibrate it to the Brazilian economy. We find that when we consider production processes that have 11 or more stages of production the productivity gains of removing the turnover tax are around 4 times larger than when we consider a production process with a single stage. But overall, even with infinitely long production processes, gains from the reform are modest and equal to around 0.05% of national income.

The results of the paper have implications for tax design, especially in developing countries. Taxes on turnover are popular in countries where the informal sector is large and tax capacity low, because they are easier to administer and enforce than more complex tax schemes such as taxes on profits (Best et al. (2013)). Since taxes on turnover are known to be inefficient (Diamond and Mirrlees (1971a,b),Keen (2009)), their diffusion must reflect a belief among policy makers that the benefits in terms of higher tax compliance more than offset the productivity loss created by these taxes. Our results suggest that productivity losses are modest, and they may justify maintaining these schemes in cases in which informality is a major concern. Moreover, the gain we estimated for Brazil may be considered as an upper bound for the productivity loss of a tax on turnover for two reasons. First, Brazil in the early 2000s was relatively developed, and had a diversified economy in which sectors were connected in relatively complex production chains. To the extent that the economies of poorer countries have simpler production structures with fewer input-output linkages, the productivity loss of a tax on turnover are likely to be smaller. Second, the combined tax rates of PIS and COFINS in Brazil in 2002 were 3.65%. This tax rate is high compared to turnover taxes observed elsewhere in the world, where tax rates for turnover taxes usually do not exceed 1%. For these reasons, the welfare gain of eliminating turnover taxes in other developing countries is likely to be smaller than the one calculated here for Brazil.

These conclusions have two limitation. First, in the model we do not consider the potential effects of turnover taxes on vertical integration, which in turn could have effects on productivity.[15] This would be an interesting avenue of future research. Second, the model derives all results by assuming perfect competition at every stage of production. Relaxing this assumption and allowing different market structures would be another promising avenue for future research.

[15]The literature here is extensive: theoretical contributions are: Grossman and Hart (1986); Hart and Moore (1990); Antràs and Helpman (2004). Empirical contributions include: Woodruff (2002); Acemoglu et al. (2009, 2010); Macchiavello (2012).

References

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Acemoglu, D., Griffith, R., Aghion, P., and Zilibotti, F. (2010). Vertical integration and technology: theory and evidence. *Journal of the European Economic Association*, 8(5):989–1033.

Acemoglu, D., Johnson, S., and Mitton, T. (2009). Determinants of vertical integration: financial development and contracting costs. *Journal of Finance*, 64(3):1251–1290.

Antràs, P. and Chor, D. (2013). Organizing the global value chain. *Econometrica*, 81(6):2127–2204.

Antràs, P., Chor, D., Fally, T., and Hillberry, R. (2012). Measuring the upstreamness of production and trade flows. *American Economic Review, Papers & Proceedings*, 102(3):412–416.

Antràs, P. and Helpman, E. (2004). Global sourcing. *Journal of Political Economy*, 112(3):552–580.

Baldwin, R. and Venables, A. J. (2013). Spiders and snakes: offshoring and agglomeration in the global economy. *Journal of International Economics*, 90(2):245–254.

Banerjee, A. V. and Duflo, E. (2005). Growth Theory through the Lens of Development Economics. In Aghion, P. and Durlauf, S., editors, *Handbook of Economic Growth*, volume 1 of *Handbook of Economic Growth*, chapter 7, pages 473–552. Elsevier.

Bartelme, D. and Gorodnichenko, Y. (2014). Linkages and economic development. Mimeo.

Basu, S. (1995). Intermediate goods and business cycles: Implications for productivity and welfare. *American Economic Review*, 85(3):512–31.

Besley, T. and Persson, T. (2013). Taxation and development. In Alan J. Auerbach, Raj Chetty, M. F. and Saez, E., editors, *Handbook of Public Economics, Vol. 5*, pages 51 – 110. Elsevier.

Best, M. C., Brockmeyer, A., Kleven, H. J., Spinnewijn, J., and Waseem, M. (2013). Production vs revenue efficiency with limited tax capacity: Theory and evidence from pakistan. Mimeo.

Bovenberg, A. L. (1987). Indirect taxation in developing countries: A general equilibrium approach. *Staff Papers-International Monetary Fund*, 34(2):333–373.

Ciccone, A. (2002). Input chains and industrialization. *Review of Economic Studies*, 69(3):565–587.

Corden, W. (1966). The structure of a tariff system and the effective protective rate. *Journal of Political Economy*, 74:221–237.

De Paula, A. and Scheinkman, J. A. (2010). Value-added taxes, chain effects, and informality. *American Economic Journal: Macroeconomics*, 2(4):195–221.

Diamond, P. A. and Mirrlees, J. A. (1971a). Optimal taxation and public production I: Production efficiency. *American Economic Review*, 61(1):8–27.

Diamond, P. A. and Mirrlees, J. A. (1971b). Optimal taxation and public production II: Tax rules. *American Economic Review*, 61(3):261–78.

Emran, M. S. and Stiglitz, J. E. (2005). On selective indirect tax reform in developing countries. *Journal of Public Economics*, 89(4):599–623.

Fally, T. (2012). Production staging: Measurement and facts. Mimeo.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772.

Gottfried, P. and Wiegard, W. (1991). Exemption versus zero rating: A hidden problem of VAT. *Journal of Public Economics*, 46(3):307–328.

Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4):691–719.

Guo, J., Lawson, A. M., and Planting, M. A. (2013). *From make-use to symmetric IO tables: an assessment of alternative technology assumptions*. BiblioGov.

Harberger, A. C. (1962). The incidence of the corporation income tax. *Journal of Political Economy*, 70(3):215–240.

Hart, O. D. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6):1119–1158.

Hirschman, A. O. (1958). *The strategy of economic development*. Yale University Press New Haven.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics*, 124(4):1403–1448.

Hulten, C. R. (1978). Growth accounting with intermediate inputs. *Review of Economic Studies*, 45(3):511–518.

IBGE (2008). *Matriz de Insumo-Produto — Brasil 2000/2005*, volume 23 of *Contas Nacionais*. Instituto Brasileiro de Geografia e Estatística - IBGE, Rio de Janeiro, Brazil.

IMF (2002). IMF approves US$ 30.4 billion stand-by credit for Brazil. Press Release 02\40.

IMF (2003). IMF approves 15-month extension, and US$6.6 billion augmentation of Brazil's stand-by credit. Press Release 03\217.

Jones, C. I. (2011). Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics*, 3(2):1–28.

Jones, C. I. (2013). Misallocation, economic growth, and input-output economics. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics, Vol. II*, pages 419–454. Cambridge University Press.

Keen, M. (2009). What do (and don't) we know about the value added tax? A review of Richard M. Bird and Pierre-Pascal Gendron's the VAT in developing and transitional countries. *Journal of Economic Literature*, 47(1):159–70.

Keen, M. (2013). *Targeting, Cascading, and Indirect Tax Design*. International Monetary Fund.

Keen, M. and Lockwood, B. (2006). Is the VAT a money machine? *National Tax Journal*, 59(4):905–28.

Leontief, W. W. (1936). Quantitative input and output relations in the economic systems of the United States. *Review of Economic Statistics*, 18(3):105–125.

Long, J. B. and Plosser, C. I. (1983). Real business cycles. *Journal of Political Economy*, 91(1):39–69.

Macchiavello, R. (2012). Financial development and vertical integration: theory and evidence. *Journal of the European Economic Association*, 10(2):255–289.

Newbery, D. M. (1986). On the desirability of input taxes. *Economics Letters*, 20(3):267–270.

Oberfield, E. (2011). Business networks, production chains, and productivity: A theory of input-output architecture. Working paper, Federal Reserve Bank of Chicago.

Piggott, J. and Whalley, J. (2001). VAT base broadening, self supply, and the informal sector. *American Economic Review*, 91(4):1084–1094.

Pomeranz, D. (2013). No taxation without information: Deterrence and self-enforcement in the value added tax. Technical Report w19199, National Bureau of Economic Research.

Restuccia, D. and Rogerson, R. (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720.

Werneck, R. L. F. (2006). An evaluation of the 2003 tax reform effort in Brazil. *Revista de Economia Política*, 26(1):75–94.

Woodruff, C. (2002). Non-contractible investments and vertical integration in the Mexican footwear industry. *International Journal of Industrial Organization*, 20(8):1197–1224.

World Bank (2006). *Doing business 2007: How to reform*. World Bank Publications.

World Bank (2012). *Doing business 2013: smarter regulations for small and medium-size enterprises*. World Bank Publications.

Yi, K.-M. (2003). Can vertical specialization explain the growth of world trade? *Journal of Political Economy*, 111(1):52–102.

Figures and tables

Figure 2.1
Two versions of the one sector economy

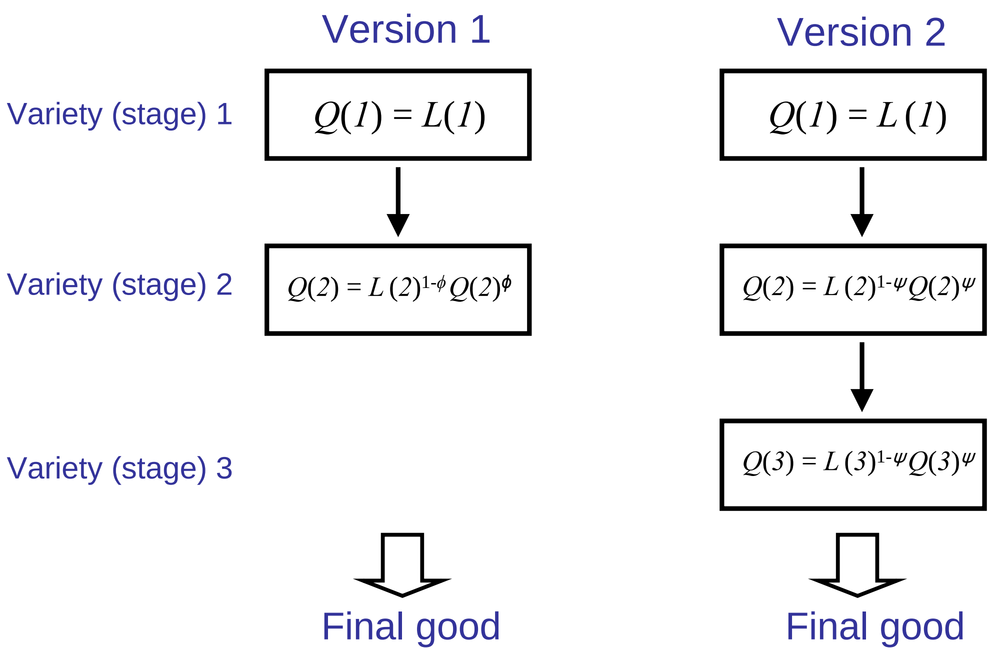

Notes: The figure illustrates the basic structure of the two versions of the economy described in section 2.2. On the left, it is represented the production process when the economy produces 2 varieties (stages). In this case, the observed input flow X is equal to the output of variety 1 sold as intermediate input to the producers of variety 2. The observed gross output V is equal to the output of variety 1 plus the output of variety 2. On the right, it is represented the production process when the economy produces 3 varieties (stages). In this case, the observed input flows X are equal to the output of variety 1 sold as intermediate input to the producers of variety 2 plus the output of variety 2 sold as intermediate input to the producers of variety 3. The observed gross output V is equal to the sum of the output of varieties 1, 2 and 3.

Figure 2.2
Welfare gain of the tax reform in the one sector example
Welfare gain with two and three stages of production

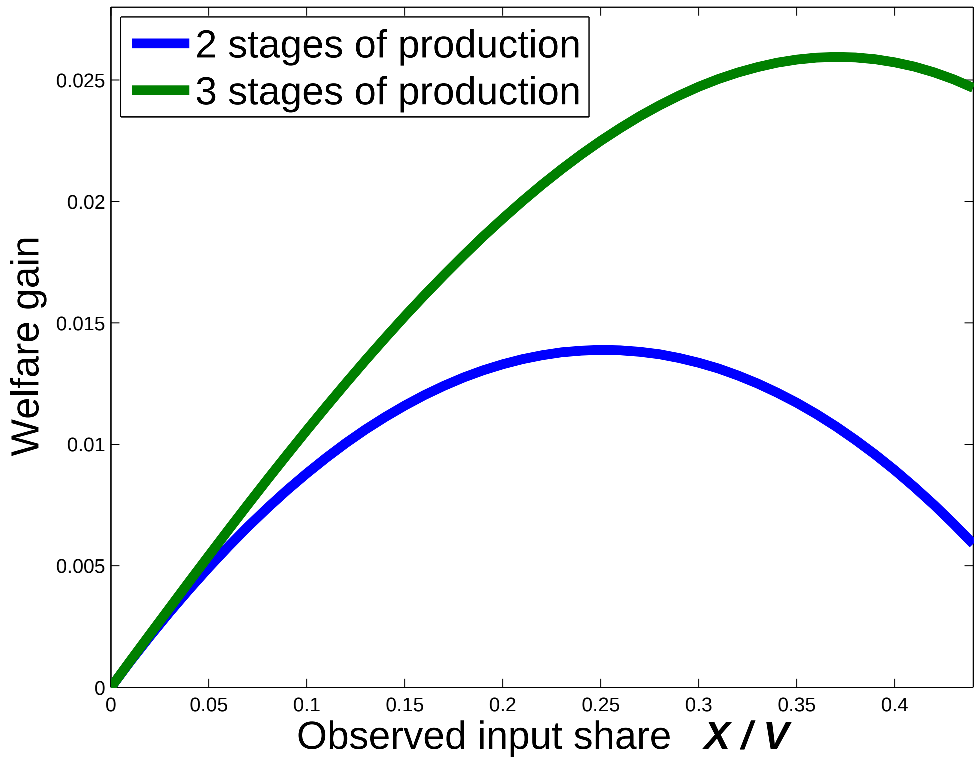

Notes: The figure shows the welfare gain (in percentage points) of removing a turnover tax as a function of the observed input share of an economy in the first two versions of the economy described in section 2.2. The blue line shows the welfare gain of the reform when the economy produces in two stages, and the green line shows the welfare gain when the economy produces in three stages. Both lines are drawn using equation (2.7) when the number of stages is either two or three. For every assumed production length, the input elasticities ω are calibrated to the observed input share shown on the horizontal axis, while the tax rate is set to match a value of tax revenue over national income equal to 4.6% (the value that is observed in Brazil before the 2002-3 reform).

Figure 2.3
Welfare gain of the tax reform in the one sector example
Welfare gain with two, three and an infinite number of production stages

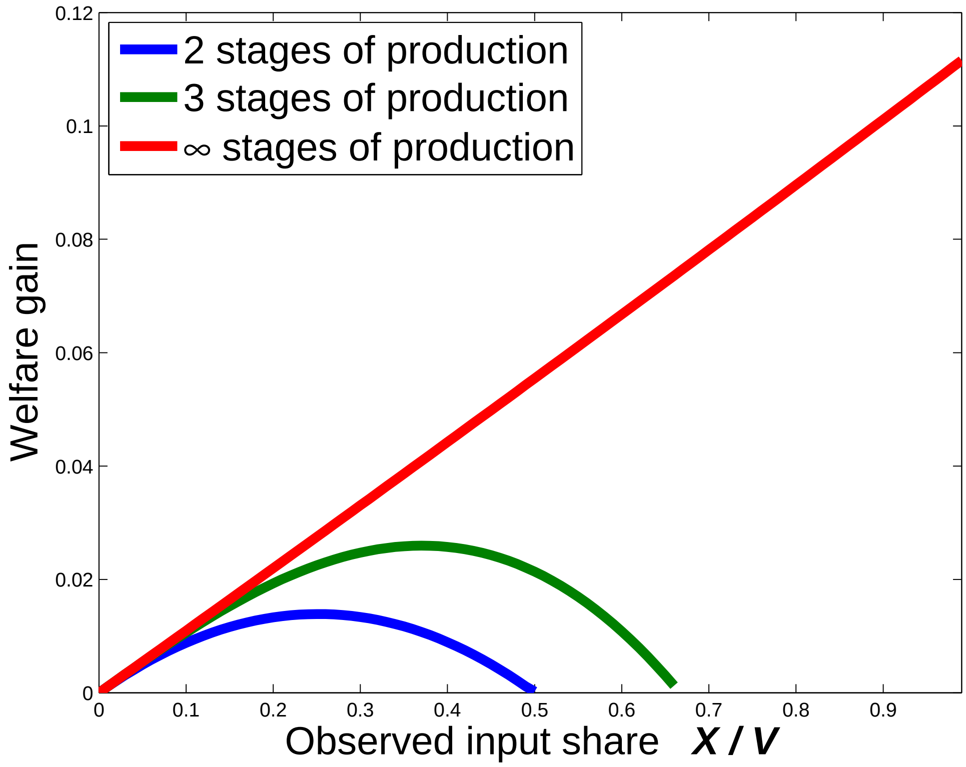

Notes: The figure shows the welfare gain (in percentage points) of removing a turnover tax as a function of the observed input share of an economy in the three versions of the economy described in section 2.2. The blue line shows the welfare gain of the reform when the economy produces in two stages, the green line shows the welfare gain when the economy produces in three stages and the red line shows the welfare gain when the economy produces with an infinite number of stages. All lines are drawn using equation (2.7) when the number of stages is either two, three or it approaches infinity. For every assumed production length, the input elasticities ω are calibrated to the observed input share shown on the horizontal axis, while the tax rate is set to match a value of tax revenue over national income equal to 4.6% (the value that is observed in Brazil before the 2002-3 reform).

Figure 2.4
Welfare gain of the tax reform in the one sector example
Tax revenues equal to 10% of national income

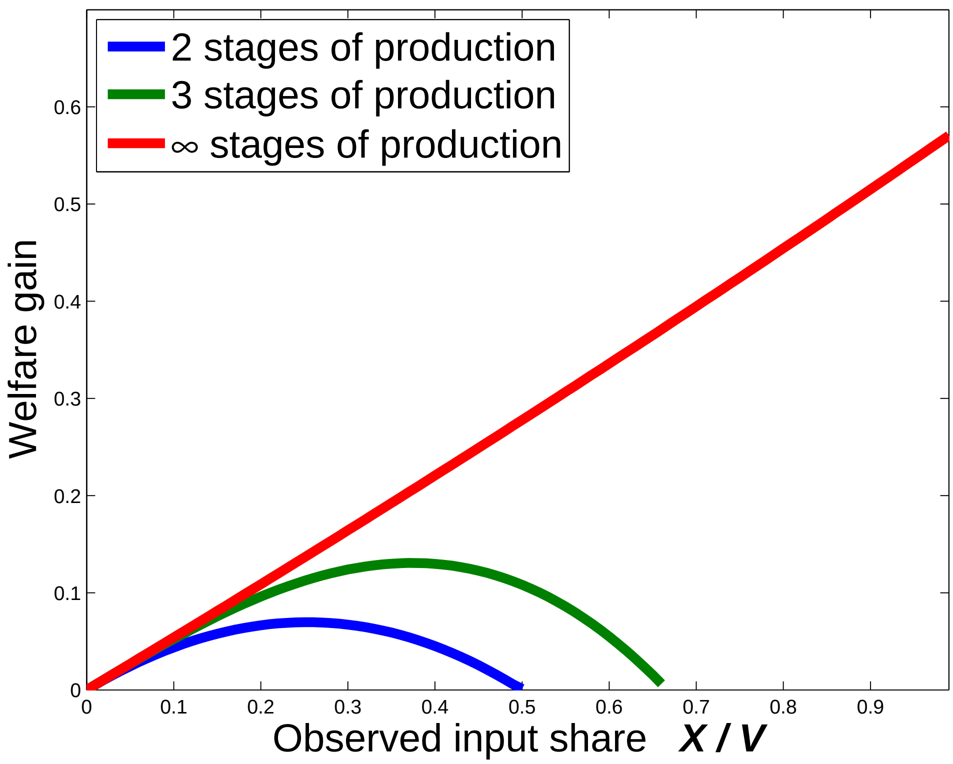

Notes: The figure shows the welfare gain (in percentage points) of removing a turnover tax as a function of the observed input share of an economy in the three versions of the economy described in section 2.2. The blue line shows the welfare gain of the reform when the economy produces in two stages, the green line shows the welfare gain when the economy produces in three stages and the red line shows the welfare gain when the economy produces with an infinite number of stages. All lines are drawn using equation (2.7) when the number of stages is either two, three or it approaches infinity. For every assumed production length, the input elasticities ω are calibrated to the observed input share shown on the horizontal axis, while the tax rate is set to match a value of tax revenue over national income equal to 10%.

Figure 2.5
Welfare gain of the tax reform in the one sector example
Tax revenues equal to 15% of national income

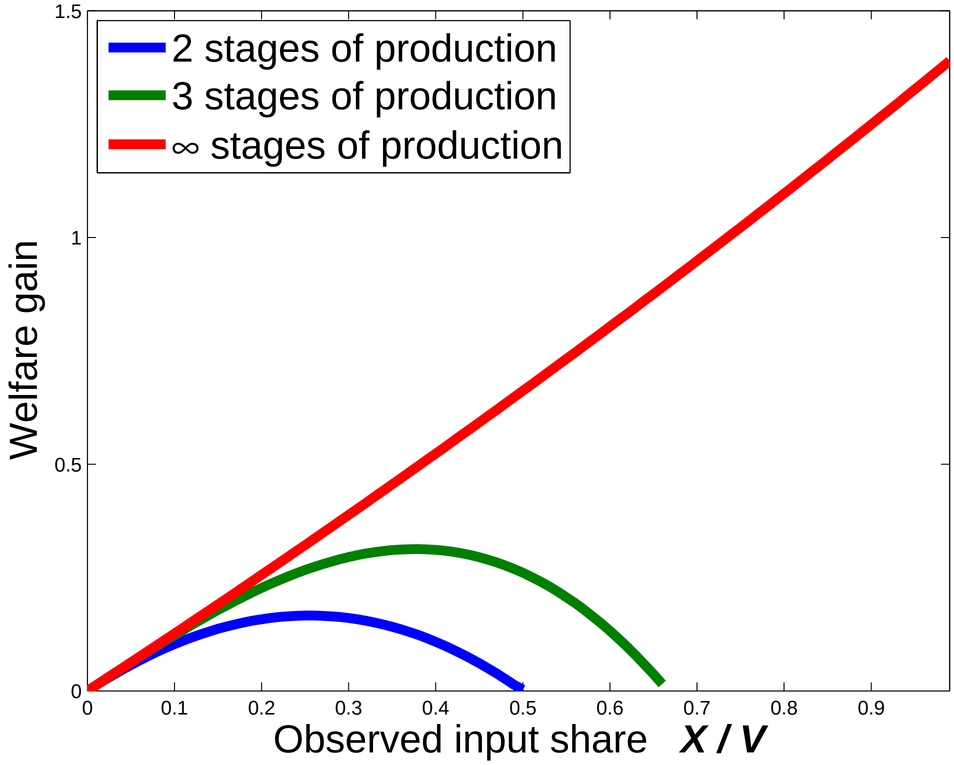

Notes: The figure shows the welfare gain (in percentage points) of removing a turnover tax as a function of the observed input share of an economy in the three versions of the economy described in section 2.2. The blue line shows the welfare gain of the reform when the economy produces in two stages, the green line shows the welfare gain when the economy produces in three stages and the red line shows the welfare gain when the economy produces with an infinite number of stages. All lines are drawn using equation (2.7) when the number of stages is either two, three or it approaches infinity. For every assumed production length, the input elasticities ω are calibrated to the observed input share shown on the horizontal axis, while the tax rate is set to match a value of tax revenue over national income equal to 15%.

Figure 2.6
Welfare gain of the tax reform in the 100 sectors economy

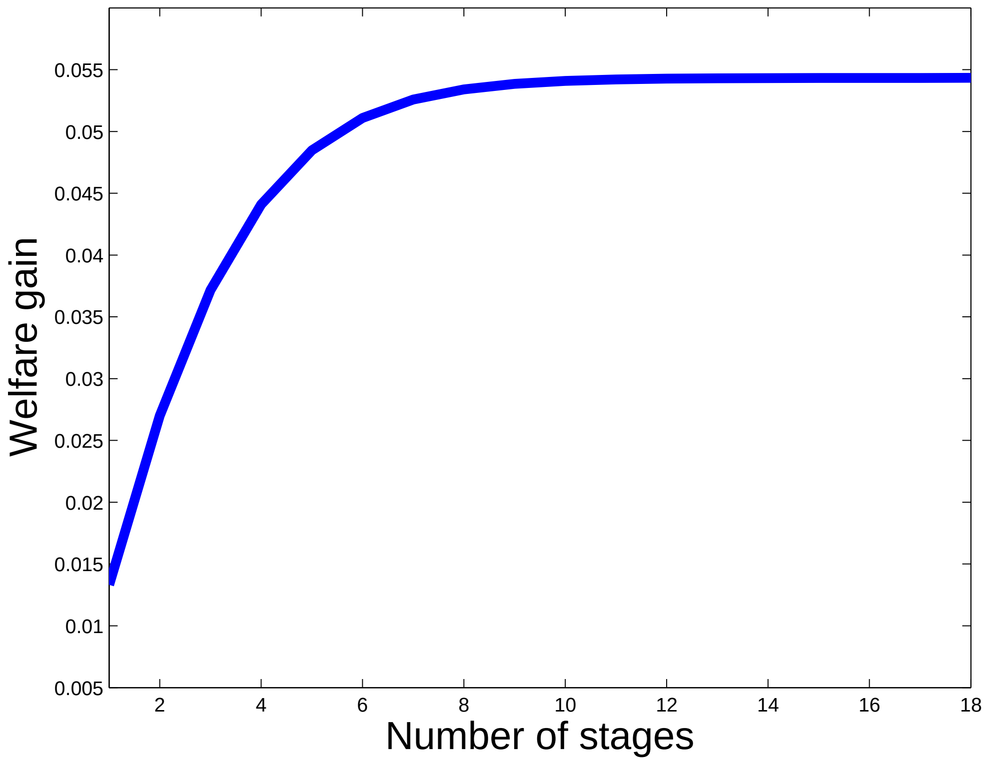

Notes: The figure shows the welfare gain (in percentage points) of the Brazilian 2002-3 reform as a function of the number of stages of production. The line is drawn using equation (2.6) adapted to an economy that produces 100 final goods. For every assumed production length, the input elasticities of the 100×100 matrix ω_{ij}, are calibrated to match the observed Brazilian input-output matrix in 2005, using the procedure described in section 2.4.3. The tax rate is set to match a value of tax revenue over national income equal to 4.6% (the value that is observed in Brazil before the 2002-3 reform) using the procedure described in the same section.

Figure 2.7
Welfare gain of the tax reform in the 100 sectors economy
Comparison with welfare gain in a single sector economy

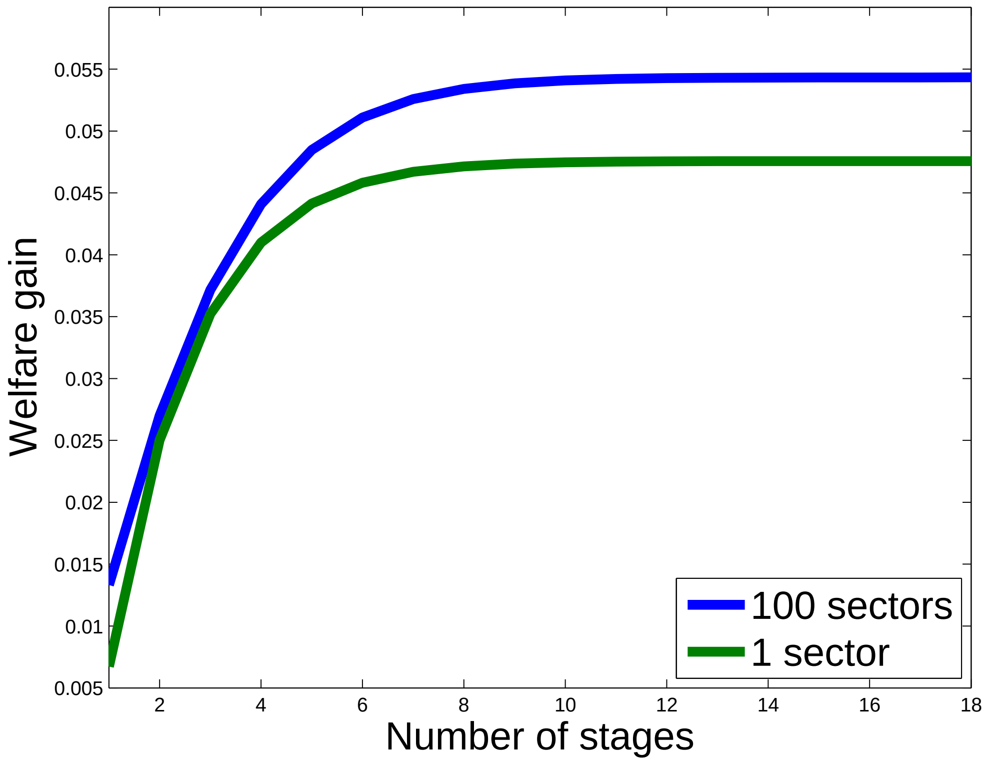

Notes: The figure shows the welfare gain (in percentage points) of the Brazilian 2002-3 reform as a function of the number of stages of production. For every given production length, it plots two lines. The blue line is drawn using equation (2.6) adapted to an economy that produces 100 final goods. The green line is drawn using equation (2.7): the welfare gain of the reform in an economy that produces a single final good. In the first case, the input elasticities of the 100×100 matrix ω_{ij}, are calibrated to match the observed Brazilian input-output matrix in 2005 using the procedure described in section 2.4.3. In the second case the input elasticity of the Cobb-Douglas production function is calibrated to 0.43: the input share observed in the aggregated Brazilian input-output matrix of 2005. In both cases, the tax rate is set to match a value of tax revenue over national income equal to 4.6% (the value that is observed in Brazil before the 2002-3 reform) using the procedure described in section 2.4.3.

Figure 2.8
Welfare gain of the tax reform in the constant elasticity of substitution model **Number of stages and welfare gain**

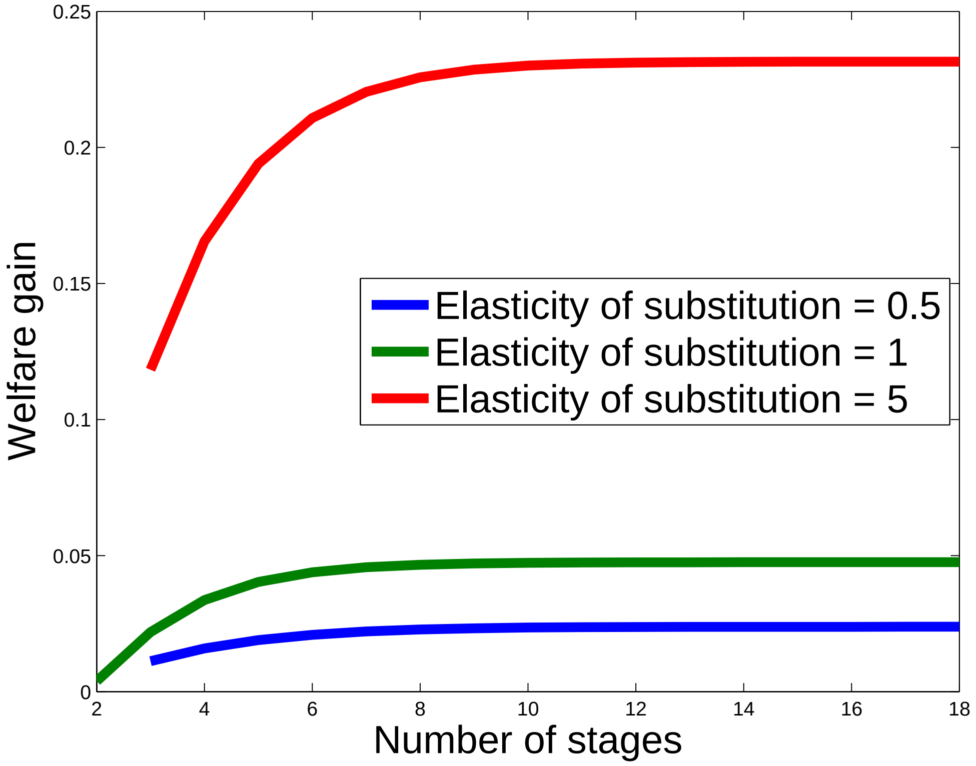

Notes: The figure shows the welfare gain (in percentage points) of the Brazilian 2002-3 reform as a function of the number of stages of production. Welfare gain are calculated using equation (2.15): the welfare gain of the reform in the constant elasticity of substitution model. For every given production length, the figure shows three lines. The blue line is drawn assuming an elasticity of substitution equal to 0.5. The green line is drawn assuming an elasticity of substitution equal to 1: because this corresponds to the Cobb-Douglas case with a single final good, this line reproduces the green line shown in figure 2.7. The red line is drawn assuming an elasticity of substitution equal to 5. For every assumed production length and elasticity of substitution, the joint parameter of the production function γ/A is calibrated using the procedure described in section 2.5.2 to match 0.43: the input share observed in the aggregated Brazilian input-output matrix of 2005. The tax rate is set to match a value of tax revenue over national income equal to 4.6% (the value that is observed in Brazil before the 2002-3 reform) using the procedure described in section 2.5.2.

Figure 2.9
Welfare gain of the tax reform in the constant elasticity of substitution model **Elasticity of substitution and welfare gain**

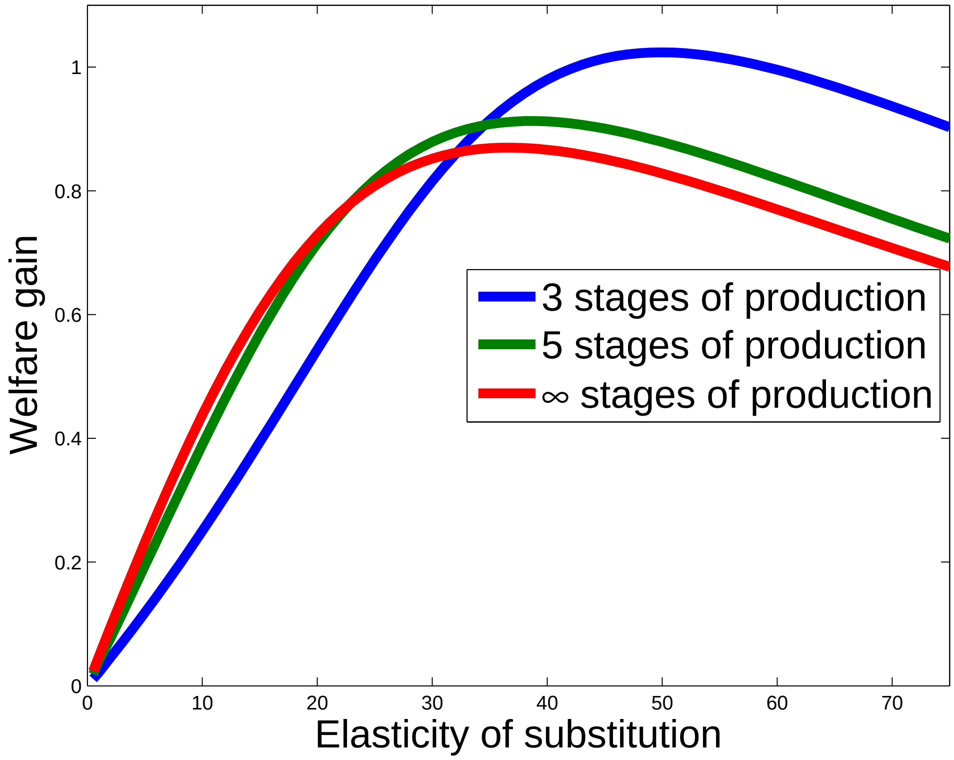

Notes: The figure shows the welfare gain (in percentage points) of the Brazilian 2002-3 reform as a function of the elasticity of substitution between labor and intermediate input. For every given production length, it plots three lines. All lines are drawn using equation (2.15): the welfare gain of the reform in the constant elasticity of substitution model. The blue line is drawn assuming a production process of 3 stages. The green line is drawn assuming a production process of 5 stages. The red line is drawn assuming a production process with an infinite number of stages. For every assumed production length and elasticity of substitution, the joint parameter of the production function γ/A is calibrated using the procedure described in section 2.5.2 to match 0.43: the input share observed in the aggregated Brazilian input-output matrix of 2005. The tax rate is set to match a value of tax revenue over GDP equal to 4.6% (the value that is observed in Brazil before the 2002-3 reform) using the procedure described in the same section.

Appendix A

Analytical derivations of Cobb-Douglas model

In this appendix we derive explicitly equation (2.6): the welfare gains of a reform that converts a turnover tax into a consumption tax in a model with a full input-output matrix.

The indirect utility formula derived in section 2.4.1 implies that the relative change in utility after the reform can be written as:

$$u^v - u^{tt} = p_C^{tt} - p_C^v$$

where a lowercase letter denotes the logarithm of a variable: $x \equiv \log X$ and we use tt and v to label equilibrium variables when taxes are collected on turnover or on consumption.

Under both scenarios P_C is the price index of the Cobb-Douglas aggregate (2.5), which is a function of the prices of the N goods sold at the last stage of production. With a turnover tax equal to t, the price of consumption good i is: $P_i^{tt}(S) = (1+t)MC_i^{tt}(S)$, where $MC_i^{tt}(S)$ is the marginal cost of variety S produced by sector i when there is a turnover tax. With a consumption tax equal to v, the price of consumption good i is: $P_i^v(S) = (1+v)MC_i^v(S)$, where $MC_i^v(S)$ is the marginal cost of variety S produced by sector i when there is no distortion at any point in the production chain. Thus, using the price index of equation (2.5), the welfare gain of the reform can be written as:

$$u^v - u^{tt} = \left[\sum_{i=1}^{N} \beta_i \left(\log(1+t) + mc_i^{tt}(S) - mc_i^v(S)\right)\right] + \log\left(\frac{1}{1+v}\right) \tag{A.1}$$

To derive a solution to (A.1) we proceed in two steps. First, we use technology and market structure to find the solution of the sum in brackets. Next, we use assumption of revenue neutrality to derive an expression for $\log[1/(1+v)]$.

As argued in section 2.4.1, under perfect competition marginal costs and prices of every variety in every sector are uniquely determined by input-output parameters and the tax scheme. This can be derived explicitly starting from stage 1, where the log of the marginal cost of sector i are the same under the two tax schemes: $mc_i^v(1) = mc_i^{tt}(1) = w - a_i$. Given the technology and the market structure, the vector $\boldsymbol{mc}(s)$ that collects the logarithm of marginal costs of all sectors at stage $s > 1$ can be written as:

$$\boldsymbol{mc}^v(s) = \boldsymbol{k} + (\mathbf{1} - \boldsymbol{\omega})w + \Omega\, \boldsymbol{mc}^v(s-1) \qquad s = 2, 3, \dots, S \quad \text{(A.2)}$$

when the tax is collected only from firms producing at stage S as a consumption tax. When there is a turnover tax equal to t, the same vector becomes:

$$\boldsymbol{mc}^{tt}(s) = \boldsymbol{k} + (\mathbf{1} - \boldsymbol{\omega})w + \Omega\, \boldsymbol{mc}^{tt}(s-1) + \Omega\, \boldsymbol{\iota} \log(1+t) \quad s = 2, 3, \dots, S \;\; \text{(A.3)}$$

In both equations $\boldsymbol{k}$ is a vector with typical element equal to: $a_i + (1-\omega_i)\log(1-\omega_i) - \sum_{j=1}^{N} \omega_{ij} \log \omega_{ij}$, $\boldsymbol{\omega}$ is a vector collecting ω_i and $\boldsymbol{\iota}$ is a vector of ones. Solving systems (A.2) and (A.3) from stage 1 forward, yields the solution for $\boldsymbol{mc}^{tt}(S) - \boldsymbol{mc}^v(S)$:

$$\boldsymbol{mc}^{tt} - \boldsymbol{mc}^v(S) = \left[\sum_{s=1}^{S} \Omega^s\right] \boldsymbol{\iota} \log(1+t) \qquad \text{(A.4)}$$

Thus, each term in the sum inside the brackets of equation (A.1) can be written as:

$$\beta_i \log(1+t) \sum_{j=1}^{N} \lambda_{ij}$$

where λ_{ij} the typical entry of the matrix Λ:

$$\Lambda = \left[\sum_{s=0}^{S-1} \Omega^s\right]$$

Next, we derive an expression for $\log[1/(1+v)]$. The condition that the new tax scheme should collect the same amount of tax revenues T^{tt} collected with the turnover tax implies that $[v/(1+v)]\,\overline{L} = T^{tt}$. Hence:

$$\begin{aligned} \frac{1}{1+v} &= 1 - \frac{T^{tt}}{\overline{L}} \\ &= 1 - \frac{t}{\overline{L}} \sum_{i=1}^{N} \sum_{s=1}^{S} MC_i^{tt}(s) Q_i^{tt}(s) \end{aligned} \tag{A.5}$$

Or, in matrix notation:

$$\frac{t}{\overline{L}} \sum_{i=1}^{N} \sum_{s=1}^{S} MC_i^{tt}(s) Q_i^{tt}(s) = \frac{t}{(1+t)\overline{L}} \sum_{s=1}^{S} \boldsymbol{\iota}' \boldsymbol{V}^{tt}(s)$$

where $\boldsymbol{V}^{tt}(s)$ is the vector collecting the value of sales of the firms producing variety s in the N sectors of the economy and $\boldsymbol{\iota}$ is a $N \times 1$ vector of ones. We use 1 to write the value of production at stage s as:

$$\boldsymbol{V}^{tt}(s) = \frac{1}{1+t} \Omega' \, \boldsymbol{V}^{tt}(s+1)$$

Thus, the revenues collected with the turnover tax can be written as:

$$\frac{t}{\overline{L}} \sum_{i=1}^{N} \sum_{s=1}^{S} MC_i^{tt}(s) Q_i^{tt}(s) = \frac{t}{(1+t)\overline{L}} \boldsymbol{\iota}' \left[\sum_{s=0}^{S-1} \left(\frac{1}{1+t} \Omega' \right)^s \right] \boldsymbol{V}^{tt}(S)$$

Given constant consumption shares implied by the Cobb-Douglas aggregator of final goods, $\boldsymbol{V}^{tt}(S) = \boldsymbol{\beta}\overline{L}$. Re-arranging equation (A.5), the tax rate $\log\left[1/(1+v)\right]$:

$$\log\left(\frac{1}{1+v}\right) = \log\left[1 - \frac{t}{1+t} \sum_{i=1}^{N} \beta_i \left(\sum_{j=1}^{N} \lambda_{ji}^{tt} \right)\right] \tag{A.6}$$

where λ_{ij}^{tt} is the generic element of the following matrix:

$$\Lambda^{tt} = \left[\sum_{s=0}^{S-1} \left(\frac{1}{1+t} \Omega' \right)^s \right]$$

Putting together (A.4) and (A.6), one obtains the welfare gain (2.6).

Appendix B

Analytical derivations of CES model

In this appendix we derive equation (2.15): the welfare gain of a reform that converts a turnover tax into a consumption tax in the constant elasticity of substitution model.

We start from the indirect utility function of the household. Because the level of the public good provided by the government and the total number of hours worked by the household are assumed to be unaffected by the reform, the welfare gain of the reform can be written as the log change of the price of the consumption good:

$$u^v - u^{tt} = p^{tt}(S) - p^v(S)$$

Because the price of the consumption good is $P^{tt}(S) = (1+t)MC^{tt}(S)$ and $P^v(S) = (1+v)MC^v(S)$ with a turnover tax and a consumption tax respectively, the welfare gain of the reform can be written as:

$$u^v - u^{tt} = \log\left[(1+t)\frac{MC^{tt}(S)}{MC^{tt}(S)}\right] + \log\left[\frac{1}{1+v}\right] \tag{B.1}$$

We solve for the two terms of this equations in turn.

The system of prices (2.14) allows to write the marginal cost of the variety produced at stage S with the turnover tax and with the consumption tax as:

$$MC^{tt}(S) = \frac{1}{B}\left[\sum_{s=0}^{S-1}\left(\frac{\gamma}{A}(1+t)^{1-\eta}\right)^{s}\right]^{\frac{1}{1-\eta}}$$
$$MC^{v}(S) = \frac{1}{B}\left[\sum_{s=0}^{S-1}\left(\frac{\gamma}{A}\right)^{s}\right]^{\frac{1}{1-\eta}}$$

Solving these sums, the first term in equation (B.1) can be written as:

$$\log\left((1+t)\frac{MC^{tt}(S)}{MC^{v}(S)}\right) = \log\left[(1+t)\cdot\frac{1-\left[\frac{\gamma}{A}(1+t)^{1-\eta}\right]^{S}}{1-\frac{\gamma}{A}(1+t)^{1-\eta}}\cdot\frac{1-\frac{\gamma}{A}}{1-\left(\frac{\gamma}{A}\right)^{S}}\right] \tag{B.2}$$

Next, we use the assumption that total revenues are not affected by the reform to express the second term of equation (B.1) as a function of the parameters of the model. Setting $T^{tt} = T^{v}$ implies:

$$\frac{1}{1+v} = 1 - \frac{T^{tt}}{\overline{L}}$$
$$= 1 - \frac{t}{\overline{L}}\left[\sum_{s=1}^{S} MC^{tt}(s)Q^{tt}(s)\right]$$

By multiplying and dividing every term in the sum in brackets by the after-tax value of all varieties produced in the stages following it, it is possible to write:

$$MC^{tt}(s)Q^{tt}(s) = \left[\prod_{k=s+1}^{S}\frac{\omega^{T}(k)}{1+t}\right]\cdot\frac{P^{tt}(S)Q^{tt}(S)}{1+t}$$
$$= \left[\prod_{k=s+1}^{S}\frac{\omega^{T}(k)}{1+t}\right]\cdot\frac{\overline{L}}{1+t} \tag{B.3}$$
$$\tag{B.4}$$

where the second line follows from the budget constraint and $\omega^{tt}(k)$ is the cost share of intermediate inputs in the production of variety k, as defined in equation (2.16). Using the price equations (2.14), these cost shares can be written as a function of parameters:

$$\omega^{tt}(k) = \frac{\gamma}{A}(1+t)^{1-\eta} \cdot \frac{1 - \left[\frac{\gamma}{A}(1+t)^{1-\eta}\right]^{k-1}}{1 - \left[\frac{\gamma}{A}(1+t)^{1-\eta}\right]^{k}}$$

Plugging this expressions for the cost share into (B.3), it is possible to write each term in the sum inside the brackets of equation (2.15) as:

$$\prod_{k=s+1}^{S} \frac{\omega^{tt}(k)}{1+t} = \left[\frac{\gamma}{A}(1+t)^{-\eta}\right]^{S-s} \left[\frac{1 - \left[\frac{\gamma}{A}(1+t)^{1-\eta}\right]^{s}}{1 - \left[\frac{\gamma}{A}(1+t)^{1-\eta}\right]^{S}}\right]$$

Summing these terms for all stages of production allows to write (2.17). Combining this equation with equation (B.2) derived above we get a formula for the welfare gain of the reform that is a function only of the parameters of the model.

Chapter 3

Team visibility and city travel. Evidence from the UEFA Champions League random draw

3.1 Introduction

Professional teams are often taken to increase their home city's visibility among tourists. As teams get extensive media coverage, their home cities make the news, and this is thought to turn anonymous places into potential tourism destinations. Such potential benefits are often behind sports teams' demands for public subsidies. Yet, so far no analysis has shown that the increased visibility generated by a team brings more visitors into town.

Estimating the effect of team visibility on the attractiveness of a city as a destination is difficult, as both the location of teams and their decision to move are likely to be related to some of the factors that also make cities attractive to tourists (Siegfried and Zimbalist (2000)). I present the first evidence of a positive causal effect of team visibility on city travel by exploiting the random draw of the Union of European Football Associations (UEFA) Champions League soccer tournament, a European competition that takes place once a year. In the first phase of the tournament, participating teams are randomly divided into groups of 4, and have to play a 3-months round-robin tournament with the teams of their group. During this phase, the teams — and their home cities — are likely to get more mention in cities hosting group rivals, and the visibility of a city should therefore increase more in cities hosting teams in the same group than in other Champions League cities. As the groups are formed randomly, this setup allows to compare air travel between cities hosting teams in the same group (treated routes), to air travel between Champions League cities hosting teams assigned to different groups (control routes). The control routes represent a valid counterfactual for treated routes because the teams in these cities could have met, and did not by pure chance.

The group phase of the UEFA Champions League ends in early December and is followed by a break in the competition that lasts until early March. Looking at monthly data of intra-European flights, I find that over the months from January to March cities hosting rivals in the group phase see an increase in monthly air traffic of 5 to 8 percentage points relative to routes across cities whose teams played in different groups. At the mean, this effect implies that between January and March around 2700 more people travel on treated routes. The effect fades away as the tournament proceeds to its final phase, and teams that were initially assigned to the same group no longer play against each other.

The UEFA Champions League brings enormous media attention to the 32 participating soccer teams. Football is by far the most popular sport in Europe, and the UEFA Champions League is the major continental competition for European soccer teams. The competition is covered widely by national and local media, and the games are watched by millions of people across the continent. Participating clubs do not play any other continental competition during the season, and Champions League games are never played on days when domestic league or other continental competitions take place, which means that no other game competes for attention. The first phase of the tournament takes place during fall and lasts for around three months. During this phase participating teams are randomly divided into groups of 4 and teams in the same group play against each other twice. My estimation strategy exploits the random formation of these groups and compares city travel between cities hosting teams in the same group with city travel between cities hosting teams in different groups. The advantage of this approach is that it eliminates the effect of participation in the UEFA Champions League, which is unlikely to be random.

Although the formation of groups is inherently random, I find that treated routes tend to be busier 9 to 3 months before the group phase starts (figure 3.1). I clean the residual heterogeneity across treated and control routes by estimating all regressions with route fixed effects. Figure 3.2 shows that air traffic has identical distribution across treated and control routes once route fixed effects are accounted for. I also estimate dyadic standard errors on incomplete networks of cities (Fafchamps and Gubert (2007a,b)) and find that results are robust to dyadic inference.

Related Literature. This paper relates to the empirical literature that analyzed the impact of professional teams on local communities (Coates (2007)). This literature has examined the effect of sport franchises on the local economy and their contribution to the recovery of depressed urban areas (Baade (1996); Coates and Humphreys (2003)). A recent strand of the literature has also examined the intangible benefits that teams bring to the local population. Carlino and Coulson (2004) argue that higher house rents in cities with NFL franchises compensate for better quality of life, and Coates and Humphreys (2006) show that, when asked to vote on the public funding of a new stadium, people living in areas closer to the proposed site of the facility support the project more than people living farther away.

In this literature, identification of a causal effect always comes from a difference-in-difference approach that exploits the relocation of some professional teams across cities. Since the decision to relocate might be related to other time-varying determinants of cities' economic success, researchers control for existing city-specific linear trends. To the best of my knowledge, this is the first paper that exploits a natural experiment to establish the causal effects of having a sports team in town.

The paper proceeds as follows. Section 3.2 explains the identification strategy and presents the data. Section 3.3 shows that observable variables are balanced across treated and control routes. Section 3.4 presents the main results and section 3.5 proves their robustness. Section 3.6 concludes.

3.2 Identification and data

3.2.1 Identification

If a team increases the visibility of its home town, then we should observe visitors travel to the city after the club appears on high-profile international games. Moreover, the greatest number of visitors should come from cities where rival teams reside, and where the matches are more salient. In this paper I exploit the special structure of the UEFA Champions League to test this second proposition.

The UEFA Champions League brings enormous media attention to participating teams. Football is the most popular sport in Europe, and the UEFA Champions League is the major continental competition for European soccer teams. The group phase is played by 32 European teams divided in 8 groups of 4 teams each and offers to these teams the opportunity to be under the spotlight repeatedly over a period of three months. Access to the group phase is reserved to the teams that performed best in their respective national leagues during the previous season. Once admitted to the group phase, every team is seeded into one of four pots according to its international standing: the eight strongest teams are seeded into pot number 1, the next eight teams into pot number 2, and so on. After seeding, each of the 8 groups is made of exactly one team randomly drawn from every pot, with the provision that teams from the same Football Federation should not play in the same group. The random draw is performed publicly in front of the press at the end of August. Once the groups are formed, between September and December each club plays twice against each of the other 3 teams in its group: once at home and once as visitor. After the conclusion of the group phase in early December, the first two teams of every group advance to the final phase: this has the knock-out format and proceeds from the round of sixteen in March to the final in May.[1]

Group phase games are always played on Tuesday or Wednesday, and they are broadcasted live on national televisions in prime time. These matches, and the media attention that they create in the cities where they take place, is my treatment. If the visibility effect is greater in the cities where opposing teams reside, it is possible to test its existence by regressing visitor arrivals from city j to city i in month m ($V_{ij,m}$) on an indicator of whether the two cities had their team matched in the previous group phase of the Champions League (G_{ij}):

$$V_{ij,m} = \beta_0 + \beta_1 G_{ij} + e_{ij,m} \tag{3.1}$$

In equation (3.1) β_1 is identified consistently because, within the population of routes across cities with at least 1 team in the Champions League, G_{ij} is randomly assigned.[2]

Note that equation (3.1) is a conservative test for the visibility effect, because taking

[1] Appendix C provides additional details on the structure and history of the competition.

[2] Note that compliance is unlikely to be an issue here, because teams assigned to different groups have no occasion to play against each other between September and December, and because only one group phase match was forfeited in the past 15 years: A.S. Roma versus Dynamo Kyiv F.C., scheduled for the 15th of September 2004.

part to the Champions League is likely to increase the visibility of a city in all participating cities, regardless of the group in which they play. Since equation (3.1) compares visitors from cities with a team in the same group to visitors from cities with teams in different groups, it tests for the presence of an effect of being in the same group above and beyond the simple effect of taking part to the same edition of the Champions League.

3.2.2 Data and Sample

I proxy $V_{ij,m}$ in equation (3.1), the number of visitors from city j to city i in month m, with (the logarithm of) $P_{ij,m}$, the number of arrivals from all airports serving city j to all airports serving city i in month m. Since most Europeans travel abroad by plane, the effect of team visibility on air travel is the most relevant effect to estimate, and considering also visitors arriving with different modes of transport are unlikely to alter the results shown here.[3]

Data on all Champions League games comes from the UEFA official website. Data on air travel are available from Eurostat at monthly frequency since 1998, so I will focus on all Champions League editions between the 1998-99 and the 2010-11. I focus only on routes across cities with at least one team playing in the group phase of the Champions League, and exclude routes that could not have been treated. These are the routes across cities that had teams seeded in the same pot and all national routes. I also omit treated and control routes to and from UEFA countries that require a passport and/or a visa to enter (Israel, Serbia, Russia, Turkey and Ukraine). Travel to these countries requires significantly more time and effort than travel within the Schengen Area, and these costs are likely to offset any boost coming from the Champions League.[4]

[3]Among the countries considered, across any pair of countries for which both air and train traffic is available, there was a median of 11.2 air travelers for every passenger arriving by train between 2004 and 2010. For countries that are connected via sea, the median ratio of air to boat arrivals over the same period was 2.9. There are no similar statistics on intra European road traffic, but given that within the countries considered the median ratio of passenger-Km transported by car (by coach) to those transported by train is 11.2 (1.6), air traffic is likely to be at least as important as car travel, and several times more relevant than either train, coach or boat. Notice moreover that these numbers are likely to be lower bounds of the relevance of air traffic within Europe, because they are computed only for country pairs for which a direct connection is active (either by train or via sea). For several country pairs in the sample no such link exists, and on many routes airplanes are simply the only practical mode of transport available.

[4]Results obtained including these routes are available upon request. Inclusion of these routes has two consequences. First, the distribution of variables across treated and control routes is more balanced when

Appendix A gives details on data construction.

3.3 Balancedness

This section documents the balance between treated and control city pairs. The first line of table 3.1 shows that the baseline value of the variable of interest is *not* balanced across treated and control routes. Average arrivals between January and June are 0.21 log points greater on routes that will be treated the following September (p-value = 0.016). Figure 3.1 shows that also the distribution of this variable is different across treated and controls (p-value < 0.001 in a Kolmogorov-Smirnov test). Since between January and June most participants of the Champions League edition that starts in September are still to be decided, these tests suggest that treated routes tend to be busier than control routes even before treatment status is defined.

Although the formation of groups is inherently random, cities that send their teams more often to the Champions League are more likely to have their teams matched together, and at the same time might be richer and have busier routes. In general, if unobservable characteristics of two cities affect both the average air traffic and the likelihood of their teams to meet during the group phase of the Champions League, the estimates of the Champions League effect will be inconsistent. This is especially true when the dependent variable is air arrivals, because air traffic is extremely persistent,[5] and in these cases Bruhn and McKenzie (2009) insist that consistency of estimates is warranted only when also the baseline value of the outcome of interest is balanced across treated and controls.

I address the heterogeneity in the baseline value of air arrivals across treated and controls by exploiting the panel structure of my data. In figure 3.2 I plot the distributions of the residuals of a regression of air arrivals between January and June before the group phase on route fixed effects (FE). The figure shows that the distribution of these residuals in treated routes is very similar to the distribution in control routes, and the Kolmogorov-Smirnov tests can not reject the null of identical distributions (p-value = 0.194). Figure

these routes are included. Second, coefficients from all regressions are less precise and somewhat smaller (but still significant).

[5] In the sample considered, a univariate regression of log arrivals on its value 12 months before explains 93 percent of total variability.

3.2 suggests that including route FE in my specifications allows to estimate consistently the effect on air arrivals of playing in the same Champions League group: for this reason I will only present results from regressions that include routes FE.

The rest of table 3.1 shows that treated and control pairs are very balanced also with respect to other economic, geographic and demographic observables that should correlate with air arrivals. Notice that the units of observation are *pairs* of cities, since both air travel and the treatment are defined over a network of European cities. Network analysis leads to dyadic regression, and appendix B discusses issues of identification and inference that arise in this context. Here it is sufficient to note that the treatment G_{ij} does not vary within a route: when a team from city i plays in the same group of a team from city j the opposite is also true. In this case the balance of treated and control routes with respect to city-specific characteristics must be tested both for the average and for the absolute difference of the variables across the two cities on a route. The intuition is that both the *levels* and the *difference* of these variables may correlate with air travel across the two cities, and consistency is warranted when treatment is uncorrelated with both. Appendix B deals with the details.

Overall treated and control routes are balanced in terms of all observables. With the exception of the test on air arrivals before the group phase starts (first two lines in table 3.1) almost all of the other tests in table 3.1 can not reject the null of identical means across treated and control routes. Out of 60 tests, three are significant at the 10 percent level and only one at the 5 percent level. Also a joint test that allows correlation across these variables cannot reject the null of identical means: when I standardize all variables and run a single regression on the treatment dummy, the coefficient of this dummy is not significant (p-value = 0.284). Since in this stacked regression I use standard errors clustered at the route-year level, this procedure tests for significant differences across treatment and control routes while allowing errors of different variables to be correlated within a single route-year.

By and large these results are reassuring. They prove that treatment is randomly assigned with respect to both time invariant and time varying characteristics, and that also pre trends do not differ significantly across treated and control routes. Finally, not observing all air routes should not be source of concern. Panel B of table 3.1 shows that out of all the routes across cities that could have met but did not, 14.9 percent have non missing air data for both ways of the route during the months of the group phase. This proportion is very similar to the proportion of routes across cities that had two teams playing in the same group (p-value = 0.248).

3.4 Results

3.4.1 Effect on the Month of the Match

I start by documenting the effect of being drawn into the same Champions League group on air arrivals in the month in which the game is played. Groups of supporters usually follow their team when this plays international competitions, so the objective of this section is to show that group-phase matches represent a relevant shock for city-to-city air travel, and that monthly air traffic picks this shock very precisely.

Figure 3.3 plots coefficients γ_l from this regression:

$$\begin{aligned}\log P_{ij,m} \quad = \quad & \alpha_{ij,m} + \delta_t + \mu_m \times t + \sum_{l=-6}^{12} \gamma_l G_{ij,m-l} + \\ & \sum_c \phi^c (c_i + c_j) \times t + \sum_c \psi^c (c_i - c_j) \times t + e_{ij,m} \qquad (3.2)\end{aligned}$$

where $\log P_{ij,m}$ is the logarithm of monthly arrivals and $G_{ij,m}$ is a dummy variable that is equal to 1 if in month m teams from the two cities play a game of the group phase of the Champions League, and 0 otherwise. I add to the regression 6 leads and 12 lags of this dummy ($G_{ij,m-l}$: $l = -6, ..., +12$): figure 3.3 plots the coefficients of these dummies.[6] Inclusion of route-month fixed effects ($\alpha_{ij,m}$) warrants consistency of γ_l. Year fixed effects (δ_t) and month-specific trends ($\mu_m \times t$) improve precision. Time trends specific to the country of origin and to the country of destination ($\sum_c (c_i + c_j) \times t$ and $\sum_c (c_i - c_j) \times t$) enter as suggested by Fafchamps and Gubert (2007b) for directional dyadic regressions.[7]

[6]The number of leads and lags was chosen in order to span half a year before and one year after the event, but figures similar to figure 3.3 can be produced with any number of leads and lags. Standard errors used to compute the confidence interval shown are clustered at the route-month level as in regression (3) below. See the discussion therein for details.

[7]Equation (3.2) is a *directional* dyadic regression because the number of monthly arrivals on the two directions of a route is not identical ($\log P_{ij,m} \neq \log P_{ji,m}$). The way country-specific trends are inserted

On average, each Champions League match increases city-to-city monthly air travel by 7.5 percent (p-value < 0.000). On a monthly average of 16545 arrivals, this equals 1238 passengers more for every match, 0.06 standard deviations or almost 7 full Airbus A320. Note moreover that all coefficients before the match are not significantly different from 0, neither alone nor jointly (F = 1.26; p-value = 0.271): this shows that the effect is not driven by existing pre-trends.

Although figure 3.3 suggests that treated routes stay busier after the game (an F-test on the first 8 lags of G_{ijm-l} rejects the null of no effect with p-value = 0.002), equation (3.2) is inappropriate to test for the visibility effect, because treats each group phase match as a separate event with identical effect after a fixed number of periods, while the "treatment" that matters is likely to be the whole group phase, from September to December. A better test of the visibility effect should consider each month after the group phase separately: this is what next section does.

3.4.2 The Effect after the Group Phase

In this section I present for every month between the end of the group phase and start of following edition estimates of separate regressions of the form:

$$\begin{aligned} \log P_{ij,m} &= \alpha_{ij} + \delta_t + \beta G_{ij} + \\ &\quad \sum_c \phi^c(c_i + c_j) \times t + \sum_c \psi^c(c_i - c_j) \times t + e_{ij,m} \end{aligned} \tag{3.3}$$

where $G_{ij} = 1$ if the teams from the two cities played in the same group in the last edition of the Champions League and the meaning of other symbols is the same as in regression (3.2). Consistency of β is warranted by random assignment of G_{ij} and the inclusion of route fixed effects α_{ij}.

Table 3.2 shows estimates of (3.3) for the month of the match and for all months

---

is intended to impose symmetry on the effect that country trends have on air traffic. To see how this works, take all arrivals to and from Italy ($c = IT$): for this country, ϕ^{IT} captures the average trend of passengers to and from all Italian airports in the sample, while ψ^{IT} picks the average trend of passengers for routes whose origin is in Italy. Symmetry in this context requires that the average trend of passengers for routes that have destination in Italy be equal to minus the average trend for routes with origin there. The terms $\sum_c \phi^c(c_i + c_j) \times t$ and $\sum_c \psi^c(c_i - c_j) \times t$ in equation (3.2) impose this condition.

from January through August. In all regressions standard errors are clustered at the route-month level.[8]

The effect of playing in the same group of the Champions League is large and significant both on the month of the match (row 1) and in the first three months following the end of the group phase. The effect is smaller and not significant by April, when the knock-out phase reaches its most important games (quarter of finals and semi finals), and it disappears before the start of the new season in September. The estimates imply that 996 more passengers fly in January, 775 in February and 921 in March, at the mean number of arrivals for these three months (11794, 12089 and 14798 respectively). Overall, two matches played during the fall attract 2692 visitors during the following three months.

Importantly, the positive effects shown in the three months after the end of the group phase is not confounded by other matches played during the same period. The last game in the group phase is played on the first week of December, and I exclude all routes across cities whose teams play a match during the Champions League knock-out stage either that year or the year before.[9] Since teams are not allowed to take part in more than one international competition a year, teams from the cities considered in the regression had no opportunity to meet until the next season.

[8]The correct unit of clustering is the route-month and not the route as a whole because here correlation of errors within a route is more troublesome than the persistence of the treatment. On the one hand, since the dependent variable varies within a route-month while the treatment does not, estimates of the standard errors are likely to be biased downward (Moulton (1990); (Angrist and Pischke, 2008, p. 313)). On the other hand, the treatment is not persistent at all (the coefficient of a regression of treatment status on its lag on yearly data gives a coefficient of 0.01, p-value = 0.75), and for this reason serial correlation is not an issue as in DD studies (Bertrand et al. (2004)).

[9]Excluding these routes avoids biasing estimates downward. The reason is that knock-out phase matches are a bigger shock to air traffic than group phase matches and teams playing in the same group are less likely to meet again at later stages: the probability to meet a team from the same group is 0 for the round of sixteen, and very low afterwards. Excluding routes across cities that met in the knock-out phase one year before avoids that mean reversion 12 months after the event results in downward bias of the β, because meeting in a knock-out phase in a year is correlated with the probability of ending in the same group the year after. Inclusion of these routes drives estimates downward and worsen precision, but is not crucial for any of the results shown. Moreover, balance of treated and control routes holds also after excluding these routes.

3.5 Robustness

In this section I show that all results shown in section 3.4 are robust to different strategies of inference. Traditional standard errors estimated in regression (3.3) are likely to be biased for several reasons. First, since the dependent variable varies within a route while the treatment does not, correlation of errors within a route is a serious issue (Moulton (1990)). Second, standard errors in (3.3) can be biased because the regression is specified on a network of cities, and shocks hitting a specific city can affect air traffic between this city and all others destinations in the network. Finally, the panel structure of data might complicate further the structure of standard errors, to the point that no single formula is appropriate to estimate them. I deal with these three issues in turn.

Panel A in table 3.3 shows that all results go through when I estimate (3.3) considering a single observation for every route-month and using $\log AvgP_{ij,m} \equiv \log[0.5 \times (P_{ij,m} + P_{ji,m})]$ as dependent variable. This is an alternative to clustering standard errors at route level when the dependent variable varies within a route and the treatment does not (Angrist and Pischke (2008)). Panel B in table 3.3 shows that results are are, if anything, stronger when standard errors are corrected with the formula proposed by Fafchamps and Gubert (2007a,b) for dyadic regressions.

Finally, I assess the overall goodness of the standard error shown in table 3.2 by exploiting the known structure of the Champions League random draw with a bootstrap-like procedure. I randomly create groups of four teams following the same rules UEFA applies to create its groups: after drawing 1000 such placebo treatments, I estimate regression (3.3) for each draw and store the p-values of the β coefficients. The first column of table 3.4 shows the percentage of simulations that had a p-value smaller than 0.05 and suggests that, although standard errors in table 3.2 are downward biased, the bias is small and does not drives the results.

3.6 Summary

Do professional teams make their home towns more visible among tourists? Using a natural experiment embedded in the European Champions League competition, I have shown that cities hosting teams receive more visitors from cities where Champions League games are more salient. My findings provide the first causal evidence that teams have the potential to increase the visibility of their home towns. To the extent that my results carry over to the US, they may help explain Carlino and Coulson (2004) finding that house rents are significantly higher in U.S. cities with NFL franchises. The greater visibility of these cities may increase the flow of visitors, which would bring direct advantages for retail business and hotels for example, and ultimately should increase rents.

In any case, it is important to stress that the visibility effect I find appears to require continuous media exposure. The effect of playing in the same group of the Champions League disappears soon after teams stop playing against each other. This implies that the visibility effect depends on the structure of the competitions. For example, teams participating in leagues with no turnover (such as major American leagues) and teams that have to play more games per season may be more valuable for the visibility of a city.

References

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Baade, R. A. (1996). Professional sports as catalysts for metropolitan economic development. *Journal of Urban Affairs*, 18(1):1–17.

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

Bruhn, M. and McKenzie, D. (2009). In pursuit of balance: Randomization in practice in development field experiments. *American Economic Journal: Applied Economics*, 1(4):200–232.

Carlino, G. and Coulson, N. E. (2004). Compensating differentials and the social benefits of the nfl. *Journal of Urban Economics*, 56(1):25–50.

Coates, D. (2007). Stadiums and arenas: economic development or economic redistribution? *Contemporary Economic Policy*, 25(4):565–577.

Coates, D. and Humphreys, B. R. (2003). The effect of professional sports on earnings and employment in the services and retail sectors in us cities. *Regional Science and Urban Economics*, 33(2):175–198.

Coates, D. and Humphreys, B. R. (2006). Proximity benefits and voting on stadium and arena subsidies. *Journal of Urban Economics*, 59(2):285–299.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of econometrics*, 92(1):1–45.

European Commission (2003). Commission decision of 23 july 2003 relating to a proceeding pursuant to article 81 of the ec treaty and article 53 of the eea agreement. 2003/778/EC. (COMP/C.2-37.398 — Joint selling of the commercial rights of the UEFA Champions League).

Fafchamps, M. and Gubert, F. (2007a). The formation of risk sharing networks. *Journal of Development Economics*, 83(2):326–350.

Fafchamps, M. and Gubert, F. (2007b). Risk sharing and network formation. *American Economic Review*, 97(2):75–79.

Moulton, B. R. (1990). An illustration of a pitfall in estimating the effects of aggregate variables on micro unit. *Review of Economics and Statistics*, 72(2):334–38.

Siegfried, J. and Zimbalist, A. (2000). The economics of sports facilities and their communities. *Journal of Economic Perspectives*, 14(3):95–114.

Figures and tables

Figure 3.1
Arrivals: January-June before Champions League (1998-2010)

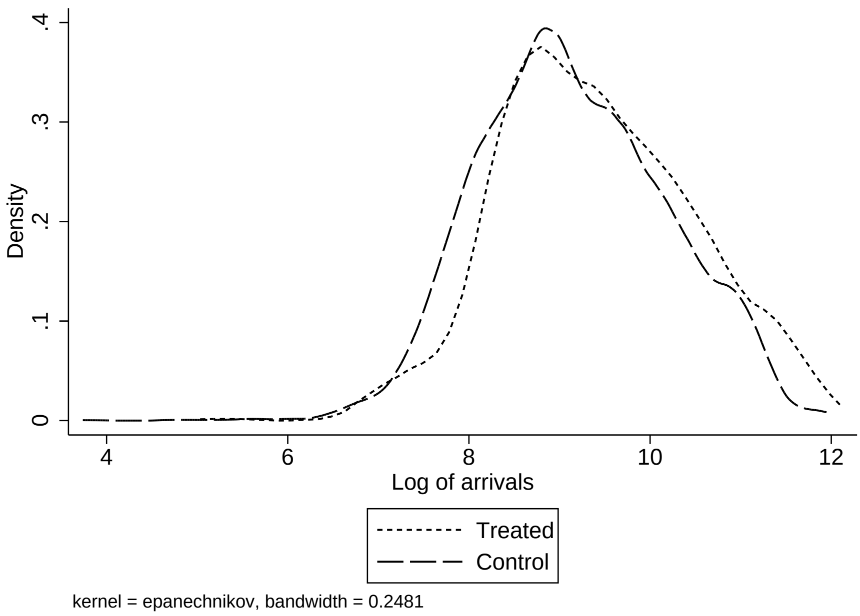

Notes: The figure plots the distribution of arrivals between January and June for all routes across cities that take part to the UEFA Champions League from the following September. There is one observation per route: number of arrivals is the average on the 2 directions. Treated routes have their team playing in the same group from September to December, control routes are routes across cities that could have met but eventually had their teams playing in different groups. I exclude all routes that: (i) could not have been treated given the seeding structure of the random draw and (ii) connect cities in Israel, Russia, Serbia, Turkey or Ukraine.

Figure 3.2
Arrivals: January-June before Champions League (1998-2010)
Residuals from a regression on route fixed effects

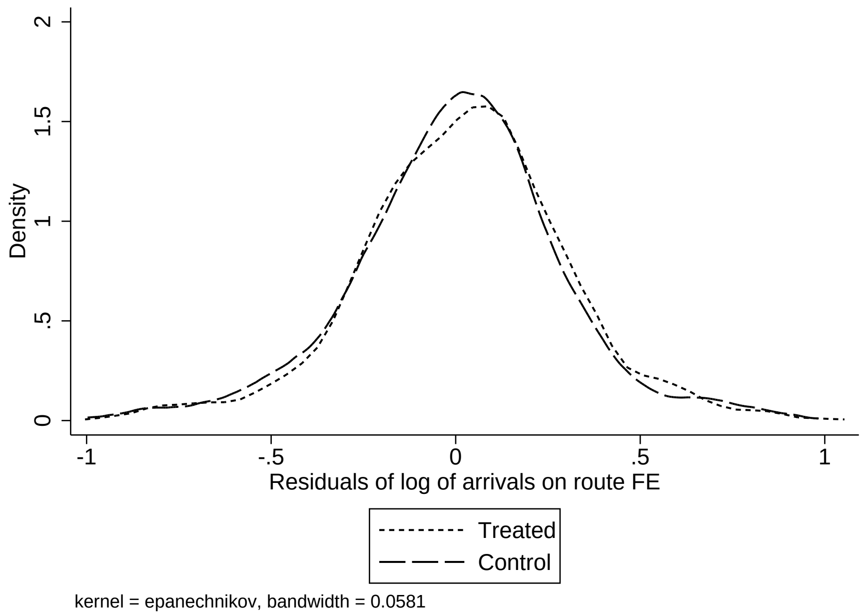

Notes: The figure plots the distribution of residuals from a regression of air arrivals on route fixed effects. Air arrivals are between January and June across all routes across cities that take part to the UEFA Champions League from the following September. There is one observation per route: number of arrivals is the average on the 2 directions. Treated routes have their team playing in the same group from September to December, control routes are routes across cities that could have met but eventually had their teams playing in different groups. I exclude all routes that: (i) could not have been treated given the seeding structure of the random draw and (ii) connect cities in Israel, Russia, Serbia, Turkey or Ukraine.

Figure 3.3
Effect on the month of the match

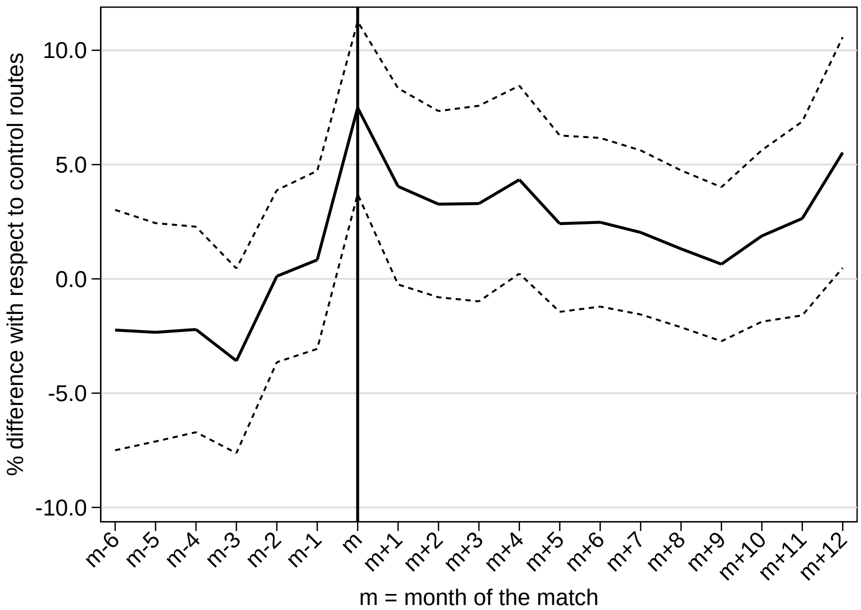

Notes: The figure plots estimates of γ_l from regression (3.2) ($l = -6, -5, \dots, +12$). These coefficients represent the proportional difference in arrivals from city j to city i on treated routes relative to control routes. The coefficient in m (γ_0) is the effect on the month in which a group phase match is played. The figure also plots effects 6 months before the game ($m - 6$) until 12 months after it ($m + 12$). 95 percent confidence intervals calculated using standard errors clustered at the route-month level are reported around the estimates. The number of observations is 17922: these are all routes across cities that had at least 1 team taking part in the Champions League group phase either in the current or in the previous year, but excludes routes that: (i) could not have been treated given the seeding structure of the random draw; (ii) connect cities in Israel, Russia, Serbia, Turkey or Ukraine and (iii) had their teams met in the later stages of the competition either in the current or in the previous edition of the Champions League. The dependent variable ($\log P_{ij,m}$) has the top and bottom 0.5 percent of observations winsorized. Additional controls are year fixed effects, and trends specific to every month and to every country of origin and of destination. See text for details.

Table 3.1
Balance of Treated and Control Routes

Variable	Observations: Treated	Observations: Control	Mean: Treated	Mean: Control	p-value (T=C)
PANEL A — Dependent variable ($\log P$)					
Arrivals (January-June before group phase, logs)					
Average	162	1061	11.1	10.9	0.016**
Absolute difference	162	1061	6.6	6.4	0.008***
Route FE regression residuals (January-June before group phase, logs)					
Average	162	1061	0.01	0.00	0.761
Absolute difference	162	1061	0.01	0.00	0.696
Change in arrivals (January-June before group phase, logs)					
Average	137	933	6.80%	4.40%	0.230
Absolute difference	137	933	10.30%	2.50%	0.081*
PANEL B — Selection					
Routes with non missing air-traffic (September-December)					
	3544	23312	15.60%	14.90%	0.248
PANEL C — Tourism					
Touristic nights by residents (year of the match, logs)					
Average	112	720	15.73	15.66	0.286
Absolute difference	112	720	1.26	1.29	0.746
Touristic nights by residents (1 year before the match, logs)					
Average	109	703	15.7	15.63	0.351
Absolute difference	109	703	1.28	1.3	0.809
Change in touristic nights by residents (year of the match)					
Average	109	703	2.50%	2.20%	0.658
Absolute difference	109	703	9.90%	10.00%	0.964
Touristic nights by non residents (year of the match, logs)					
Average	112	720	15.89	15.85	0.621
Absolute difference	112	720	1.44	1.52	0.465
Touristic nights by non residents (1 year before the match, logs)					
Average	109	703	15.84	15.81	0.586
Absolute difference	109	703	1.41	1.53	0.311
Change in touristic nights by non residents (year of the match)					
Average	109	703	4.20%	3.70%	0.537
Absolute difference	109	703	6.50%	7.60%	0.138

Note: The sample include all treated and control routes for which information on reported variable is available but excludes all routes to and from Israel, Russia, Serbia, Turkey and Ukraine. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table 3.1
Balance of Treated and Control Routes (continued)

Variable	Observations Treated	Observations Control	Mean Treated	Mean Control	p-value (T=C)
PANEL D — Economy					
Income per capita (year of the match, logs)					
Average	8	49	9.8	9.77	0.577
Absolute difference	8	49	0.28	0.28	0.887
Income per capita (3 years before the match, logs)					
Average	17	105	9.81	9.78	0.630
Absolute difference	17	105	0.36	0.41	0.454
Income per capita growth (during the match)					
Average	5	33	1.80%	0.40%	0.548
Absolute difference	5	33	10.20%	13.10%	0.436
Unemployment rate (year of the match)					
Average	54	294	9.10%	8.90%	0.636
Absolute difference	54	294	5.20%	4.30%	0.135
Unemployment rate (3 years before the match)					
Average	69	387	8.60%	8.30%	0.485
Absolute difference	69	387	4.80%	4.40%	0.434
Change in unemployment rate (year of the match)					
Average	35	190	1.10%	0.60%	0.323
Absolute difference	35	190	5.50%	4.20%	0.106
PANEL E — Geography					
Distance (logs)	113	668	6.76	6.76	1.000
Cities are capital					
Both	139	860	18.00%	17.20%	0.823
Only 1	139	860	50.40%	46.50%	0.399
Cities are on an island (e.g. Great Britain, Cyprus)					
Both	139	860	0.00%	0.50%	0.421
Only 1	139	860	29.50%	31.40%	0.654
Cities are in a landlocked country					
Both	139	860	0.00%	0.10%	0.688
Only 1	139	860	10.10%	10.00%	0.979
Cities are in a Mediterranean country					
Both	139	860	14.40%	11.30%	0.290
Only 1	139	860	48.20%	47.60%	0.888

Table 3.1
Balance of Treated and Control Routes (continued)

	Observations		Mean		
Variable	Treated	Control	Treated	Control	p-value (T=C)
PANEL F — Demography					
Population (year of the game, logs)					
Average	95	525	14.11	14.07	0.546
Absolute difference	95	525	1.33	1.33	0.993
Population (3 year before the game, logs)					
Average	91	530	14.1	14.06	0.629
Absolute difference	91	530	1.49	1.43	0.643
Population growth (year of the match)					
Average	75	426	1.70%	2.00%	0.292
Absolute difference	75	426	3.80%	3.50%	0.516
Percent population aged 20 to 35 (year of the match)					
Average	81	475	23.40%	23.30%	0.706
Absolute difference	81	475	4.10%	4.70%	0.116
Percent population aged 20 to 35 (3 years before the match)					
Average	63	358	23.20%	23.30%	0.701
Absolute difference	63	358	3.80%	4.00%	0.538
Change in percent population aged 20 to 35 (year of the match)					
Average	46	269	-0.70%	-0.70%	0.976
Absolute difference	46	269	1.30%	1.40%	0.577
Percent non national, EU residents (year of the match)					
Average	40	205	4.60%	4.40%	0.603
Absolute difference	40	205	5.70%	5.30%	0.644
Percent non national, EU residents (3 years before the match)					
Average	45	225	4.00%	3.60%	0.371
Absolute difference	45	225	5.00%	4.60%	0.636
Change in percent non national, EU residents (year of the match)					
Average	28	153	0.90%	0.70%	0.313
Absolute difference	28	153	1.00%	0.70%	0.063*
Cities speak a romance language (e.g. Italian, Spanish,...)					
Both	139	860	31.70%	26.20%	0.176
Only 1	139	860	41.70%	41.50%	0.962
Cities speak a germanic language (e.g. German, English,...)					
Both	139	860	11.50%	15.30%	0.237
Only 1	139	860	51.80%	43.30%	0.060*

Table 3.2
Effect of playing in the same group of the Champions League

	β	*s.e.*	*Obs.*	Route FE	Year FE & country trends	Month FE & month trends
Month of the match[a]	0.067***	(0.016)	6136	Yes	Yes	Yes
January	0.084**	(0.035)	1344	Yes	Yes	No
February	0.064**	(0.032)	1346	Yes	Yes	No
March[b]	0.062**	(0.028)	1340	Yes	Yes	No
April[b]	0.045	(0.030)	1358	Yes	Yes	No
May[b]	0.042	(0.030)	1354	Yes	Yes	No
June	0.030	(0.023)	1584	Yes	Yes	No
July	0.033	(0.024)	1584	Yes	Yes	No
August	0.019	(0.025)	1584	Yes	Yes	No

Note: The table reports estimates of β in equation (3.3). The sample includes all routes across cities that had at least 1 team taking part in the Champions League group phase during the current edition. I exclude all routes that: (i) could not have been treated given the seeding structure of the random draw; (ii) connect cities in Israel, Russia, Serbia, Turkey or Ukraine and (iii) had their teams met in the later stages of the competition either in the current or in the previous edition of the Champions League. The dependent variable ($\log P_{ij,m}$) has the top and bottom 0.5 percent of observations winsorized. Standard errors in parentheses are clustered at the route-month level. See text for details. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.
[a] September through December. [b] Knock-out phase.

Table 3.3
Robustness for the effect of playing in the same group

PANEL A — Route FE with data collapsed at route-month level						
	β	*s.e.*	*Obs.*	Route FE	Year FE & country trends	Month FE & month trends
Month of the match[a]	0.063***	(0.016)	3068	Yes	Yes	Yes
January	0.084**	(0.041)	672	Yes	Yes	No
February	0.064*	(0.038)	673	Yes	Yes	No
March[b]	0.063*	(0.034)	670	Yes	Yes	No
April[b]	0.045	(0.035)	679	Yes	Yes	No
May[b]	0.042	(0.036)	677	Yes	Yes	No
June	0.030	(0.027)	792	Yes	Yes	No
July	0.033	(0.028)	792	Yes	Yes	No
August	0.019	(0.029)	792	Yes	Yes	No

PANEL B — Route FE with dyadic standard errors						
	β	*s.e.*	*Obs.*	Route FE	Year FE & country trends	Month FE & month trends
Month of the match[a]	0.067***	(0.016)	6136	Yes	Yes	Yes
January	0.084***	(0.031)	1344	Yes	Yes	No
February	0.064**	(0.029)	1346	Yes	Yes	No
March[b]	0.062**	(0.025)	1340	Yes	Yes	No
April[b]	0.045	(0.031)	1358	Yes	Yes	No
May[b]	0.042	(0.030)	1354	Yes	Yes	No
June	0.030	(0.022)	1584	Yes	Yes	No
July	0.033	(0.023)	1584	Yes	Yes	No
August	0.019	(0.024)	1584	Yes	Yes	No

Note: The table reports estimates of β in equation (3.3). The sample includes all routes across cities that had at least 1 team taking part in the Champions League group phase during the current edition. I exclude all routes that: (i) could not have been treated given the seeding structure of the random draw; (ii) connect cities in Israel, Russia, Serbia, Turkey or Ukraine and (iii) had their teams met in the later stages of the competition either in the current or in the previous edition of the Champions League. Panel A reports estimates when observations are collapsed at the route-month level: standard error here are corrected for heteroschedasticity. Panel B reports the same estimates of table 2, but standard errors are dyadic. In both panels the dependent variable has the top and bottom 0.5 percent of observations winsorized. See text for details. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.
[a] September through December. [b] Knock-out phase.

Table 3.4
Results from 1000 placebo simulations of the treatment

	Simulations with p-value < 0.05	*95*th *percentile of estimates in simulations*	β^{a}	*Simulations with* estimates $> \beta^{a}$
	(1)	(2)	(3)	(4)
Month of the match[b]	8.60%	0.038	0.067	0.10%
January	5.20%	0.068	0.084	1.50%
February	6.20%	0.067	0.064	5.70%
March[b]	6.30%	0.061	0.062	4.60%
April[b]	6.20%	0.051	0.045	6.90%
May[b]	5.50%	0.048	0.042	8.40%
June	4.90%	0.039	0.030	10.30%
July	4.10%	0.037	0.033	6.70%
August	4.40%	0.039	0.019	22.50%

Note: Column (1) reports the percentage of simulations in which the effect of a placebo treatment was estimated to be different from 0 at the 5 percent confidence level. The placebo treatment is assigned according to the same rules used to form the Champions League groups; in each simulation I ran a regression identical to (3.3) substituting the true treatment G_{ij} with this placebo treatment randomly assigned. The number of simulations on which these statistics are computed is 1000. Column (2) reports the 95$^{\text{th}}$ percentile of the β estimated on the 1000 simulations of the placebo treatment. Column (3) reports the estimates of β from table 3.2. Column (4) reports the percentage of simulation with an estimated β larger than the one reported in column 3.
[a] From table 3.2. [b] September through December. [b] Knock-out phase.

Appendix A

Data description

I source air traffic data from Eurostat's "Detailed air passenger transport by reporting country and routes" tables.[1] For every airport in each European country, this database contains information on monthly air travel on every route from 1998 to 2010. On every route both arrivals and departures are available. Moreover, two different variables are available both for arrivals and for departures: total number of passengers carried and total number of passengers onboard. Passengers onboard equals passengers carried plus passengers that stop over and proceed to a different destination on the same aircraft, but in practice the two measures are almost identical (between January and August the correlation is 0.9994, p-value < 0.0001). Since for some countries only one between passengers carried and passengers onboard is reported, I pool all available information as follows. There are 4 measures of number of arrivals on each direction of a route: passengers carried and passengers onboard recorded as arrivals in the airport of destination, and passengers carried and passengers onboard recorded as departure in the airport of origin. The dependent variable used in the paper is the simple average of all measures available on every direction of a route. In the regressions shown in table 3.2, 67.9 percent of routes have all 4 measures available between January and August. Using information from the other end of a route cleans some of the noise present at the end of every month, when some passengers are recorded as flying on one month in one airport and on the following one at the other end.[2] In all regressions I use the natural logarithm

[1] Data are available online at: `http://epp.eurostat.ec.europa.eu/portal/page/portal/transport/data/database`.

[2] Although also arrivals from j to i are very correlated with departures from j to i (0.9942, p-value $<$ 0.0001) every month during which the former are greater than the latter are followed by a month in which

of the dependent variable so defined, and I winsorize the top and bottom 0.005 percent of observations to avoid extreme values to drive results.

Both tourism data (night spent in every NUTS 2 region) and demographic and economic data for European cities come from Eurostat.[3] Data on rail, maritime and road travel across and within European countries also come from Eurostat, and refer only to the countries used in the regressions.[4] Data on geographic coordinates of every European city used to compute distances come from Wikipedia. I hand-collected every match of the UEFA Champions League from the 1997-98 to the 2010-11 edition from the UEFA official website.[5]

the opposite happens, by exactly the same number of passengers.

[3] Tourism data are available at: http://epp.eurostat.ec.europa.eu/portal/page/portal/tourism/data/database. Data on cities are available at: http://epp.eurostat.ec.europa.eu/portal/page/portal/region_cities/city_urban/data_cities/database_sub1.

[4] Transport data are available at: http://epp.eurostat.ec.europa.eu/portal/page/portal/transport/data/database.

[5] Online at: http://www.uefa.com/uefachampionsleague/history/.

Appendix B

Dyadic regression

This appendix is based on Fafchamps and Gubert (2007a,b): refer to these papers for details. Both air traffic ($\log P_{ij,m}$) and the Champions League treatment (G_{ij}) are observed on networks in which observations are city-pairs, and every city appears on several different pairs. Regression analysis on network data requires to specify a dyadic model, and both identification and inference need to be adjusted: I discuss these issues in turn.

Identification.—A simpler version of the dyadic regression analyzed in the text takes the form:

$$\log P_{ij} = \alpha + \beta G_{ij} + e_{ij} \tag{B.1}$$

where time subscripts are omitted for simplicity. In (B.1) the treatment G_{ij} is specific to the route across city i and city j: in this case identification does not require any correction. When characteristics specific to the two cities enter a dyadic regression however, it is important that they affect both ways of the route symmetrically: this means that the effect of city characteristic c on air travel must be such that the effect of c_i and c_j on $\log P_{ij}$ is the same as the effect of c_j and c_i on $\log P_{ji}$. In order to impose this symmetry, Fafchamps and Gubert (2007a,b) propose two different solutions, depending on whether the dyadic relationship is directional (as with air traffic, for which $\log P_{ij} \neq \log P_{ji}$) or un-directional (as the average arrivals across a route $\log AvgP_{ij} = \log AvgP_{ji}$). When the relationship is directional, symmetry is imposed by specifying the model:

$$\log P_{ij} = \alpha + \beta G_{ij} + \phi(c_i + c_j) + \psi(c_i - c_j) + e_{ij} \tag{B.2}$$

When the dyadic relationship is un-directional, symmetry is satisfied with:

$$\log AvgP_{ij} = \alpha + \beta G_{ij} + \phi(c_i + c_j) + \psi|c_i - c_j| + e_{ij} \tag{B.3}$$

Regressions (3.2) and (3.3) are directional dyadic regressions (since arrivals from j to i in month m need not be equal to arrivals from i to j during the same period): in these regressions country-trends must enter the equation as in (B.2). The dependent variables in the regressions on data collapsed at the route-month level shown in panel A of table 3.3 and the treatment G_{ij} define un-directional relationships ($\log AvgP_{ij} = \log AvgP_{ji}$ and $G_{ij} = G_{ji}$). For this reason country dummies enter regressions on collapsed data as in (B.3), and for every city-specific variable for which I test the equality of means in table 1 I do so both for the sum and for the absolute difference across the two cities on a route.

Inference.—Standard errors in model (B.1) need to take into account that shocks affecting city i will have an impact on all routes connecting i, and that this is true for all cities on all routes. This implies that in general, for the errors in (B.2) and (B.3) $E(e_{ij}, e_{kl}) \neq 0$ whenever $i = k$ or $i = l$ or $j = k$ or $j = l$. Fafchamps and Gubert (2007a,b) propose to correct the variance-covariance matrix of coefficients in a dyadic regression with a formula similar to the one proposed by Conley (1999) for spatially correlated errors:

$$AVar(\hat{\beta}) = \frac{1}{N-K}(X'X)^{-1}\left(\sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N}\sum_{l=1}^{N}\frac{m_{ijkl}}{2N}X_{ij}e_{ij}e'_{kl}X_{kl}\right)(X'X)^{-1} \tag{B.4}$$

where $\hat{\beta}$ is the $K \times 1$ vector of estimated coefficient, N is the number of observations X is the matrix of all regressors, e_{ij} is the error in equation (B.1) and $m_{ijkl} = 1$ if either $i = k$ or $i = l$ or $j = k$ or $j = l$.

Fafchamps and Gubert (2007a,b) estimate (B.4) on complete networks, i.e. networks in which every node is connected to every other node in the network. However, the air traffic network analyzed here is not complete: first, not every city has a direct connection to every other city taking part in the Champions League (there might exist both a Rome-Lille and a Rome-Valencia route, but no Lille-Valencia connection). Second, even if all routes existed, some routes are not valid “controls” for my treated routes, because cities that had teams in the same pot could not meet, and teams from the same country can not end up in the same group. In order to estimate (B.4) on an incomplete network I coded a new option in the Stata program provided by Fafchamps: this is available on my personal website.

Appendix C

The UEFA Champions League

The first edition of the UEFA Champions League was held in 1955-56. Since then the number of participating teams and the general format of the competition have changed many times. The group phase was introduced in the 1991-92 edition, and since 1994-95 all the matches in the group phase have been played between September and the first week of December. The number of groups has grown overtime, but these have always been formed randomly. Since the 1999-2000 edition there have always been 8 groups.

The rules to admit teams to the group phase vary by country and by year. "Major" leagues send the first 2 or 3 teams of the previous season directly to the group stage. Teams that ended first and second in one of the "minor" leagues, and teams that ended third or fourth in one of the major leagues take part to a "preliminary phase". Official country rankings determine the number of teams that every country can send to the group phase or to the preliminary phase. These rankings are updated by UEFA every season according to 5-year moving average of the performance of national teams in all European competitions. The preliminary phase, played between July and August, consist of a series of knock-out matches that selects 10 of the 32 teams participating in the group phase. These games are not very popular and, since UEFA does not manage directly the TV rights for these games, they are only occasionally broadcasted, even in interested countries (European Commission (2003)).

Note that the format of the competition implies that controlling for year and country fixed effects should improve precision. Year matters because the rules to qualify changed overtime (most notably in 1999, when the number of participating teams became 32, and in 2009, when the rules to access the group phase were renewed). Country specific trends are important because the number of participating teams from any country, and the pots where these teams are seeded depend on national UEFA coefficients, which in turn are updated every year, according to the current and past performance of national teams in UEFA competitions. Country rankings have evolved very differently over the last decade, often trailing domestic economic growth. For this reason, they might correlate with both the probability of treatment and the evolution of air traffic. In order to control for this confounders I include in every specification a set of dummies for both years and country-specific time trends.