Industrialization without Innovation

Paula Bustos
ICREA-UPF, CREI, BSE and CEPR

Juanma Castro-Vincenzi
University of Chicago and NBER

Marti Mestieri
IAE-CSIC, CREI, BSE and CEPR

Joan Monras
UPF, CREI, BSE, and CEPR

Jacopo Ponticelli*
Northwestern Kellogg, NBER and CEPR

August 4, 2026

Abstract

Labor-saving technologies in agriculture can foster structural transformation by releasing workers into manufacturing. The traditional view is that this reallocation generates productivity growth, both statically and, through scale effects, dynamically. However, productivity growth may also depend on how the labor inflow from agriculture reshapes the organization of production in manufacturing. We propose a model where scientists generate ideas that must be implemented by managers, who also supervise production workers. When the inflow is biased toward low-skill workers, managerial bandwidth tightens, inducing firms to reallocate skilled workers from innovation to management, reducing long-run productivity growth despite economies of scale. We test the model exploiting the adoption of genetically engineered soy in Brazil, which released predominantly unskilled workers from agriculture. Consistent with the model, affected regions experienced a rise in the skill premium, a decline in the ratio of innovation to managerial workers, and lower manufacturing productivity. Cross-industry evidence points to the tightening of managerial bandwidth as the primary driver of the productivity decline.

Keywords: Skill-Biased Technical Change, Managerial bandwidth, GE Soy, Brazil.

JEL codes: F16, J43, O13, O14, O33, O41.

*We received valuable comments from Thibault Fally, Manuel García-Santana, Donald Davis, Gene Grossman, Michael Peters, Diego Puga, Sergio Rebelo, Andres Rodriguez-Clare, Chris Tonetti, Chris Udry, Jose P. Vasquez, and seminar participants at CEMFI, CREI, Lugano, Columbia, Northwestern, BdE, the NBER Economic Consequences of Trade, McGill, Princeton Growth Conference, GEP/CEPR Summer Trade Workshop, and the Workshop on Structural Transformation and Macroeconomic Dynamics. This project received funding from the European Research Council (ERC) under the European Union's Horizon 2020 & Horizon Europe research and innovation programmes (grant agreement no. 716388 and no. 101088060). We are also grateful to acknowledge financial support from the UK Department for International Development (DFID) and the Centre for Economic Policy Research (CEPR) (PEDL Scale-Up ERG no. 3787). Monras acknowledges the Fundacion BBVA Beca Leonardo (LEO24-1-12346) and the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2024-001476-S), for their generous support. Mark He and Erick Molina provided excellent research assistance. An earlier version of this paper circulated under the title "Structural Transformation, Industrial Specialization, and Endogenous Growth".

1 Introduction

Early development economists argued that the reallocation of workers from agriculture to manufacturing was fundamental to sustain long-run growth (Lewis 1954, Kuznets 1973). This structural transformation process can lead to higher output because labor productivity is lower in agriculture than in the rest of the economy (Clark 1940, McMillan and Rodrik 2011). In addition, industrialization can lead to higher long-run growth if the manufacturing sector is characterized by economies of scale and knowledge spillovers (Krugman 1987, Lucas 1988, Matsuyama 1992a).[1] However, these arguments typically treat manufacturing as a homogeneous sector and do not consider how the *composition* of the labor inflow shapes the internal organization of production. If workers leaving the agricultural sector are mostly unskilled, the resulting organizational adjustments within manufacturing firms can reduce productivity growth even as the sector expands.

In this paper, we investigate these ideas both empirically and theoretically. We introduce a model of manufacturing development with high-skill workers, who can work either as scientists or as managers, and low-skill workers, who work in production. Scientists generate potential new ideas; managers implement those ideas and supervise production workers. The model features both scale and composition effects. When manufacturing expands in a skill-neutral way, the number of scientists rises while managerial bandwidth remains unchanged, generating the standard scale effect emphasized by Romer (1990). When the expansion is biased toward low-skill workers, however, managerial bandwidth tightens. Firms then devote more high-skill labor to management and less to innovation; fewer ideas are generated and a smaller fraction are implemented. Productivity growth can therefore slow down even as manufacturing employment expands.

We provide empirical evidence in the context of a large increase in agricultural productivity due to the adoption of genetically engineered (GE) soy in Brazil. To identify the causal effects of this new technology on industrial development, we exploit its heterogeneous effects on potential yields across regions with different weather and soil characteristics. We think of these regions as small open economies that were differently exposed to the new technology, which permits to estimate the effects of local agricultural technical change on local labor reallocation, industrial organization, and productivity growth.

To trace the effects of technical change in soy from the agricultural to the manufacturing sector, we rely primarily on annual social security data (RAIS), which covers the universe of formal workers in Brazil and contains detailed occupation information that allows us to identify workers in innovation-intensive and managerial occupations. We complement RAIS

[1] Recent empirical studies analyzing these mechanisms include McCaig and Pavcnik (2013), Atkin, Costinot, and Fukui (2021), Goldberg and Reed (2020), Choi and Levchenko (2021).

with data from the Population Census, which covers both formal and informal workers, and the Annual Industrial Survey (PIA), which provides firm-level measures of manufacturing productivity.

In the first part of the paper, we introduce a stylized model of manufacturing development that embeds innovation into an organization with limited managerial capacity, building on the insight from Garicano (2000) and Garicano and Rossi-Hansberg (2004) that managerial knowledge and time shape the organization of production. The model has two types of workers: high- and low-skilled. High-skilled workers work in two types of tasks: innovation and management. Innovators develop ideas that lead to new input varieties. Input varieties are produced combining a team of managers and low-skilled workers. Managers supervise low-skilled work. Because managerial time is limited, the ability of firms to translate ideas into productive input varieties is constrained by managerial bandwidth, which we proxy by the ratio of managers to production workers. A reduction in this ratio lowers the probability that scientific ideas are successfully implemented, directly reducing productivity.

A flow of workers into manufacturing activates different forces depending on its skill composition. A skill-neutral expansion increases the number of scientists and managers proportionally, leaving managerial bandwidth and the occupational composition of high-skill labor unchanged. More scientists generate more ideas, and those ideas are implemented at the same rate. Hence, skill-neutral structural transformation leads to the classic result of long-run growth (Lewis 1954, Kuznets 1973).

On the other hand, however, a skill-biased inflow of workers into manufacturing affects productivity growth through mechanisms other than the scale effect just described. For example, an inflow of low-skilled workers creates bottlenecks in management, as each manager must supervise more production workers, directly reducing the rate at which ideas are converted into usable technologies. In addition, the tighter bottleneck also raises the marginal value of managerial time relative to scientific work, reallocating high-skill labor from innovation toward management and further reducing the flow of new ideas. Together, these forces can dominate the scale effect and cause productivity to decline even as manufacturing expands. The reverse happens with high-skilled inflows. Hence, the model suggests that whether structural transformation leads to long-run growth or not depends on the nature of the labor inflow, and the subsequent changes in the organization of production activities.

In the second part of the paper, we provide empirical evidence that supports the model's predictions. We start by documenting that GE soy represented a labor-saving and skill-biased technical change in agriculture. Using Population Census data, Bustos, Caprettini, and Ponticelli (2016) showed that microregions more exposed to soy technical change experienced a reallocation of workers from agriculture to manufacturing. We show that the

reallocation was driven predominantly by unskilled workers. Using annual data from RAIS, which allows us to study the dynamics and timing of the effects, we show that low-skill formal employment in manufacturing started increasing around the introduction of GE seeds, while high-skill employment did not change significantly. A one standard deviation higher increase in soy technical change raised formal manufacturing employment by about 9 percent, entirely explained by low-skill workers.

From the point of view of the manufacturing sector, this labor reallocation amounts to an increase in the relative supply of low-skill workers, which we can analyze guided by our model. The model predicts that the skill premium should rise, managerial bandwidth should tighten, and skilled workers should reallocate from innovation to management. We find evidence consistent with each of these predictions. Using composition-adjusted wages from Mincerian regressions on RAIS data, we document that regions with faster soy technical change experienced a significant increase in the skill premium (1.6 percent for one standard deviation). Turning to the organization of production, we find that micro-regions with a standard deviation higher increase in soy technical change experienced a 7% relative decline in manager and a 20% relative decline in scientists relative to low-skill labor costs, respectively. Moreover, we document a large and significant reallocation of resources from scientists to management in affected regions: a one standard deviation higher increase in soy technical change leads to a 14% decline in the wage bill dedicated to scientists relative to managers. Event-type graphs confirm that changes in the organization of production coincide with the introduction of GE soy seeds and the subsequent move of unskilled labor towards manufacturing.

Given the inflow of low-skilled workers and the subsequent reorganization of production, the model has clear predictions for the subsequent evolution of manufacturing productivity. We investigate manufacturing productivity using three different measures: value added per worker, value added per efficiency unit, and total factor productivity. All these measures point to the same result. Despite the fact that soy technical change led to structural transformation, productivity in manufacturing slowed down in the years following GE soy legalization. Microregions with a standard deviation higher increase in soy technical change experienced substantial declines in manufacturing productivity of 7 to 9%.

We propose a two-sector extension of the model to investigate how the productivity decline is distributed across industries. We classify manufacturing industries into high-R&D (H) and low-R&D (L) industries based on R&D intensity from the Brazilian Innovation Survey (PINTEC). Consistent with the extended model, the low-skill labor inflow is absorbed disproportionately by L-industries, with no significant effect on H-industry employment. Crucially, productivity declines *within* both industry types, ruling out a pure between-industry composition story. The TFP decline is larger and statistically significant in L-industries (11.5

percent) compared to H-industries (6 percent). This cross-industry pattern identifies the implementation bottleneck as the dominant mechanism: managerial bandwidth is most severely diluted in L-industries, which absorb the bulk of the low-skill inflow, and this direct effect on idea implementation outweighs the reallocation of scientists away from H-industries.

We test directly this channel by exploiting the longitudinal structure of RAIS. We follow individual innovators employed in the L-industry in the pre-period and classify them each year along two dimensions: their occupation (innovation vs management) and their industry (the same L-industry as at baseline, or a different industry). We focus on four trajectories: remain innovators in the L-industry, become innovators in other sectors, transition to managerial occupations in the L-industry, or transition to managerial occupations in other sectors. Consistent with the two-sector model, the bulk of the reallocation is concentrated in the third trajectory: in microregions with a one-standard-deviation higher increase in soy technical change, the share of the wage bill paid to innovators who remained in the L-industry and transitioned into managerial occupations rises by 2.6 percentage points. We also find that innovators are retained in the L-industry rather than moving to other sectors, consistent with the complementarity between high- and low-skill workers in the industries absorbing the inflow. The same exercise applied to innovators initially employed in H-industries shows no such reallocation. This evidence provides direct support for the central mechanism of the model, namely, that low-skill labor inflows reshape the internal organization of firms by pulling skilled workers out of innovation and into management.

Our findings suggest that different forces driving structural transformation can lead to different types of industrial development. In most countries, the process of labor reallocation from agriculture to manufacturing can be ascribed to one of two forces: "push" forces, such as new agricultural technologies that push workers out of agriculture, or "pull" forces, such as industrial productivity growth, that pull workers into manufacturing. We show that when labor reallocation from agriculture to manufacturing is driven by "push" forces that affect disproportionately the least skilled workers, the resulting organizational adjustments within manufacturing firms can slow down aggregate productivity growth.

In this sense, our results are a cautionary tale on the effects of structural change on productivity growth. The adoption of new technologies in agriculture may result in static productivity gains in the agricultural sector but dynamic losses in manufacturing productivity. The mechanism we identify suggests that these losses arise not from industrialization per se, but from the skill composition of the labor inflow and its effects on the internal organization of manufacturing firms. We think that the experience of Brazil documented in this paper is informative for low- to middle-income countries where a large share of the labor force is employed in agriculture, and who import new agricultural technologies from more

developed countries.

Related Literature

Our paper is related to the classic literature arguing that labor reallocation from agriculture to manufacturing increases aggregate productivity. A first set of papers in this literature argues that, because of the sizable productivity and wage gaps between agriculture and manufacturing in developing countries, there are large static productivity gains when countries industrialize (e.g., Caselli 2005, Restuccia, Yang, and Zhu 2008, Lagakos and Waugh 2013, Gollin, Lagakos, and Waugh 2014).[2] Second, there can also be dynamic productivity gains when labor reallocates towards manufacturing if manufacturing is subject to agglomeration externalities and knowledge spillovers (Krugman 1987, Lucas 1988, Matsuyama 1992a).[3] Our paper contributes to this literature by empirically showing how the impact of structural transformation on manufacturing productivity growth depends on the nature of the structural transformation process and how this affects industrial development. To make this point, we build on insights from the organizational and the endogenous growth literatures. In particular, we follow Garicano (2000) and Garicano and Rossi-Hansberg (2004) by incorporating organizational structure into production. We also draw on ideas from Grossman and Helpman (1991) who study open economy endogenous growth models and show how growth may depend on the relative cost of innovation.[4] As a result, we obtain a parsimonious framework, new in the literature, that allows for both scale economies and composition effects to shape manufacturing productivity growth.

Recent empirical studies exploring the theoretical mechanisms highlighted above include Franck and Galor (2019) who use historical variation in the diffusion of steam engines to argue that early adoption of technologies that do not lead to human capital development hinders long-run growth. In a similar vein, McMillan and Rodrik (2011) and McMillan, Rodrik, and Sepulveda (2017) emphasize that structural change can be growth-enhancing or

[2]More recently, Herrendorf and Schoellman (2018) measure and compare agricultural wage gaps in countries in different stages of the structural transformation process. They find that the implied barriers to labor reallocation from agriculture are smaller than usually thought in the macro-development literature, and argue that labor heterogeneity and selection are important drivers of such gaps.

[3]Recent evidence suggests that this channel may be operative in some circumstances. Peters (2022) uses the displacement of Eastern Germans towards Western Germany to show that places experiencing larger population growth specialized in manufacturing and saw GDP per capita grow over the long run. This channel, however, may not be the norm. Herrendorf, Rogerson, and Valentinyi (2022) find limited evidence in support of dynamic productivity gains from labor reallocation from agriculture into manufacturing using cross-country data for the period 1990-2018. See also De Vries, Timmer, and De Vries (2015) which proposes a decomposition of the effect of structural change on the productivity growth of Sub-Saharan African countries between static and dynamic components.

[4]Some of our findings can be rationalized with a small open economy growth model with three sectors – agriculture, low- and high-skill manufacturing, and two factors, where R&D investment depends on the relative size of the high-skill industries. In this sense, our paper is also related to the classical literature on endogenous growth, which was mainly theoretical, see Romer (1990); Aghion and Howitt (1992, 2008).

growth-reducing depending on the correlation between changes in employment shares and productivity levels, comparing industrialization experiences across a number of countries. Recent studies of the effects of structural transformation on productivity include McCaig and Pavcnik (2013) who document how trade-induced labor reallocation from informal to formal manufacturing increased productivity in Vietnam; and Imbert, Seror, Zhang, and Zylberberg (2020) who exploit short-run agricultural shocks in China to document how migration from rural to urban areas reduces labor costs for firms leading to an expansion in labor usage and a reduction in capital-biased technology adoption. Relative to this work, we provide evidence on how labor-saving agricultural productivity shocks can affect the patterns of industrial organization and, as a result, manufacturing productivity growth.

Our paper also builds on the empirical literature studying the effects of agricultural technical change, particularly the papers that argue that technological advancements in agriculture are skill-biased. For instance, Foster and Rosenzweig (1996) study the effects of the introduction of high-yield varieties in India, and show that technological innovations in agriculture increased the relative demand for skill in agriculture and, thus, returns to primary schooling.[5] We contribute to this literature by showing that the recent introduction of GE soy was also skill-biased. More importantly, we focus on the implications of skill-biased agricultural technical change for industrialization, which have not previously been explored.

Finally, it is worth noting that this paper is part of a broader research agenda that studies the effects of agricultural productivity on structural transformation in the context of the adoption of GE crops in Brazil. A first study in this agenda, Bustos et al. (2016) shows that, if agricultural technical change is labor-saving, increases in agricultural productivity can lead to a reallocation of labor towards the manufacturing sector, even in open economies.[6] Relative to Bustos et al. (2016), our paper studies how the reallocation of workers from agriculture to manufacturing can shape the growth prospects of the industrial sector. A second study in this agenda analyzes the effects of the agricultural boom in Brazil on capital markets. Bustos, Garber, and Ponticelli (2020) document that regions with faster technical change in soy experienced an increase in local savings deposits which were mostly not lent locally, leading to an increase in capital outflows. Banks redirected agricultural savings to urban areas outside soy-producing regions where they were invested in the manufacturing

[5] In related work, Bragança (2014) shows that investments in soybean adaptation in Central Brazil in the 1970s induced positive selection of labor in agriculture.

[6] As noted by Matsuyama (1992b), the mechanisms through which agricultural productivity can speed up industrial growth proposed in the classical models of structural transformation are only operative in closed economies, while in open economies high agricultural productivity induces a reallocation of labor toward agriculture, the comparative advantage sector. Bustos et al. (2016) proposes a model where the effect of agricultural productivity on industrial development in open economies depends on the factor-bias of technical change.

and service sectors. Those findings describe the effects of agricultural technical change on structural transformation through a capital supply channel and are consistent with a high level of financial integration across regions. In contrast, the current paper documents the effects of agricultural technical change through a labor supply channel which operates in local labor markets with limited migration responses.

2 Model

We develop a model to study how the size and *composition* of manufacturing employment affect aggregate productivity growth. The model builds on two central mechanisms in growth theory. First, the scale effects of the canonical frameworks of Romer (1990) and Aghion and Howitt (1992): a larger manufacturing sector makes innovation more profitable and increases productivity growth. Second, human capital effects emphasized by Grossman and Helpman (1991) in an open economy setting: if innovation is more skill-intensive than manufacturing production, an expansion of manufacturing production driven by a larger supply of unskilled labor can crowd out innovative activities because high-skill workers are reallocated toward production tasks.

Our framework integrates these classic mechanisms in a model in which, in addition, the successful implementation of new technologies requires skilled managers to train unskilled workers. Importantly, knowledge transmission is decreasing in the number of unskilled workers per manager, leading to decreasing returns in technology implementation. The central implication of the model is that structural transformation driven by inflows of low-skill labor into manufacturing can reduce manufacturing productivity, even as manufacturing employment expands. This mechanism contrasts with the cases of skill-neutral or high-skill labor inflows, which yield opposite conclusions.

In what follows, we present the baseline one-sector model and derive the key comparative statics results that guide our empirical analysis (Sections 2.1 and 2.2). Next, we extend the model to two sectors to generate industry-level predictions which will allow us to empirically identify the main theoretical mechanisms at work in the data (Section 2.3). Appendix A contains detailed derivations and proofs.

2.1 Environment

We consider a collection of small open economies with a representative manufacturing firm in each location. Firms sell differentiated products in international markets under monopolistic competition and choose employment of scientists (S), managers (M), and production

workers (L). We characterize the equilibrium in one such small open economy, without explicitly indexing locations.

Production Technology The production function for final goods is:

$$Y = A\left(\int_0^N y_i^{\frac{\sigma-1}{\sigma}}\,\mathrm{d}i\right)^{\frac{\sigma}{\sigma-1}},$$

where y_i are differentiated intermediate inputs which are combined into the final good and $\sigma > 1$ is the elasticity of substitution across varieties of intermediates. Each intermediate input is produced using a Cobb–Douglas technology that combines low-skill workers and managers:

$$y_i = \tilde{A}\, l_i^{\theta} m_i^{1-\theta}, \qquad \tilde{A} = \theta^{-\theta}(1-\theta)^{-(1-\theta)}. \tag{1}$$

By symmetry, we have $y_i = y$ for every i. Hence, final output can be expressed as

$$Y = A\, N^{\frac{1}{\sigma-1}} M^{1-\theta} L^{\theta}.$$

The N intermediates are varieties of differentiated production inputs that facilitate a finer division of labor, thereby enhancing overall productivity. New ideas for intermediates are generated by scientists and implemented by managers, as described below.

Innovation Technology Scientists generate a stock of potentially implementable ideas according to

$$I = S^{\phi}, \qquad \phi > 0.$$

Managers perform two distinct tasks: production and innovation. For the first task, managers work together with low-skill workers in the production of intermediates, as specified in equation 1. Note that the production task is standard in the sense that an increase in the number of low-skill workers raises the marginal productivity of each manager in the production of already adopted technologies. For second task, managers teach low-skill workers how to manufacture new intermediates developed by scientists. In contrast to the production task, the innovation task requires the transmission of knowledge, which becomes less effective as the number of low-skill workers per manager increases.

In particular, we assume that each idea i has a different level of complexity Z_i and is usable in production if and only if managerial training time per worker exceeds its complexity, i.e., $M/L \geq Z_i$. In that case, the share of ideas successfully implemented is increasing in the number of managers per worker, which we call managerial bandwidth: $B \equiv M/L$. Let $q(B)$ denote the probability that a given idea is successfully implemented. We assume that

$q'(B) > 0$, i.e., more managers per worker increase implementation success.[7] For tractability, we adopt the logistic probability functional form

$$q(B) = \frac{B^\chi}{1 + B^\chi}, \qquad \chi > 0.$$

The parameter χ governs how responsive implementation success is to managerial resources. However, the predictions of the model do not depend on this specific functional form.[8]

Given these assumptions, the mass of implemented ideas is

$$N = S^\phi \, q(B). \tag{2}$$

In words, the number of implemented ideas is the total stock of potential ideas, S^ϕ, times the probability of implementation success, $q(B)$, since all ideas are drawn independently from the same distribution.

Occupational Choice High-skill workers can be employed as scientists or managers, while low-skill workers can only perform production work.

Demand and Profits Each firm sells differentiated final products on international markets under monopolistic competition. Demand in global markets is isoelastic with elasticity ϵ. A firm chooses (S, M, L) to maximize

$$\Pi = Y(S, M, L)^\mu - w_H(S + M) - w_L L,$$

with $\mu = 1 - 1/\epsilon$.

Equilibrium allocations and prices Using the profit function above and the production functions for intermediate goods and ideas, firms choose the optimal number of workers L,

[7] An alternative microfoundation is one in which managers allocate time to implementation and production, $M_{\text{imp}} + M_{\text{prod}} = M$, and implementation success depends on implementation intensity per worker, $q(M_{\text{imp}}/L)$. For a broad class of concave success functions, the optimal split implies that implementation intensity is increasing in overall managerial resources per worker; our reduced-form $q(B)$ summarizes (and locally approximates) this mapping.

[8] The elasticity of implementation success with respect to bandwidth is $\frac{\partial \ln q(B)}{\partial \ln B} = \chi\,[1 - q(B)]$, which is increasing in χ for given $q(B)$. For $\chi \leq 1$, q is globally concave in B, so there are diminishing returns to bandwidth everywhere. For $\chi > 1$, q is S-shaped: returns to bandwidth are initially increasing and then diminishing, with an inflection at $B^\chi = (\chi - 1)/(\chi + 1)$. Empirically, $B = M/L$ is well below one (in our calibration $\chi = 1$ and $B \approx 0.15$), so diminishing returns hold throughout the relevant range.

managers M, and scientists S:

$$\max_{S,L,M} A^{\mu}\big[\,S^{\phi}q(B)\big]^{\kappa\mu}L^{\theta\mu}M^{(1-\theta)\mu} - w_H(S+M) - w_L L,$$

with $\kappa = \frac{1}{\sigma-1}$. The first-order conditions imply

$$\frac{w_H S}{\mu Y^{\mu}} = \kappa\phi, \tag{3}$$

$$\frac{w_H M}{\mu Y^{\mu}} = (1-\theta) + \kappa\chi\big(1-q(B)\big), \tag{4}$$

$$\frac{w_L L}{\mu Y^{\mu}} = \theta - \kappa\chi\big(1-q(B)\big). \tag{5}$$

These conditions have a simple interpretation. The cost share of scientists, described by Equation (3), is constant and increasing in both the degree of input differentiation, $\kappa = 1/(1-\sigma)$, and the elasticity of new ideas for input varieties to scientists, ϕ. This is because when scientists are more productive at generating new ideas and input varieties are more differentiated, the demand for scientists is higher. In turn, the managerial cost share, as characterized by Equation (4), reflects both the contribution of managers to current production and their contribution to innovation. The former is represented by the standard Cobb–Douglas factor share, $(1-\theta)$. The latter component increases when inputs are more differentiated and when the productivity parameter governing managers' effectiveness in implementing ideas (χ) is high. Conversely, when managerial time per worker (B) is scarce, the probability of successful implementation of new ideas is low, which implies high marginal returns to hiring additional managers. Finally, the last condition, Equation (5), characterizes the cost share of low-skill labor, which falls below the standard Cobb–Douglas benchmark (θ). This deviation arises because additional low-skill labor tightens managerial bandwidth, thereby lowering the probability of successful implementation of ideas and, consequently, the effective marginal productivity of low-skill labor.

Factor Market Clearing and Key Equilibrium Objects Let $\bar{H}$ and $\bar{L}$ denote the total endowments of high- and low-skilled workers, respectively. Market clearing implies $S + M = \bar{H}$ and $L = \bar{L}$. Optimality requires that the marginal revenue product of each input equal its corresponding wage, with high-skilled workers assigned between S and M to equalize their marginal returns. Combining first-order conditions with factor market clearing

yields:

$$B = \frac{\bar{H}}{\bar{L}} \frac{(1-\theta)+\kappa\chi[1-q(B)]}{(1-\theta)+\kappa\chi[1-q(B)]+\kappa\phi}, \tag{6}$$

$$\frac{S}{M} = \frac{\kappa\phi}{(1-\theta)+\kappa\chi(1-q(B))}. \tag{7}$$

Equation (6) defines equilibrium managerial bandwidth as the unique solution to a fixed-point equation. Since implementation success $q(B)$ increases with managerial bandwidth, the ratio on the right-hand side of Equation (6) decreases in B, while the left-hand side increases in B. The equilibrium is therefore unique. In addition, because the right-hand side is proportional to $\bar{H}/\bar{L}$, an increase in the supply of low-skill labor shifts it downward and lowers the equilibrium value of B.

Similarly, Equation (7) indicates that the equilibrium ratio of scientists to managers depends on their relative returns. The denominator in Equation (7) represents the effective return to managers: it comprises the managers' direct contribution to production, as well as their contribution to innovation, captured by $\kappa\chi[1-q(B)]$. Note that as managers train additional workers to produce new varieties, a reduction in their bandwidth B increases their marginal value relative to scientists. Consequently, an increase in the supply of low-skill labor, $\bar{L}$, lowers the equilibrium bandwidth and the S/M ratio, implying a reallocation of high-skill labor away from scientific tasks and toward managerial activities. In the next subsection, we provide a formal analysis of these mechanisms.

2.2 Manufacturing Productivity and Labor Inflows

This section analyzes the effects of changes in labor supply on aggregate productivity and the organization of scientific, managerial, and productive activities. We interpret the equilibrium changes in the mass of implemented varieties as changes in productivity growth, following the standard variety-expansion logic in Romer (1990). We define $N^{\kappa} = [S^{\phi}q(B)]^{\kappa}$ as total factor productivity (TFP), since it corresponds to the efficiency gains arising from the expansion in product variety, as formalized in Equation (2).

First, let us note that, in line with the frameworks of Romer (1990) and Lucas (1988), our model features scale effects as key drivers of productivity growth. When labor supply expands without changing the skill composition of the workforce, the model delivers standard scale effects. A proportional increase in production employment raises the number of scientists one-for-one, increases the mass of ideas, and leaves organizational constraints unchanged. As a result, a larger manufacturing sector accelerates innovation and increases aggregate productivity, as in canonical growth models. This benchmark case is formally characterized

in Proposition A2 in Appendix A and provides a useful point of comparison for the analysis that follows.

We now turn to the central case of interest, in which manufacturing expansion is driven by an inflow of low-skill labor, holding the supply of high-skill workers fixed. The next Proposition demonstrates that, in this case, additional mechanisms arise beyond scale effects.

Proposition 1 (Comparative Statics with respect to $\bar{L}$)**.** *Holding $\bar{H}$ fixed, an increase in the supply of low-skill labor $\bar{L}$ implies the following changes in equilibrium outcomes:*

1. *The skill premium, w_H/w_L, increases;*
2. *Managerial bandwidth, $B = M/L$, decreases;*
3. *High-skill labor reallocates toward management, i.e., M/S increases;*
4. *Total factor productivity (TFP), N^κ, falls:*

$$\frac{d\ln \mathrm{TFP}}{d\ln \bar{L}} < 0, \qquad \frac{d\ln(Y/\bar{L})}{d\ln \bar{L}} < 0.$$

5. *Total output Y may either rise or decline, depending on which of the two effects—the scale effect or the labor-force skill-composition effect—dominates.*

To sum up, this proposition states that when the relative supply of low-skilled workers rises, the skill premium increases. As a result, both managers and scientists become expensive relative to production workers. Firms respond by reducing the demand for managers and scientists relative to unskilled workers. This implies that firms operate with fewer managers per worker, or lower bandwidth (B), which generates a reduction in the probability of successful implementation of ideas, $q(B)$. In an effort to improve idea implementation, firms reallocate high-skill workers away from innovation into managerial tasks.

The reorganization of production described above has direct consequences for aggregate productivity. First, firms operate with fewer managers per worker, which reduces the effective stock of implemented varieties, $N = S^\phi q(B)$, and thereby lowers aggregate productivity. Second, the reallocation of high-skill workers toward managerial tasks leads to a reduction in innovation, S, further diminishing aggregate productivity. To clarify the mechanisms at play, we derive a simple expression characterizing the impact of an inflow of low-skilled labor on aggregate productivity:

$$\frac{d\ln \mathrm{TFP}}{d\ln \bar{L}} = \left[1 + \underbrace{\frac{\phi\,\kappa\chi\,q(B)}{(1-\theta)+\kappa\chi(1-q(B))+\kappa\phi}}_{\text{Innovation effect}}\right] \times \underbrace{\kappa\,\chi\ \times(1-q(B))\times\frac{d\ln B}{d\ln \bar{L}}}_{\text{Implementation effect}}$$

The term labeled as *implementation effect* represents the *direct effect* of an increase in the supply of low-skill labor: the managerial bandwidth B contracts in response to the inflow of low-skill workers, which in turn lowers the probability of successful idea implementation, $q(B)$, and thus diminishes the number of effective varieties. The second term, labeled as *innovation effect*, reflects the endogenous reallocation of high-skill labor between research and management triggered by the reduction in the marginal product of scientists relative to managers. Both effects combine to reduce the number of new varieties and thus TFP.

Finally, we analyze the case of an increase in the supply of high-skill labor which relaxes managerial constraints and leads to opposite comparative statics than Proposition 1 in Appendix A.

2.3 Two-Sector Model

In this section, we extend the baseline model to two industries to understand how the mechanisms discussed above translate into systematic cross-industry heterogeneity that we can explore in the data. We describe the model in detail in Appendix B and summarize the setup and its main implications below.

We consider an empirically motivated partition of the manufacturing sector in two industries: H and L. We assume that production and innovation in each of the two industries follow the same structure described above for the one-sector baseline model, but one of the industries is more R&D and skill-intensive. More specifically, the H-industry has a higher elasticity of ideas with respect to scientists, $\phi_H > \phi_L$ while the L-industry has a higher low-skill production share, $\theta_L > \theta_H$. While we allow for both differences in innovation and skill intensity in this exposition, we show in Appendix C that differences in R&D intensity map into differences in effective skill intensity when innovation is skill-biased. Therefore, we view differences in innovation intensity as the fundamental difference across industries and use this dimension in our empirical test of the model predictions.

The next proposition summarizes the effect of an expansion in low-skill labor supply, under the empirically plausible assumption that in the initial equilibrium the L-industry has a higher output elasticity to low-skill labor. In the model, this elasticity has two components: the direct production share of low-skill labor in output, which is higher in the L-industry by definition, and an endogenous bottleneck effect that captures how additional low-skill workers dilute managerial bandwidth and reduce the probability of successful idea implementation.[9]

Proposition 2. *(Two-sector comparative statics with respect to $\bar{L}$). Consider a manufacturing sector composed of two industries, H and L, that differ in innovation and skill intensity,*

[9]See Proposition A1 in Appendix B for the formal statement and conditions.

$\phi_L < \phi_H$ and $\theta_L > \theta_H$. Let hats denote elasticities with respect to the aggregate low-skill labor endowment, so that $\hat{x} \equiv d\ln x / d\ln \bar{L}$. An increase in the supply of low-skill labor $\bar{L}$ implies the following changes in equilibrium outcomes:

1. *The skill premium w_H/w_L increases.*
2. *$\hat{L}_L > \hat{L}_H$: the L-industry absorbs the low-skill inflow more than proportionally.*
3. *$\widehat{TFP}_H < 0$ and $\widehat{TFP}_L < 0$. TFP declines in both industries due to the two mechanisms discussed in the baseline model:*

   (a) implementation effect: the tightening of managerial bandwidth reduces the probability of successfully implementing ideas in each industry and is larger in industry L due to its larger absorption of low-skilled labor;

   (b) innovation effect: the reallocation of high-skill labor away from innovation, whose TFP impact is governed by the elasticity of ideas to scientists (ϕ) and is therefore larger in industry H.

   The cross-industry ordering of TFP declines reveals information about the relative strength of these two forces. In particular, $\widehat{TFP}_L < \widehat{TFP}_H$ if and only if the gap in the reduction in productivity between industries L and H due to the implementation effect (3.a) is larger than the gap due to the innovation effect (3.b). For a formal statement see condition A3 *in Appendix B).*

As in our baseline model, a low-skill labor inflow rises the skill premium economy-wide, tightening managerial bandwidth in all industries and decreasing the equilibrium rate of successful idea implementation in both industries. As in the one-sector model, the inflow of low-skill labor expands production capacity. In the two-industry economy, however, this expansion is uneven: employment grows disproportionately in the L-industry. The intuition mirrors the logic of the Rybczynski theorem: because the L-industry is less skill-intensive, its effective marginal product of low-skill labor remains relatively high even as managerial bandwidth tightens, making it the primary absorber of the inflow: $\hat{L}_L > \hat{L}_H$.

The tighter implementation constraint also affects the allocation of high-skill labor across occupations. In the one-sector economy, the rise in the skill premium reduces the demand for managers and scientists relative to unskilled workers. As firms operate with fewer managers per worker, the probability of successful implementation of ideas falls. As a consequence, the demand for managers increases and high-skill workers reallocate away from innovation into managerial tasks. These forces combine to generate a decline in TFP.

As in the one-sector model, TFP declines in each industry through two channels: the direct reduction in implementation success (the Implementation effect) and the endogenous

reallocation of high-skill labor away from innovation (the Innovation effect). The cross-industry ordering of TFP reductions reveals the relative strength of these two channels. The implementation effect generates a larger contraction of TFP in the L-industry, where managerial bandwidth is more diluted by the larger inflow of low-skill workers. In contrast, the innovation effect generates a larger TFP decline in industry H because it has a higher elasticity of ideas to scientists ($\phi_H > \phi_L$). As a result, the reallocation of scientists towards management generates a larger loss of new ideas in industry H.

Taking Stock The two-sector model generates two additional empirical predictions relative to the baseline framework. First, following an inflow of low-skill labor, employment reallocates towards industries that are less skill and R&D intensive. Second, the heterogeneous effect of low-skill labor inflows on sectoral productivity provides an empirical test to identify the relative strenght of the two mechanisms operating in the data. If TFP declines more in *less* R&D-intensive industries, the direct implementation effect is quantitatively more important in the data. By contrast, if TFP reductions are larger in *more* R&D-intensive industries, this indicates that the loss of scientists in innovation-intensive sectors is the prevailing mechanism (innovation effect). The subsequent sections take these predictions to the data.

3 Empirical strategy and data

Our empirical strategy aims at tracing the effects of a "push" factor of structural transformation — labor-saving technical change in agriculture — on industrial employment composition, innovation and productivity. For this, we exploit the legalization of genetically engineered (GE) soy in Brazil as a quasi-natural experiment, using the potential increase in soy yields from GE seeds in each region — which is a function of its soil and weather characteristics — as a plausibly exogenous measure of agricultural technical change.

In what follows, we provide background information on GE soy (Section 3.1), describe the identification strategy (Section 3.2), present the data sources (Section 3.3), and detail the construction of our measures of innovation workers and managers (Section 3.4).

3.1 Background Information on genetically engineered Soy

Previous research has documented that genetically engineered (GE) agricultural technologies to produce soy are labor-saving and that their legalization in Brazil induced a reallocation of labor from agriculture to manufacturing (Bustos et al. (2016)). In this section, we present additional evidence indicating that the adoption of GE technology is also skill-biased.

Consequently, it triggered an inflow of low-skill labor into the manufacturing sector, which corresponds to the main comparative statics exercise in the model.

GE soy seeds are engineered to resist a specific herbicide (glyphosate), allowing farmers to spray their fields without harming soy plants and reducing labor requirements for weed control.[10] Traditional soy planting requires tillage, the labor-intensive removal of weeds from the seedbed, whereas GE seeds require no tillage, as herbicide application selectively eliminates weeds. Because weed control is mostly performed by unskilled workers, the introduction of GE soy is likely to displace unskilled labor relatively more than skilled labor.

The first generation of GE soy seeds (Monsanto's Roundup Ready) was commercially released in the U.S. in 1996 and legalized in Brazil in 2003.[11] However, smuggling of GE seeds from Argentina was detected as early as 2001 and 2002 (USDA 2001), which explains why some of the empirical effects we document below begin before formal legalization. Once legalized, adoption was rapid: by 2006, 46.4% of Brazilian soy farmers were using GE seeds (IBGE 2006, p.144), and by 2011–2012, GE seeds covered 85% of soy area (USDA 2012).[12]

The adoption of GE soy had two effects on agricultural labor demand. First, a *within-soy* effect: as documented in Panel (a) of Figure 1, legalization was followed by a rapid expansion of soy area, from 11 to 19 million hectares between 2000 and 2010.[13]. In the same period, employment in the soy sector declined substantially (Panel (b)). This finding is consistent with the adoption of GE seeds reducing the number of agricultural workers per hectare required to cultivate soy. Bustos et al. (2016) document that labor intensity in soy production fell from 28.6 workers per 1000 hectares in 1996 to 17.1 workers per 1000 hectares in 2006.

Second, an *across-crop* effect: soy is less labor-intensive than all other major agricultural activities[14]. Thus, whenever soy displaced other agricultural activities, labor intensity in agriculture decreased.

Crucially, both the within-soy and across-crop effects are skill-biased. To show this, we divide agricultural workers by their number of years of education, where those who completed 8th grade are defined as skilled. Panel (c) of Figure 1 shows that the decline in soy

[10]Other advantages of GE soy seeds are that they require fewer herbicide applications (Duffy and Smith 2001; Fernandez-Cornejo, Klotz-Ingram, and Jans 2002), allow a higher density of the crop on the field (Huggins and Reganold 2008) and reduce the time between cultivation and harvest.

[11]See Law 10.688 of 2003 and Law 11.105 – the New Bio-Safety Law – of 2005 (art. 35).

[12]Although the initial patent was filed in the US by Monsanto, the product available in Brazil resulted from a collaboration with Embrapa—the Brazilian agricultural research institute—which crossed Monsanto's herbicide-tolerant variety with seeds adapted to Brazilian agro-ecological conditions.

[13]According to the two most recent agricultural censuses, the area planted with soy increased from 9.2 to 15.6 million hectares between 1996 and 2006 (IBGE 2006, p.144).

[14]Cereals require 94.9 workers per 1,000 hectares, other seasonal crops 129.8, and permanent crops 126.7. According to the 2006 Agricultural Census, even cattle ranching uses more workers per unit of land than soy production (30.6 per 1,000 hectares).

employment is driven by low-skilled workers, while skilled workers were retained. Moreover, soy is more skill-intensive than other crops: the share of skilled workers in soy exceeds 20 percent, compared to 5-15 percent in other agricultural activities (Panel (d)). Thus, as soy expanded, agricultural labor demand contracted and shifted toward skilled workers, pushing predominantly low-skill workers out of agriculture.

3.2 Identification strategy

Our identification strategy exploits the legalization of GE soy seeds in Brazil as a quasi-natural experiment. The key idea is that some regions benefited more from GE soy adoption due to exogenous differences in soil and weather characteristics. Thus, we exploit the year of legalization of GE soy seeds in Brazil as a source of variation across time and differences across regions in the potential increase in soy yields generated by adoption of the new technology as a source of cross-sectional variation.

We measure the potential gains from GE soy adoption using the FAO-GAEZ database, which reports maximum attainable crop yields under different input technologies. Following Bustos et al. (2016), we define technical change in soy production in microregion k as:

$$\Delta A_k^{soy} = A_k^{soy,High} - A_k^{soy,Low}$$

where $A_k^{soy,Low}$ and $A_k^{soy,High}$ are potential soy yields under traditional seeds and improved (GE) technology, respectively.[15] Figure D.1 in the Appendix shows substantial geographical variation in this measure across microregions.

We use two main empirical specifications. First, we estimate long-run effects using decennial Population Census data through a differences-in-differences specification. Second, we use annual data from RAIS and PIA, to estimate a dynamic difference-in-differences specification which permits to trace the timing of effects and test for pre-trends. In particular, when analyzing decennial Census data, we use the following first-difference specification:

$$\Delta Y_{k(r)} = \alpha + \beta \Delta A_{k(r)}^{soy} + \varphi X_{k(r)} + \delta_r + \varepsilon_k \tag{8}$$

where $\Delta Y_{k(r)}$ is the change in the outcome of interest in microregion k (located in macrore-

[15]The FAO-GAEZ database reports the maximum attainable yields of each crop under different technologies. Yields under the *low* technology are described as those obtained from planting traditional seeds, with no use of chemicals or mechanization. Yields under the *high* technology are obtained using improved high-yielding varieties, with optimum application of fertilizers, herbicides, and mechanization. In the case of soy, the improved seed varieties used are GE soy seeds according to our direct consultation with FAO-GAEZ researchers. The predictive power of the change in potential yields on GE soy seeds adoption documented in what follows supports this claim.

gion r) between 2000 and 2010 – the years of the last two Population Censuses –, and $X_{k(r)}$ is a vector of controls of microregion k. δ_r indicates macroregion fixed effects that account flexibly for trends across the five major geographical regions of the country: North, Northeast, South, Southeast and Central-West. Our identification strategy relies on the fact that the GE soybeans seeds were legalized in Brazil in 2003, and that this new technology disproportionately favored microregions with certain soil and weather characteristics (as captured by $\Delta A^{soy}_{k(r)}$), which was not anticipated as of 2000.

In all our specifications we include controls for the initial level of income per capita, the share of rural population, the alphabetization rate, and population density at the microregion level, all sourced from the 1991 Population Census, and a measure of technical change in maize. These controls are meant to absorb differential trends across microregions with different initial levels of development.

Second, we estimate the timing of effects through the following dynamic difference-in-differences specification using annual data from RAIS and PIA:

$$\ln y_{k(r),t} = \delta_k + \delta_{rt} + \sum_{\substack{j=1998 \\ j\neq 2001}}^{2009} \beta_j 1[j=t]\Delta A^{soy}_{k(r)} + \Gamma X_{k(r)} \times year_t + \varepsilon_{k(r),t} \tag{9}$$

where β_j captures the effect of soy technical change in each year relative to 2001, δ_k are microregion fixed effects, and δ_{rt} are macroregion-by-year fixed effects. We use 2001 as reference year to at least partially account for early adoption of GE soy via smuggling of seeds from other countries before their official legalization in Brazil, which occurred in 2003. $X_{k(r)}$ are baseline controls, which we interact with year fixed effects.

We also estimate a panel specification that summarizes the post-treatment effect in a single coefficient:

$$\ln y_{k(r),t} = \delta_k + \delta_{rt} + \beta A^{soy}_{k(r),t} + \Gamma X_{k(r)} \times year_t + \varepsilon_{k(r),t} \tag{10}$$

where $A^{soy}_{k(r),t}$ switches from potential yield under low inputs (1998–2002) to potential yield under high inputs (2003–2009), so that β is a continuous difference-in-differences estimate.[16]

Table E.1 in the Appendix validates this measure of agricultural technical change using data from the 1996 and 2006 Agricultural Censuses. First, in Panel A, we show that the measure of soy technical change strongly predicts variation in the actual adoption of GE seeds by Brazilian farmers across microregions (columns 1 and 2). Importantly, it does not predict the expansion of area farmed with traditional soy (columns 3 and 4). This indicates

[16] In these specifications, we use a balanced panel of microregions that includes all the microregions for which we have observations in each year of the decade.

that this measure of the effect of technical change on potential soy yields is a good proxy of the actual benefits of GE soy adoption given soil and weather characteristics of different areas. Second, in Panel B, we show that our measure of soy technical change predicts the expansion of agricultural area farmed with soy, but not the one farmed with maize, the other main temporary crop which experienced significant technological innovation in this period (columns 1 and 2).[17] If we build a measure of maize technical change using the same methodology, we find that such measure predicts the expansion in maize area between 1996 and 2006, but not the expansion of soy area (columns 3 and 4). This indicates that our measure of technical change is a good proxy of technological innovation at the crop level.

3.3 Data

Testing the model's predictions requires measuring several key variables in the data. The first step is to establish that technical change in soy generates an inflow of low-skilled workers into the manufacturing sector. We measure employment by skill level across sectors using both the Population Census (which covers formal and informal workers at decennial frequency) and social security records from RAIS (which provide annual frequency for formal workers). Similarly, we measure the skill premium using composition-adjusted wages from the Population Census and RAIS. Second, we test the predictions of the model regarding the effects of the inflow of low-skilled workers on the organization of production and innovation tasks within the manufacturing sector. In particular, we focus on the allocation of high-skill workers between scientists and managers using occupation-level data from RAIS. We construct novel measures of workers in innovation-intensive occupations and in managerial occupations using task descriptions from RAIS. Third, we test the predictions of the model regarding productivity using TFP estimates constructed from the Annual Industrial Survey (PIA). Finally, the cross-industry predictions of the two-sector extension are tested by exploiting variation in industry-level R&D intensity from the Innovation Survey (PINTEC).

The geographical units of observation in the empirical analysis are Brazilian microregions, which attempt to approximate local labor markets or the small open economies described in the model. The Brazilian Institute of Geography and Statistics (IBGE) defines these microregions by combining economically integrated municipalities with similar production and geographic characteristics.

Population Census We use the Population Censuses of 2000 and 2010 to measure employment and wages across sectors. The Census covers both formal and informal workers, which

[17] See Bustos et al. (2016) for a detailed discussion of second-season maize.

is essential for studying movements out of agriculture—where most employment is informal. We focus on individuals aged 25–55 working more than 35 hours per week. Table 1 reports summary statistics at micro-region level of employment shares in agriculture, manufacturing, and services from the Census data.[18]

RAIS The Annual Social Information System (RAIS) is an employer-employee dataset covering the universe of formal workers in Brazil.[19] We use RAIS from 1998 to 2009 to study within-manufacturing reallocation at yearly frequency. As in the Population Census, we focus on individuals aged between 25 and 55 that work more than 35 hours a week.[20] RAIS contains detailed occupation information, which we use to identify workers in innovation-intensive occupations (our empirical counterpart of scientists S) and managerial occupations (our empirical counterpart of managers M), as described in Section 3.4.

We also use RAIS to construct composition-adjusted wages and skill premia. We regress log hourly wages of formal, male, full-time workers employed as of December 31st on microregion fixed effects, age-group fixed effects, sector fixed effects, and age-group–sector interactions. The estimated microregion fixed effects capture local wage levels net of observable worker characteristics. We classify workers as skilled if they have completed high school or more, and as unskilled otherwise, and estimate separate regressions for each group. The skill premium in each microregion is defined as the difference between the skilled and unskilled microregion wage premia.[21]

PIA We use the Annual Industrial Survey (PIA) to construct measures of manufacturing productivity. PIA covers all firms with 30 or more employees, ensuring representativeness at the microregion-industry level. We measure employment as end-of-year workers and value added as value of production minus intermediate inputs. We construct capital stock using

[18] We define agriculture, manufacturing and services by following the classification of the CNAE Domiciliar of the 2000 Census. Agriculture includes Sections A and B (agriculture, cattle, forestry, and fishing). Manufacturing includes Section D, which corresponds to the transformation industries. Services include: construction, commerce, lodging and restaurants, transportation, finance, housing services, domestic workers, and other personal services. We exclude the following sectors because they are mostly under government control: public administration, education, health, international organizations, extraction, and public utilities.

[19] Employers are required by law to report worker information to the Ministry of Labor (Decree n. 76.900, December 23^{rd} 1975).

[20] Following Helpman, Itskhoki, Muendler, and Redding (2017), our data cleaning procedure includes: (i) restricting to workers employed as of December 31st in each year; (ii) restricting to the highest-paying job for each worker that appears more than once in the data during one year (randomly dropping ties).

[21] For the wage analysis we focus on the period starting in 1999, as the level of detail in the reporting of wages in RAIS changed in that year. Before 1999, wages in RAIS were reported only as multiples of the national minimum wage.

the perpetual inventory method.[22] We focus on manufacturing industries (CNAE 1.0 codes 15–37) for 2000–2009.

PINTEC We use data from the Survey of Innovation (PINTEC) to classify manufacturing industries by R&D intensity, which maps directly onto the two-sector extension of the model (Section 2.3). We compute R&D intensity as total R&D expenditure over sales in 2000, and define high-R&D industries (H) as those above the median and low-R&D industries (L) as those below.[23]

Table E.2 provides summary statistics for the main variables at the microregion level. Brazil has 557 microregions, with an average population of around 300,000 inhabitants. We use microregions as approximations of local labor markets—small open economies where goods are traded but labor is relatively immobile.

3.4 Measuring Innovative Workers and Managers

In the model introduced in Section 2, high-skill workers can be employed either as scientists (S), who generate new ideas, or as managers (M), who supervise low-skill workers (L) and implement innovations. The model predicts that an inflow of low-skill workers into manufacturing creates management bottlenecks, inducing firms to reallocate high-skill workers away from innovation activities and toward managerial tasks. Testing these predictions requires measures of workers employed in innovation-intensive occupations and in managerial occupations that are representative at fine levels of geographical disaggregation. We construct both measures using occupation data from RAIS, as described below.

Measuring Innovation Workers In the model, scientists (S) are high-skill workers who generate new ideas. To construct an empirical counterpart, we develop a measure of workers employed in innovation-intensive occupations using textual analysis of the task descriptions of more than 2,500 occupations in the Brazilian Classification of Occupations (CBO 2002). Specifically, we define a set of 39 keywords capturing tasks related to innovative activities—such as developing new products and processes, creating prototypes, or optimizing production methods—and classify an occupation as innovation-intensive if its description contains at least one of these keywords.

The list of keywords used is reported in Appendix Table E.4. As shown, most entries are a combination of a verb and a noun describing a task associated with innovation. These com-

[22] We use data on investment, depreciation, and book value of assets in the first year a firm appears in the sample.

[23] Table E.3 in the Appendix reports the full list of industries by R&D intensity.

binations can be grouped in those capturing innovation of products (e.g. "develop/improve product/s"), innovation of processes (e.g. "develop/improve/test process/es"), innovation of machinery and equipment (e.g. "develop device/s", "develop equipment"). We also include single nouns, combinations of nouns, or combinations of nouns and adjectives that are often found in the description of innovation intensive tasks (e.g. "innovation", "prototypes", "research and development", "new technologies"). This methodology identifies 251 innovation-intensive occupations.[24]

We perform a set of consistency tests on our measure of innovative labor. Figure D.2 shows the correlation between employment share in innovation-intensive occupations and other measures of innovation that are available at the industry level from the PINTEC survey. We include measures that capture the amount of inputs devoted to the innovation process – such as R&D expenditure per worker – as well as measures capturing the output of the innovation process – such as the share of firms in a given sector that have filed patents and the share of firms that have introduced new processes or products. As shown, the share of innovation-intensive workers is highly correlated with all these alternative measures, with the additional advantage of being available not only at the sector level but also at fine levels of geographical aggregation. Table E.5 reports the magnitude of the correlations between the share of innovation-intensive workers in each industry and the alternative measures reported in Figure D.2.

An important advantage of this measure is that it is representative at fine levels of geographical disaggregation, which is essential for our empirical analysis at the microregion level. Standard innovation surveys, such as PINTEC, are based on samples that are not representative at low levels of geographic aggregation. Similarly, patent data—while capturing innovation output—miss the large share of innovations that firms do not patent. Data from PINTEC show that only 7 percent of firms introducing new products or processes file patent applications.[25]

Finally, it may also be useful to clarify the nature of innovative activities performed by workers in our sample. Aggregate data from the PINTEC suvey indicate that most innovations introduced by Brazilian firms happen through adaptation of technologies that are new to the firm but already in use elsewhere. In particular, about 84% of new products and 94% of new processes introduced are an innovation for the firm but already exist in some form either in Brazil or in the rest of the World according to the PINTEC firm survey. In this sense, we think of innovation workers in our sample as mostly involved in adapting technologies developed elsewhere to the Brazilian context.

[24]See Lagaras (2017) for an application of this methodology at the firm level.

[25]Based on Table 6497 of the PINTEC surveys covering 1997–2008. Cohen, Nelson, and Walsh (2000) document similar patterns in US data.

Measuring Managers In the model, managers (M) are high-skill workers who supervise low-skill workers and teach them how to implement innovations. To construct an empirical counterpart, we identify managerial occupations using keyword matching on occupation names in the Brazilian Classification of Occupations (CBO 2002). Unlike innovation-intensive occupations, which require analyzing detailed task descriptions, managerial roles can be more readily identified from job titles, as these typically contain explicit references to supervisory or administrative responsibilities.

Specifically, we classify an occupation as managerial if its name contains terms such as *gerente* (manager), *supervisor* (supervisor), *coordenador* (coordinator), *chefe* (chief), or *diretor* (director). The complete list of 14 keywords is reported in Appendix Table E.6. For occupations containing *diretor*, we exclude artistic and creative directors by filtering out those whose names also contain terms related to arts and entertainment.[26] In some cases, there is overlap between innovation-intensive and managerial occupations, as roles such as R&D managers involve both supervisory responsibilities and innovation-related tasks. These occupations are included in both categories.

4 Industrialization, the organization of production, and productivity

According to the model introduced in section 2, the first step in the analysis of the effect of structural transformation on economic development is to study the nature of the shock that leads to the relocation of workers toward manufacturing. The model suggests that an (exogenous) inflow of workers into manufacturing may have very different consequences depending on whether the labor inflow is mainly low- or high-skilled. The second step involves studying how the labor inflow affects the organization of production. Finally, the model has implications for the evolution of productivity and for the heterogeneous effects of unskilled labor inflows on the productivity of different industries within manufacturing.

In particular, the model delivers three empirical predictions: (i) skill-biased labor inflows raise the skill premium; (ii) firms reallocate high-skill labor from innovation to management; and (iii) manufacturing productivity growth declines despite industrial expansion. We test empirically each of these points in what follows. The two-sector extension generates additional cross-industry predictions that help identify the dominant mechanism, which we also test in Section 4.4.

[26] Specifically, we exclude occupations containing terms such as *cinema*, *teatro*, *televisão*, or *artístico*.

4.1 Reallocation of low-skilled labor from agriculture to manufacturing

GE soy technical change led to a reallocation of labor from agriculture to manufacturing, as documented using Population Census data by Bustos et al. 2016. We replicate this finding in Table E.7 in the Appendix, where we additionally show that labor reallocation was predominantly driven by unskilled workers. Estimated coefficients from equation (8)indicate that microregions with a standard deviation higher increase in potential soy yields experienced a 2.5 percentage points larger decrease in the share of low-skilled workers employed in agriculture, and a corresponding 2 percentage points larger increase in the share of low-skilled workers employed in manufacturing, while the effects on high-skilled workers were smaller in magnitude.

Our empirical analysis primarily relies on annual social security data from RAIS, which offers several advantages for testing the model's predictions relative to Population Census data: RAIS allows us to observe manufacturing employment at yearly frequency, and thus to study the dynamics and timing of the effects; informality is less of a concern in the manufacturing sector than in agriculture; and it contains detailed occupation information, which is essential for identifying workers in innovation-intensive and managerial occupations and tracking them over time.

Figure 2 plots estimated coefficients of the effect of technical change in soy on formal manufacturing employment by skill group from the event-study specification described by equation (9). Low-skilled workers started to reallocate towards manufacturing in microregions more exposed to soy technical change around 2002, while there is no systematic difference in the trends in previous years. We find no differential increase in skilled labor moving towards manufacturing, as shown in Figure 2 (b). The timing of the effect suggests that changes were permanent. Reallocation of unskilled labor towards manufacturing started around 2002, one year after the first reported smuggling of the GE soy seeds in Brazil and the year when the area planted with soy started expanding at a fast rate (Figure 1). The reallocation then accentuated around 2004, one year after the formal legalization of GE soy in Brazil, and stabilized during the second half of the decade.

We summarize the results shown in Figure 2 by estimating the panel specification that averages the post-treatment effect in a single coefficient described in equation (10). Estimates reported in Table 2 imply that that microregions with a standard deviation higher increase in potential soy yields experienced an increase in formal manufacturing employment of about 9 percent. This inflow of workers is explained by an increase in low-skilled labor, which

increased by about 10.4 percent for one standard deviation. Coefficient estimates are stable when we control for initial micro-region characteristics interacted with year fixed effects.

Taken together, the estimates presented in Table E.7, Table 2, and Figure 2 show that low-skill labor reallocated from agriculture to manufacturing in regions more exposed to GE soy adoption. These findings indicate that labor-saving technical change in agriculture driven by the adoption of GE soy was skill-biased and led low-skill workers to reallocate towards manufacturing.

4.2 Industrialization and the organization of production

From the point of view of the manufacturing sector, the reallocation of low skilled workers from agriculture documented above amounts to an increase in the supply of unskilled labor. In the model, an inflow of low-skilled workers shapes the organization of production in the industrial sector: the skill premium rises, making both managers and scientists more costly relative to production workers. Firms respond by reducing the demand for managers and scientists relative to unskilled workers. Consequently, each firm employs fewer managers per worker, which lowers the likelihood that new ideas are implemented successfully. To counteract this decline in implementation success, firms reassign high-skill workers from innovation activities toward managerial roles.

Consistent with prediction (1) of Proposition 1, we find that microregions with faster agricultural technical change experienced a rise in the skill premium. To measure skill premia we use composition-adjusted wages—that is, wages net of observable worker characteristics—obtained from Mincerian regressions as described in Section 3.3. Using annual RAIS data on formal workers and the panel specification in equation (10), Table 3 shows that a one standard deviation increase in soy technical change raises composition-adjusted wages of skilled workers by 2.4 percent (column 2), while the effect on unskilled wages is smaller and not statistically significant (column 4). As a result, the skill premium increases: microregions with a one standard deviation larger increase in soy technical change experienced a 1.6 percent higher increase in the skill premium (column 6). We confirm these findings using Population Census data where we find a larger skill premium increase of 3.6 percentage points for formal workers.[27] Finally, we analyze the timing of these changes in Figure 3 , which reports yearly

[27]We use the Population Census to examine the effect of agricultural technical change on the skill premium for both formal and informal workers, as the latter are not covered by RAIS. Estimates are reported in Appendix Table E.9. Among formal workers, composition-adjusted wages of skilled workers rise by 5.6 percent while those of unskilled workers rise by 2.0 percent, yielding a significant skill premium increase of 3.6 percentage points. Among informal workers, however, composition-adjusted wages for both skilled and unskilled labor increase by comparable magnitudes (2.5 and 2.2 percent, respectively), and the implied change in the skill premium is not statistically distinguishable from zero.

estimates of the effect of soy technical change on the skill premium around the legalization of GE soy in Brazil using the event-study specification in equation (9). The pattern is consistent with the model's prediction that a low-skill labor inflow raises the relative price of high-skill workers, who become scarcer relative to unskilled workers.

Next, we investigate the predictions of the model regarding the reorganization of production and innovation in response to the increase in the skill premium (Proposition 1, parts 2 and 3). Table 4 reports estimates of the effects of agricultural technical change on the relative wage bill shares of managers and scientists using the panel specification in equation (10). Estimates show that the wage bill devoted to high-skill management and high-skill innovation workers compared to low-skill workers decreased. This is consistent with the prediction that firms respond to the increase in the skill premium by decreasing the ratio of managers (and innovators) per unskilled worker. Perhaps more interestingly, columns (5) and (6) show that there is a reallocation of skilled labor use towards high-skilled managers and away from innovation. This is one of the mechanisms highlighted in the model. As high-skilled labor is relatively more scarce and expensive, managerial bandwidth tightens and there is larger demand for supervision of low-skilled labor. As a result, scientific or innovation work declines relative to managerial work. Estimates imply that regions with a standard deviation higher increase in soy technical change experienced fairly large changes in the organization of production: a 7% decline in the wage bill of managers relative to low-skill workers, a 20% decline in the wage bill of scientists relative to low-skill labor and a 14% decline in the wage bill of scientists relative to managers.

Finally, we investigate the timing of these organizational changes in Figure 4. Panels (a) and (b) show that the wage bill devoted to high-skill innovation and high-skill management workers compared to low-skill workers decreased sharply around 2002 and especially in 2003, coinciding with the adoption of GE soy and the subsequent inflow of low-skilled labor into manufacturing. In turn, panel (c) shows a similar timing for the change in factor use towards high-skilled managers relative to innovation work. Note that while some of the coefficients are marginally significant in this event-study specification, when pulling all post and all pre-shock observations together in the panel-specification reported just above, we obtain statistically significant results at conventional levels (Table 4).

4.3 Industrialization without productivity growth

In this section, we investigate the predictions of the model regarding productivity growth. Note that the model predicts that a larger manufacturing sector accelerates innovation and

increases aggregate productivity, as in canonical growth models, only when the increase in labor supply is proportional for skilled and unskilled labor. However, we just documented that the inflow of agricultural workers into manufacturing was driven by low-skilled labor. In this case, the model predicts that the lower number of managers per worker reduces the probability of successful implementation of new ideas. In turn, the reallocation of scientists towards management reduces innovation, and thus the number of new varieties. Both effects reduce total factor productivity.

To test this prediction, we construct three measures of productivity using data from the manufacturing survey (PIA): value added per worker, valued added over wage bill, and total factor productivity.[28] Table 5 shows estimates of the effects of technical change in soy on productivity using the panel specification in equation (10). Estimates imply that Microregions with a standard deviation higher increase in soy technical change experience substantial declines in manufacturing productivity of about 7.5% across all three measures of productivity. Estimates are stable to the inclusion of various sets of controls.

To analyze the timing of these productivity changes, Figure 5 shows the corresponding yearly coefficient estimates using the event-study specification in equation (9). We see that microregions more exposed to soy technical change were, if anything, experiencing manufacturing productivity improvements prior to 2001. These trends stop and even reverse coinciding with the introduction of GE soy beans and the subsequent flow of low skilled workers towards manufacturing. The effects become larger later in the decade. We observe the same patterns with the three measures of manufacturing productivity.

Taken together, the results presented in Figure 5 and Table 5 indicate that, despite the fact that soy technical change reallocated labor from agriculture towards manufacturing, increasing the scale of production, productivity growth in manufacturing slowed down. The model and empirical evidence presented so far suggests that this is because skill-biased technical change in soy released low-skill workers, which led to higher scarcity of skilled labor in manufacturing. As a result, firms operated with less managers per worker, reducing their effectiveness in the implementation of ideas, and reallocated innovators towards management, reducing the rate at which they introduced new varieties.

Let us discuss a plausible alternative interpretation of our findings, namely that the de-

[28] We compute total factor productivity as the Solow residual of a Cobb-Douglas production function that combines skilled labor, unskilled labor, and capital in constant returns to scale fashion. We compute the factor shares for skilled and unskilled labor by combining the share of the aggregate wage bill that corresponds to each type of labor in each type of industry with measures of the labor share at the industry-level for the US retrieved from Becker, Gray, and Marvakov (2021). The capital share is calibrated by leveraging the constant returns to scale assumption. These assumptions imply that for a given industry the production technology is the same across micro-regions and periods, and thus, changes in the TFP are dictated by changes in the allocation of production factors.

cline in manufacturing productivity is driven by a change in the composition of workers within manufacturing. Our measures of productivity take educational differences across workers into account as we measure changes in value added per wage bill, and account for the two skill types separately in our measure of TFP. However, it is possible that workers with the same measured level of education have different levels of human capital. In particular, former agricultural workers may lack the skills needed to thrive in the manufacturing sector, and thus be less productive than workers already employed in manufacturing, at least temporarily. This, in turn, could lower the productivity of the sectors where they enter. We explore this channel using individual-level data from RAIS which permit to follow workers as they switch sectors. RAIS data shows that formal agricultural workers entering L manufacturing industries earn on average around 27% lower wages than workers already employed in the L industries, and 12% lower wages when compared with workers with similar individual characteristics such as age and education. Assuming wage differences across sectors are a good measure of differences in worker productivity, and given that manufacturing productivity decreased by around 7.5% for a one standard deviation increase in exposure to soy technical change, and that employment of low-skilled labor increased by around 10.4%, the worker composition effect can explain at most 37% (10.4% $\times$ 27% / 7.5%) of the decline in productivity.

4.4 Heterogeneous Effects by Industry and Discussion of Mechanism

The two-sector extension of the model generates two additional testable predictions described in Proposition 2. First, the inflow of low-skill labor should be absorbed disproportionately by less R&D-intensive industries (L-industries). Second, productivity should decline in both industries, and the relative size of the declines across industries can be used to identify the main mechanism at work in the data. If productivity declines more in L-industries, the implementation effect dominates; if instead productivity declines more in H-industries, the innovation effect dominates. We test these predictions using the classification of manufacturing industries into H and L based on R&D intensity from PINTEC, as described in Section 3.3.

Table 6, Panel A, reports the effect of soy technical change on employment in L and H industries using annual RAIS data and the panel specification in equation (10). Consistent with Proposition 2, labor flows into L-industries: microregions with a one standard deviation higher increase in soy technical change experience a 10.3 percent increase in employment in L-industries, with no significant effect on H-industries (columns 1 and 2). Columns 3–6 decompose employment by skill. The expansion of the L-industry is entirely driven by low-

skill workers (column 3), while high-skill employment does not change significantly in either industry. This pattern mirrors the Rybczynski logic embedded in the model: because L-industries use low-skill labor more intensively, their effective marginal product of low-skill labor remains relatively high even as the number of managers per worker falls, making them the primary absorber of the inflow.

Panel B of Table 6 reports the effect of soy technical change on productivity separately for L and H industries. For this analysis, we rely on the same three measures of productivity described above. We find a negative effect of GE soy adoption on TFP in both industries, but the decline is larger and statistically significant only in L-industries. In terms of magnitude, a one standard deviation increase in soy technical change reduces TFP by 11.5 percent in L-industries (column 1) compared to 6 percent in H-industries (column 2), where the effect is not statistically significant. Similar patterns hold for value added over the wage bill and value added per worker, with negative and significant effects for L-industries and smaller and non-significant effects for H-industries.

The fact that productivity declines *within* both industries has an important implication: the aggregate manufacturing productivity decline documented in Section 4.3 is not simply a composition effect driven by the reallocation of employment toward less productive L-industries. If the decline were purely a between-industry phenomenon, we would observe no within-industry productivity changes. Instead, the within-industry declines indicate that the mechanisms emphasized by the model operate within industries and are the primary drivers of the aggregate productivity slowdown.

In addition, as discussed in Section 2.3, the heterogeneous TFP decline across industries can be used as a test the relative strenght of the model's mechanisms. The fact that TFP declines are concentrated in L-industries indicates that the implementation effect—whereby managerial bandwidth is more severely diluted in industries that absorb more low-skill workers—is the dominant force.

Taken together, the results in this section provide a coherent picture that maps closely into the model's predictions. The low-skill labor inflow expands L-industries. The within-industry nature of the productivity declines rules out a pure composition explanation and points instead to the organizational mechanisms at the core of the model. Finally, the fact that TFP declines are larger in L-industries reveals that the direct managerial implementation effect dominates the innovation channel, providing direct evidence for the centrality of the implementation mechanism in this setting.

4.5 Reallocation of Skilled Workers across Occupations and Sectors

In this section, we investigate the predictions of the two-sector extension of the model presented in Section 2.3 on the reallocation of skilled workers across occupations. Because low-skill agricultural workers enter primarily into L-industries, the rise in the demand for managerial supervision is concentrated there. Thus, the model predicts that high-skilled workers that initially perform innovative tasks will have an incentive to transition into management in L-industries, while no comparable reallocation should arise towards management in H-industries.

The longitudinal structure of RAIS data allows us to zoom into this prediction of the two sector model by tracking the movement of innovation workers across occupations and sectors over time. We start by identifying as innovation workers those who hold an innovation occupation in each industry in the baseline years 2000, 2001, or 2002. We then follow innovators' occupational trajectories over time. In particular, we are interested in studying the impact of soy technical change on four potential trajectories for such workers: remain innovators in their original industry, become innovators in other sectors, transition to managerial occupations in their original industry, or transition to managerial occupations in other sectors.

For every microregion-year, we compute the share of the wage bill paid to workers in each trajectory.[29] Table 7 quantifies the effect of soy technical change on the share of the wage bill paid to innovators who initially worked in the L industry and followed each of the four trajectories described above. Column (3) provides the central piece of evidence. The model predicts that innovation workers should reallocate towards managerial occupations in the L-industry in regions with high exposure to the soy shock. Indeed, estimates in column (3) show that in regions with a one-standard-deviation higher increase in soy technical change, the wage bill share of workers who transition from innovation to managerial occupations within the L-industry is 2.6 percentage points higher. In turn, Column (2) shows that the same shock reduces by 1.9 percentage points the wage bill share of innovators who leave the L-industry while remaining in innovation. This indicates that innovators transition to management within the industries that absorb the low-skill inflow rather than reallocating to other sectors.

Next, we study the timing of these transitions across occupations. Figure 6 reports event-study estimates for each of these four trajectories. Panel (a) of Figure 6 shows that the share of innovation workers who stay in innovation in the L-industry does not vary differentially

[29] We use the wage bill rather than worker counts to account for productivity differences across workers.

with GE soy exposure. Panel (b) shows that innovation workers in the L-industry are less likely to move to become innovators in other industries in regions more exposed to GE soy. This might be dute to their transition into management within the L industry: panel (c) documents that the share of innovators who remain in the L industry but transition into managerial occupations rises disproportionately with exposure to the soy shock. For completeness, Panel (d) of Figure 6 shows no differential movement towards managerial occupations for innovation workers who moved away from the L industry.

We perform an analogous analysis for innovators initially employed in the H-industry, as reported in Table E.10 and Figure D.3 in the Appendix. Note that in this case, the model does not predict a reallocation of innovation workers towards management because H-industries did not receive a low-skill labor inflow and thus no managerial bottleneck arises. Indeed, we do not find any effect of soy technical change on reallocation towards management within these industries (Column 3). Consistent with the model, we do find some reallocation of innovation workers initially in the H industry towards management in other industries but due to large standard errors these differences are not statistically significant (Column 4). Note that worker reallocation from innovation in the H-industry to management in the L-industry is possibly weaker than the within L-industry reallocation documented above due to frictions in labor reallocation across industries present in the data but abstracted from in the model.

Overall, these mobility patterns are strongly in line with the main mechanism in our model. As the supply of low-skill labor expands, managerial bottlenecks tighten in the industries that absorb the inflow, and skilled workers in those industries shift out of innovation and into management.

5 Conclusions

This paper shows that the skill composition of the labor inflow into manufacturing is a key determinant of whether structural transformation leads to productivity growth. We develop a model in which an inflow of low-skill workers tightens managerial bandwidth — reducing the rate at which new ideas are implemented — and raises the skill premium, inducing firms to reallocate high-skill workers from innovation to management. Both forces reduce manufacturing productivity even as the sector expands. Exploiting the adoption of GE soy in Brazil as a natural experiment, we find empirical support for each of the model's predictions: the soy shock generated a predominantly low-skill labor inflow into manufacturing, raised the skill premium, shifted the organization of production away from innovation and toward management, and reduced manufacturing TFP. The fact that the negative effects are stronger

in less R&D-intensive industries suggest that the implementation bottleneck as the dominant mechanism.

Our findings suggest that labor-saving agricultural technologies — common when developing countries import technologies from more advanced ones — can generate static gains in agriculture alongside dynamic losses in manufacturing. These losses arise not from industrialization per se, but from the skill composition of the labor released from agriculture and its effects on the internal organization of firms.

References

Aghion, P. and P. Howitt (1992). "A Model of Growth through Creative Destruction". *Econometrica 60*(2), 323–351.

Aghion, P. and P. Howitt (2008). *The Economics of Growth*. The MIT Press.

Atkin, D., A. Costinot, and M. Fukui (2021). Globalization and the ladder of development: Pushed to the top or held at the bottom? Technical report, National Bureau of Economic Research.

Becker, R., W. Gray, and J. Marvakov (2021). Nber-ces manufacturing industry database: Technical notes. Technical report, National Bureau of Economic Research.

Bragança, A. (2014). *Three Essays on Rural Development in Brazil*. Ph. D. thesis, PUC-Rio.

Bustos, P., B. Caprettini, and J. Ponticelli (2016). Agricultural Productivity and Structural Transformation: Evidence from Brazil. *American Economic Review*.

Bustos, P., G. Garber, and J. Ponticelli (2020). Capital accumulation and structural transformation. *The Quarterly Journal of Economics 135*(2), 1037–1094.

Caselli, F. (2005). Chapter 9 Accounting for Cross-Country Income Differences. Volume 1, Part A of *Handbook of Economic Growth*, pp. 679 – 741. Elsevier.

Choi, J. and A. A. Levchenko (2021). The long-term effects of industrial policy. Technical report, National Bureau of Economic Research.

Clark, C. (1940). *The Conditions of Economic Progress*. London: Macmillan.

Cohen, W. M., R. R. Nelson, and J. P. Walsh (2000). Protecting their intellectual assets: Appropriability conditions and why us manufacturing firms patent (or not).

De Vries, G., M. Timmer, and K. De Vries (2015). Structural transformation in africa: Static gains, dynamic losses. *The Journal of Development Studies 51*(6), 674–688.

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

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

Foster, A. D. and M. R. Rosenzweig (1996). Technical Change and Human-Capital Returns and Investments: Evidence from the Green Revolution. *The American Economic Review 86*(4), 931–953.

Franck, R. and O. Galor (2019). Flowers of evil? industrialization and long run development. *Journal of Monetary Economics*.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy 108*(5), 874–904.

Garicano, L. and E. Rossi-Hansberg (2004). Inequality and the organization of knowledge. *American Economic Review 94*(2), 197–202.

Goldberg, P. K. and T. Reed (2020). Income distribution, international integration, and sustained poverty reduction. Technical report, National Bureau of Economic Research.

Gollin, D., D. Lagakos, and M. E. Waugh (2014). The Agricultural Productivity Gap. *Quarterly Journal of Economics*.

Grossman, G. and E. Helpman (1991). *Innovation and Growth in the Global Economy*. The MIT Press.

Helpman, E., O. Itskhoki, M.-A. Muendler, and S. J. Redding (2017). Trade and inequality: From theory to estimation. *The Review of Economic Studies 84*(1), 357–405.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2022). New evidence on sectoral labor productivity: Implications for industrialization and development. Technical report, National Bureau of Economic Research.

Herrendorf, B. and T. Schoellman (2018). Wages, Human Capital and Structural Transformation. *American Economic Journal: Macroeconomics*.

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

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

Imbert, C., M. Seror, Y. Zhang, and Y. Zylberberg (2020). Migrants and firms: Evidence from china.

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*(1-2), 41–55.

Kuznets, S. (1973). Modern Economic Growth: Findings and Reflections. *American Economic Review*.

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

Lagaras, S. (2017). Corporate takeovers and labor restructuring. *Available at SSRN 3222950*.

Lewis, A. W. (1954). Economic Development with Unlimited Supplies of Labor. *The Manchester School*.

Lucas, R. (1988). On the mechanics of economic development. *Journal of Monetary Economics*.

Matsuyama, K. (1992a). A Simple Model of Sectoral Adjustment. *The Review of Economic Studies*, 375–388.

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

McCaig, B. and N. Pavcnik (2013). Moving out of agriculture: structural change in vietnam. Technical report, National Bureau of Economic Research.

McMillan, M. and D. Rodrik (2011). Globalization, Structural Change and Productivity Growth. In *NBER Working Paper 17143*. ILO/WTO.

McMillan, M., D. Rodrik, and C. Sepulveda (2017). Structural Change, Fundamentals and Growth: A Framework and Case Studies. *National Bureau of Economic Research Working Paper Series No. 23378*.

Peters, M. (2022). Market size and spatial growth—evidence from germany's post-war population expulsions. *Econometrica 90*(5), 2357–2396.

Restuccia, D., D. T. Yang, and X. Zhu (2008). Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis. *Journal of Monetary Economics 55*(2), 234 – 250.

Romer, P. (1990). Endogenous Technological Change. *Journal of Political Economy 98*(5).

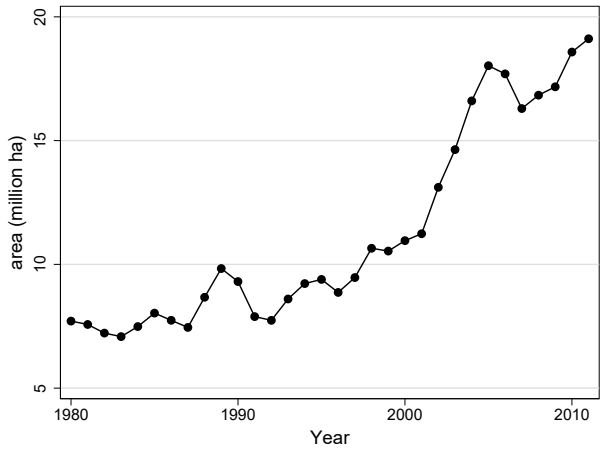
USDA (2001). "Agriculture in Brazil and Argentina: Developments and Prospects for Major Field Crops". United States Department of Agriculture, Economic Research Service.

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

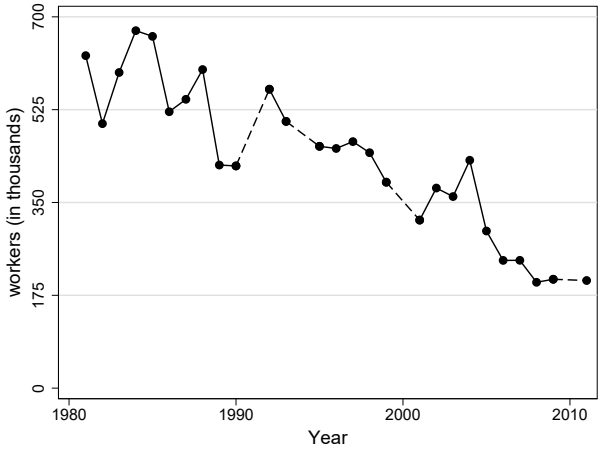
Figures and Tables

Figure 1: Soy Production and Employment

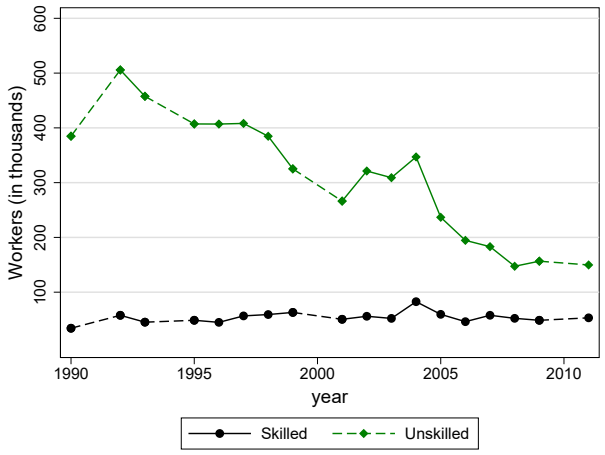
(a) Soy: Area planted



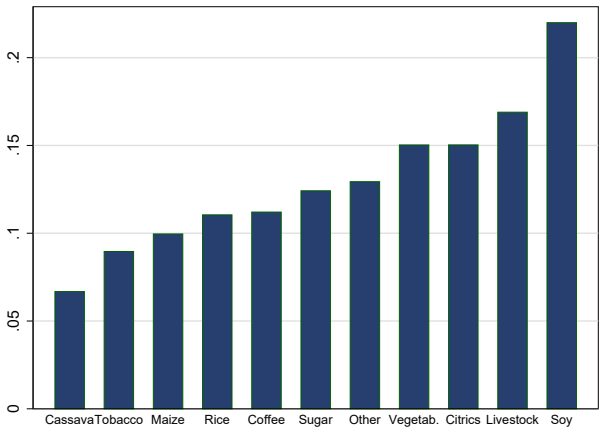
(b) Soy: Employment



(c) Soy: Employment by Skill Group

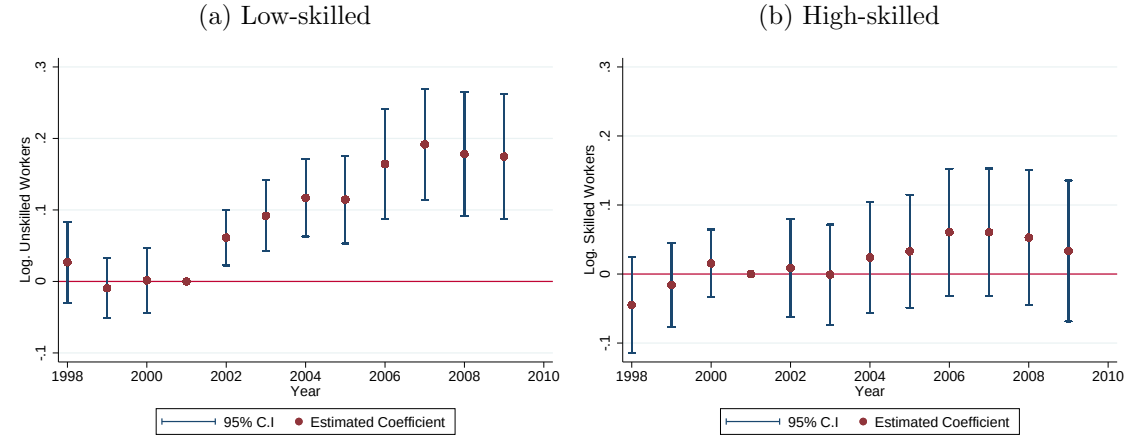


(d) Share of Skilled Workers by Agricultural Activity



Notes: Figures in Panels (a) and (b) are from Bustos et al. (2016). Data sources are CONAB (Panel A), PNAD (Panel B and C) and 2000 Population Census (Panel D). CONAB is the Companhia Nacional de Abastecimento, an agency within the Brazilian Ministry of Agriculture, which runs surveys of farmers and agronomists to monitor the annual harvests of major crops in Brazil. PNAD is the Brazilian National Household Sample Survey. The states of Rondonia, Acre, Amazonas, Roraima, Pará, Amapá, Tocantins, Mato Grosso do Sul, Goias, and Distrito Federal are excluded due to incomplete coverage by PNAD in the early years of the sample. In Panels C and D, an individual is classified as skilled if she has completed at least the 8^{th} grade.

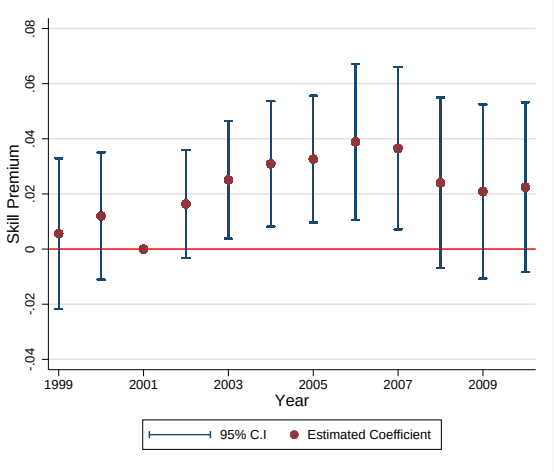
Figure 2: Effect of agricultural technical change on manufacturing labor by skill level



Notes: The figure shows the point estimates and the 95% confidence intervals for the estimates of the β_j of equation (9) where $\ln y_{k,r,t}$ corresponds to aggregate log employment of unskilled and skilled labor in microregion k located in region r at the end of year t in manufacturing. An individual is classified as high skilled if she has completed high school. (Source: RAIS). Standard errors are clustered at the microregion level.

Figure 3: Effect of agricultural technical change on skill premium
RAIS (1999-2010)

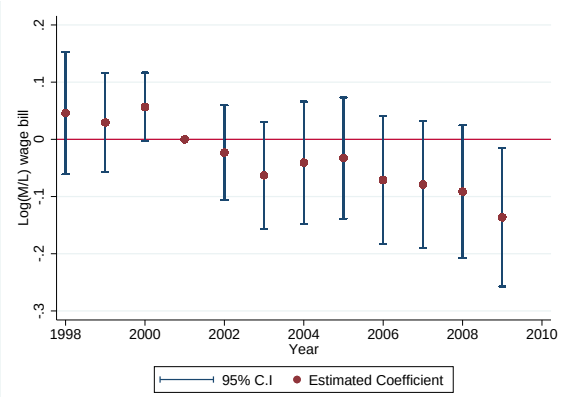
(a) Skill Premium



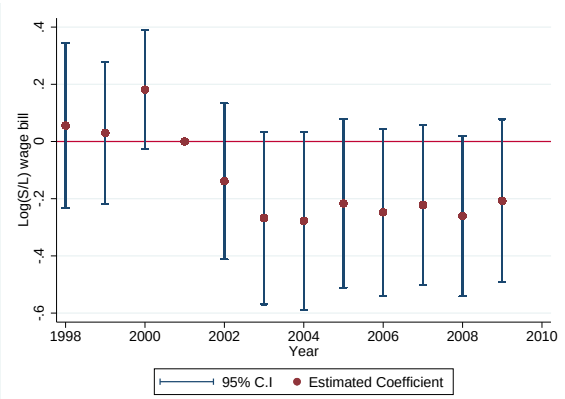
Notes: The figure shows the point estimates and the 95% confidence intervals for the estimates of the β_j from equation (9). The premia are constructed from yearly Mincer regressions of log hourly wages on microregion fixed effects, age-group fixed effects, sector fixed effects, and age-group–sector interactions. The skill premium is defined as the difference between the skilled and unskilled microregion wage premia. Skilled workers are defined as those with a high school degree or more. Standard errors are clustered at the microregion level.

Figure 4: Effect of agricultural technical change on labor cost ratios

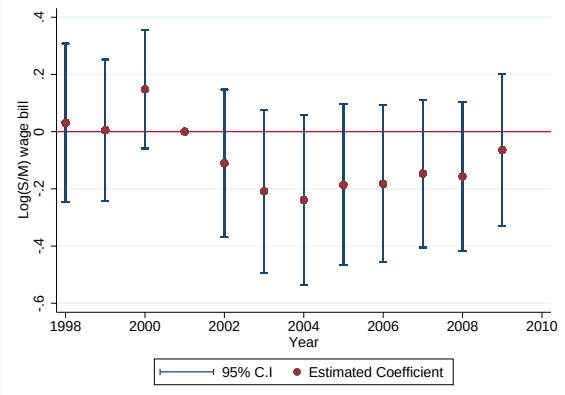
(a) High-Skill Managers / Low-Skill workers



(b) High-Skill Scientists/ Low-Skill workers



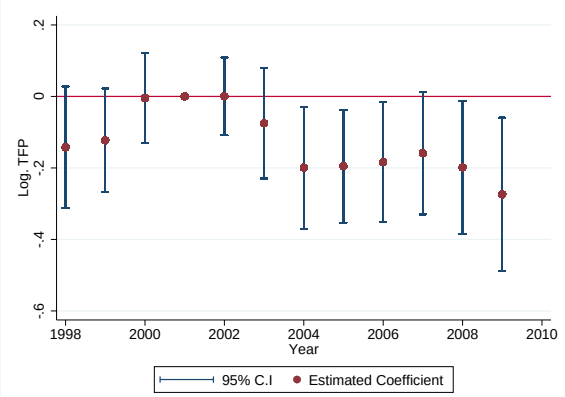
(c) High-Skill Scientists/High-Skill Managers



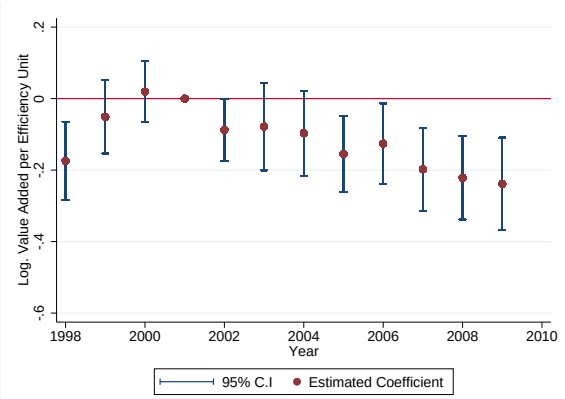
Notes: The figure shows the point estimates and the 95% confidence intervals for the e stimates of the β_j of equation (9) where the outcome variable is the log of the ratio of wage bill of the indicated categories of workers in microregion k located in region r at the end of year t in the manufacturing sector. Standard errors are clustered at the microregion level.

Figure 5: The effect of agricultural technical change on manufacturing productivity

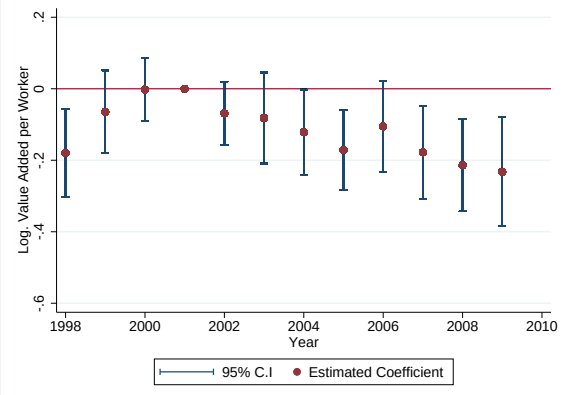
(a) **Total Factor Productivity**



(b) **Value Added per Efficiency Unit**



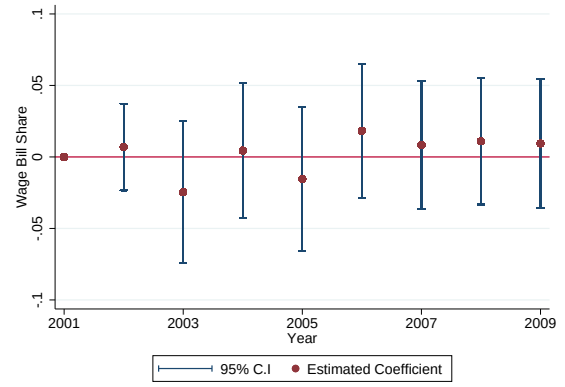
(c) **Value Added per Worker**



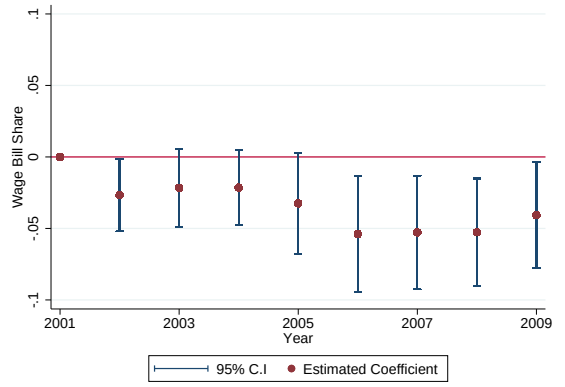
Notes: The figure shows the point estimates and the 95% confidence intervals for the estimates of the β_j of equation (9) where the outcome variables are different measures of productivity in the manufacturing sector. Standard errors are clustered at the microregion level

Figure 6: The effect of agricultural technical change on occupational transitions of innovation workers in L industries

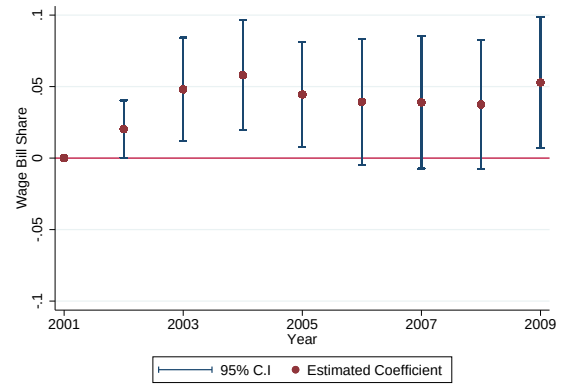
(a) **Stay in Innovation, Stay in L Industry**



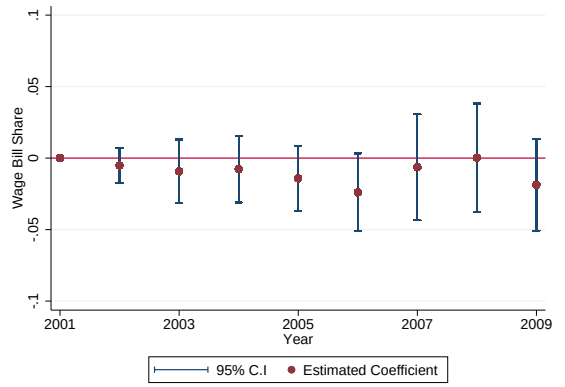
(b) **Stay in Innovation, Move to Other**



(c) **Move to Managerial, Stay in L Industry**



(d) **Move to Managerial, Move to Other**



Notes: The figure shows the point estimates and the 95% confidence intervals for the estimates of the β_j of equation (9) where the outcome variable is the wage bill share of the indicated category of skilled innovation workers initially employed in low-R&D (L) industries in microregion k located in region r at the end of year t in the manufacturing sector. Each share is computed as the wage bill of the indicated worker category divided by the total wage bill of all L-industry skilled innovation workers in that microregion and year. Workers are identified as innovators in 2000, 2001, or 2002. Standard errors are clustered at the microregion level.

Table 1: Summary Statistics: Micro-region Level

		Panel A: Decadal Variables				
		2000		Δ2000-2010		
	Source	Mean	SD	Mean	SD	Observations
Potential Yields	*FAO-GAEZ*					
Soy		0.287	0.136	1.787	0.740	557
Maize		1.847	0.998	3.082	1.639	557
Overall Employment Shares	*Population Census*					
Agriculture		0.279	0.140	−0.050	0.055	557
Manufacturing		0.147	0.077	0.006	0.042	557
Services		0.573	0.118	0.044	0.057	557
Unskilled Employment Shares	*Population Census*					
Agriculture		0.313	0.149	−0.031	0.057	557
Manufacturing		0.145	0.077	0.000	0.043	557
Services		0.542	0.126	0.030	0.054	557
Skilled Employment Shares	*Population Census*					
Agriculture		0.088	0.061	0.003	0.040	557
Manufacturing		0.159	0.085	0.011	0.055	557
Services		0.752	0.088	−0.014	0.061	557
Skill Premium & Wages	*Population Census*					
Skill Premium		0.602	0.154	−0.210	0.127	557
Skilled Wage		1.278	0.246	0.089	0.134	557
Unskilled Wage		0.676	0.246	0.299	0.095	557

	Panel B: Yearly Variables			
	Source	Mean	SD	Observations
Employment by Skill	*RAIS (1998-2009)*			
Log. Total Employment		7.205	2.002	6,496
Log. Unskilled Employment		6.891	1.952	6,496
Log. Skilled Employment		5.664	2.280	6,496
Skill Premium & Wages	*RAIS (1999-2009)*			
Skill Premium		0.472	0.194	6,124
Skilled Wage Premium		1.196	0.357	6,124
Unskilled Wage Premium		0.722	0.427	6,124
Innovation Wage Bills	*RAIS (1998-2009)*			
Log. M/L Wage Bill		−1.876	0.778	3,785
Log. S/L Wage Bill		−3.262	1.422	3,785
Log. S/M Wage Bill		−1.386	0.976	3,785
Manufacturing Productivity	*PIA (1998-2009)*			
Log. Value Added per Worker		10.721	0.807	3,575
Log. Value Added per Wage Bill		1.563	0.495	3,575
Log. Total Factor Productivity		6.416	0.627	3,575

Notes: The data sources are the Population Census (2000 and 2010), RAIS (1998-2009) and PIA (1998-2009). A worker is classified as skilled if he or she has completed high school.

Table 2: Effect of agricultural technical change on manufacturing employment, by skill level

	All		Low-skill		High-skill	
	(1)	(2)	(3)	(4)	(5)	(6)
A_{soy}	0.128***	0.121***	0.144***	0.140***	0.034	0.032
	[0.038]	[0.039]	[0.039]	[0.039]	[0.053]	[0.055]
Observations	6,496	6,496	6,496	6,496	6,496	6,496
R-squared	0.978	0.978	0.974	0.975	0.971	0.972
Microregion FE	Yes	Yes	Yes	Yes	Yes	Yes
Region x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Rural x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls x Year FE	No	Yes	No	Yes	No	Yes

Notes: The dependent variable is the log of manufacturing employment in the microregion, measured using the Yearly Social Security Data (RAIS), 1998–2009. Columns (1)–(2) use total manufacturing employment, columns (3)–(4) use unskilled manufacturing employment, and columns (5)–(6) use skilled manufacturing employment. A worker is classified as skilled if he or she has completed high school. All regressions include microregion fixed effects, region × year fixed effects, and the share of rural population interacted with year fixed effects. "Other Controls × Year FE" include income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census and interacted with year fixed effects. All regressions control for a measure of technical change in maize. Standard errors clustered at the microregion level are reported in brackets. Significance levels: $^{***}p < 0.01$,** $p < 0.05$,* $p < 0.1$.

Table 3: Effect of agricultural technical change on skill premium and wages of skilled vs unskilled labor

	Wage skilled		Wage unskilled		Skill premium	
	(1)	(2)	(3)	(4)	(5)	(6)
A_{soy}	0.020	0.033***	0.003	0.009	0.015	0.021**
	[0.013]	[0.012]	[0.008]	[0.007]	[0.011]	[0.010]
Observations	6,124	6,124	6,124	6,124	6,124	6,124
R-squared	0.928	0.932	0.978	0.979	0.806	0.810
Microregion FE	Yes	Yes	Yes	Yes	Yes	Yes
Region x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Rural x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls x Year FE	No	Yes	No	Yes	No	Yes

Notes: The dependent variables are the Mincerian wage premium for skilled workers (columns 1–2), the Mincerian wage premium for unskilled workers (columns 3–4), and the skill premium defined as the difference between the two (columns 5–6). Microregion-level wage premia are estimated from Mincerian regressions of log hourly wages on microregion fixed effects, controlling for skill, age, and sector interactions, using the Yearly Social Security Data (RAIS), 1999–2009. The sample is restricted to male workers. A worker is classified as skilled if he or she has completed high school. All regressions include microregion fixed effects, region × year fixed effects, and the share of rural population interacted with year fixed effects. "Other Controls × Year FE" include income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census and interacted with year fixed effects. All regressions control for a measure of technical change in maize. Standard errors clustered at the microregion level are reported in brackets. Standard errors clustered at the microregion level reported in parentheses. Significance levels: $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

Table 4: Effect of agricultural technical change on innovation and managerial occupation ratios in manufacturing

	$\log(\frac{\text{wage bill } M}{\text{wage bill } L})$		$\log(\frac{\text{wage bill } S}{\text{wage bill } L})$		$\log(\frac{\text{wage bill } S}{\text{wage bill } M})$	
	(1)	(2)	(3)	(4)	(5)	(6)
A_{soy}	-0.085*	-0.096**	-0.270***	-0.268***	-0.194**	-0.184**
	[0.043]	[0.044]	[0.091]	[0.091]	[0.077]	[0.078]
Observations	3,785	3,785	3,785	3,785	3,785	3,785
R-squared	0.875	0.878	0.832	0.836	0.672	0.677
Microregion FE	Yes	Yes	Yes	Yes	Yes	Yes
Region x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Rural x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls x Year FE	No	Yes	No	Yes	No	Yes

Notes: The dependent variables are the log ratio of the wage bill of skilled management workers to unskilled, non-innovative workers ($\log(M/L)$, columns 1–2), the log ratio of the wage bill of skilled innovation workers to unskilled, non-innovative workers ($\log(S/L)$, columns 3–4), and the log ratio of the wage bill of skilled innovation workers to skilled management workers ($\log(S/M)$, columns 5–6), measured using the Yearly Social Security Data (RAIS). All regressions include microregion fixed effects, region × year fixed effects, and the share of rural population interacted with year fixed effects. "Other Controls × Year FE" include income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census and interacted with year fixed effects. All regressions control for a measure of technical change in maize. Standard errors clustered at the microregion level are reported in brackets. Significance levels: $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

Table 5: Effect of agricultural technical change on manufacturing productivity

	TFP		VA/wage bill		VA/employment	
	(1)	(2)	(3)	(4)	(5)	(6)
A_{soy}	-0.115**	-0.126**	-0.093**	-0.099***	-0.089**	-0.093**
	[0.058]	[0.058]	[0.037]	[0.038]	[0.041]	[0.042]
Observations	3,575	3,575	3,575	3,575	3,575	3,575
R-squared	0.513	0.522	0.695	0.699	0.862	0.865
Microregion FE	Yes	Yes	Yes	Yes	Yes	Yes
Region x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Rural x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls x Year FE	No	Yes	No	Yes	No	Yes

Notes: The dependent variables are log total factor productivity (TFP, columns 1–2), log value added per wage bill (VA/WL, columns 3–4), and log value added per worker (VA/L, columns 5–6), measured using the Annual Industrial Survey (PIA). All regressions include microregion fixed effects, region × year fixed effects, and the share of rural population interacted with year fixed effects. "Other Controls × Year FE" include income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census and interacted with year fixed effects. All regressions control for a measure of technical change in maize. Standard errors clustered at the microregion level are reported in brackets. Significance levels: $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

Table 6: Effect of agricultural technical change on manufacturing employment and productivity: by industry

Panel A: Manufacturing employment by industry

	All		Low-skill		High-skill	
	L-industry	H-industry	L-industry	H-industry	L-industry	H-industry
	(1)	(2)	(3)	(4)	(5)	(6)
A_{soy}	0.140***	0.022	0.171***	0.002	0.015	0.020
	[0.052]	[0.038]	[0.053]	[0.039]	[0.064]	[0.051]
Observations	6,064	6,064	6,064	6,064	6,064	6,064
R-squared	0.965	0.975	0.960	0.969	0.958	0.967
Microregion FE	Yes	Yes	Yes	Yes	Yes	Yes
Region x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Rural x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls x Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Manufacturing productivity by industry

	TFP		VA/wage bill		VA/employment	
	L-industry	H-industry	L-industry	H-industry	L-industry	H-industry
	(1)	(2)	(3)	(4)	(5)	(6)
A_{soy}	-0.155**	-0.081	-0.081*	-0.037	-0.084*	-0.001
	[0.074]	[0.071]	[0.043]	[0.053]	[0.049]	[0.066]
Observations	3,360	3,360	3,360	3,360	3,360	3,360
R-squared	0.546	0.541	0.558	0.674	0.760	0.812
Microregion FE	Yes	Yes	Yes	Yes	Yes	Yes
Region x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Rural x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls x Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Panel A reports the effect of agricultural technical change on log manufacturing employment by industry, using the Yearly Social Security Data (RAIS). Panel B reports the effect on manufacturing productivity by industry, using the Annual Industrial Survey (PIA). L-industry refers to low-R&D intensive industries and H-industry refers to high-R&D intensive industries. In Panel A, columns (1)–(2) use total employment, columns (3)–(4) use unskilled employment, and columns (5)–(6) use skilled employment. A worker is classified as skilled if he or she has completed high school. In Panel B, columns (1)–(2) use log TFP, columns (3)–(4) use log value added per wage bill (VA/WL), and columns (5)–(6) use log value added per worker (VA/L). All regressions include microregion fixed effects, region × year fixed effects, and the share of rural population interacted with year fixed effects. "Other Controls × Year FE" include income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census and interacted with year fixed effects. All regressions control for a measure of technical change in maize. Standard errors clustered at the microregion level are reported in brackets. Significance levels: $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

Table 7: Effect of agricultural technical change on wage bill shares in innovation and managerial occupations: L industry

	Stay in Innovation		Move to Managerial	
	Stay in L	Move to other	Stay in L	Move to other
	(1)	(2)	(3)	(4)
A_{soy}	-0.002	-0.026**	0.035**	-0.009
	[0.018]	[0.011]	[0.016]	[0.011]
Observations	2,862	2,862	2,862	2,862
R-squared	0.678	0.677	0.603	0.587
Microregion FE	Yes	Yes	Yes	Yes
Region x Year FE	Yes	Yes	Yes	Yes
Rural x Year FE	Yes	Yes	Yes	Yes
Other Controls x Year FE	Yes	Yes	Yes	Yes

Notes: The dependent variables are the wage bill share of skilled innovation workers who stay in innovation occupations and stay in low-R&D (L) industries (column 1), the wage bill share of skilled innovation workers who stay in innovation occupations but move to other industries or outside manufacturing (column 2), the wage bill share of skilled innovation workers who move to managerial occupations and stay in L industries (column 3), and the wage bill share of skilled innovation workers who move to managerial occupations and move to other industries or outside manufacturing (column 4), all measured using the Yearly Social Security Data (RAIS). Each share is computed as the wage bill of the indicated worker category divided by the total wage bill of all L-industry skilled innovation workers in that microregion and year. The sample includes workers identified as innovators in 2000, 2001, or 2002. All regressions include microregion fixed effects, region × year fixed effects, and the share of rural population interacted with year fixed effects. "Other Controls × Year FE" include income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census and interacted with year fixed effects. All regressions control for a measure of technical change in maize. Standard errors clustered at the microregion level are reported in brackets. Significance levels: $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

Appendix

The theory appendix is organized in three parts. Appendix A contains the baseline one-sector model and proofs. Appendix B derives the two-sector model and its comparative statics. Appendix C provides a CES foundation and local mapping to the reduced-form representation used in the text.

A Baseline One-Sector Model

This appendix introduces the baseline one-sector framework used in Section 2.2. It defines notation, states the key comparative statics, and provides proofs.

A.1 Setup

The economy has three occupations: scientists S, managers M, and low-skill production workers L. Endowments satisfy

$$S + M = \bar{H}, \qquad L = \bar{L}.$$

Innovation and implementation technolog are

$$I = S^{\phi}, \qquad B = \frac{M}{L}, \qquad q(B) = \frac{B^{\chi}}{1 + B^{\chi}}, \qquad N = S^{\phi} q(B),$$

with parameters $\phi > 0$ and $\chi > 0$. Output is given by

$$Y = N^{\kappa} L^{\theta} M^{1-\theta}, \qquad \kappa > 0,\ \theta \in (0, 1),$$

and we define TFP as

$$TFP \propto N^{\kappa}.$$

Let w_H and w_L denote high- and low-skill wages, and define $\omega \equiv w_H / w_L$. Let y denote

firm revenue (up to constant markup terms). FOCs imply[30]:

$$w_H S = y\,\kappa\phi, \tag{11}$$

$$w_H M = y\,\widetilde{A}^H(B), \tag{12}$$

$$w_L L = y\,A^L(B), \tag{13}$$

where

$$\widetilde{A}^H(B) \equiv (1-\theta) + \kappa\chi\big(1 - q(B)\big), \tag{14}$$

$$A^L(B) \equiv \theta - \kappa\chi\big(1 - q(B)\big), \tag{15}$$

and $\widetilde{A}^H(B) + A^L(B) = 1$.

Combining (12) and (13):

$$\omega = \Omega(B) \equiv \frac{1}{B}\frac{\widetilde{A}^H(B)}{A^L(B)}. \tag{16}$$

Define

$$X(B) \equiv \kappa\chi^2 \frac{B^\chi}{(1+B^\chi)^2} > 0.$$

Then

$$-\frac{d\ln B}{d\ln\omega} := \frac{1}{1 + \dfrac{X(B)}{\widetilde{A}^H(B)A^L(B)}} \in (0,1). \tag{17}$$

Hence $dB/d\omega < 0$: a higher skill premium tightens managerial bandwidth.

A.2 Comparative statics

Proposition A1 (Low-skill labor inflow)**.** *Holding $\bar{H}$ fixed, an increase in $\bar{L}$ implies:*

1. *the skill premium ω rises;*
2. *bandwidth B falls and implementation success $q(B)$ falls;*
3. *high-skill labor reallocates toward management (M/S rises);*
4. *TFP and output per worker decline;*
5. *total output is ambiguous.*

[30]The log of revenue, $\ln\Pi + \text{const}$, is concave in $(\ln S, \ln M, \ln L)$: the Hessian has rank one with a single negative eigenvalue $-2\kappa\mu\chi^2 q(B)(1-q(B))$. Hence revenue is log-concave in (S, M, L), and the first-order conditions characterize the unique profit maximum.

Proof. Aggregate factor payments imply

$$\omega \frac{\bar{H}}{\bar{L}} := \frac{\kappa\phi + \widetilde{A}^H(B(\omega))}{A^L(B(\omega))} \equiv \Psi(\omega). \tag{18}$$

To sign the skill-premium response, we proceed in two steps. First, we derive the bandwidth as a function of endowments. From Equations (11) and (12),

$$\frac{M}{S} = \frac{\widetilde{A}^H(B)}{\kappa\phi}.$$

Using $S + M = \bar{H}$ and $B = M/\bar{L}$:

$$B = \frac{\bar{H}}{\bar{L}}\, G(B), \qquad G(B) \equiv \frac{\widetilde{A}^H(B)}{\kappa\phi + \widetilde{A}^H(B)}.$$

Because $q'(B) > 0$, $1 - q(B)$ decreases in B, so $\widetilde{A}^{H\prime}(B) < 0$ and therefore $G'(B) < 0$. Log-differentiation gives

$$\frac{d\ln B}{d\ln \bar{L}} = -\frac{1}{1 - \frac{d\ln G}{d\ln B}} < 0.$$

Hence an increase in $\bar{L}$ decreases bandwidth. Second, we define B in terms of the skill premium via Equation (16):

$$\omega(B) = \frac{1}{B}\, F(B), \qquad F(B) \equiv \frac{\widetilde{A}^H(B)}{A^L(B)}.$$

Let $u(B) \equiv 1 - q(B)$, with $u'(B) < 0$. Then

$$F(u) = \frac{(1-\theta) + \kappa\chi u}{\theta - \kappa\chi u}, \qquad \frac{dF}{du} = \frac{\kappa\chi}{\big(\theta - \kappa\chi u\big)^2} > 0,$$

so $F'(B) < 0$. Therefore

$$\frac{d\omega}{dB} = \frac{BF'(B) - F(B)}{B^2} < 0.$$

Combining $dB/d\bar{L} < 0$ and $d\omega/dB < 0$ yields $d\omega/d\bar{L} > 0$.

From Equation (16), when ω increases, B decreases and thus $q(B)$ decreases. From Equations (11) and (12), $M/S = \widetilde{A}^H(B)/(\kappa\phi)$, so a decrease in B implies an increase in M/S. Finally,

$$\ln TFP = \text{const} + \kappa\big(\phi \ln S + \ln q(B)\big),$$

and both channels (lower S relative to M, and lower $q(B)$) reduce productivity. This lowers

output per worker. Total output combines a scale effect from larger L with a productivity effect from lower N, so its sign is ambiguous. □

Proposition A2 (Balanced scale benchmark)**.** *Consider a proportional scaling of factor endowments:*

$$(\bar{H}, \bar{L}) \mapsto (\lambda \bar{H}, \lambda \bar{L}), \qquad \lambda > 1.$$

Then ω, B, and S/M are unchanged, while

$$\frac{d \ln TFP}{d \ln \lambda} = \kappa\phi, \qquad \frac{d \ln Y}{d \ln \lambda} = 1 + \kappa\phi.$$

Proof. Under $(\bar{H}, \bar{L}) \mapsto (\lambda \bar{H}, \lambda \bar{L})$, the ratio $\bar{H}/\bar{L}$ is unchanged. By (18), ω and therefore B are unchanged. Then S scales with $\bar{H}$, while $q(B)$ is constant, so $N = S^{\phi} q(B)$ scales with λ^{ϕ}. Hence $TFP \propto N^{\kappa}$ scales with elasticity $\kappa\phi$, and output gains an additional unit elasticity from factors. □

Proposition A3 (High-skill inflow benchmark)**.** *Holding $\bar{L}$ fixed, an increase in $\bar{H}$ lowers ω, raises B, reallocates high-skill labor toward science, and increases TFP and output per low-skill worker.*

Proof. With $\bar{L}$ fixed, increasing $\bar{H}$ raises the left-hand side of (18) for any given ω, so equilibrium requires that ω decreases. By (16), this implies that B and $q(B)$ increase. From (12)/(11), $M/S = \widetilde{A}^{H}(B)/(\kappa\phi)$ falls, so labor reallocates toward science. Both higher S and higher $q(B)$ increase $N = S^{\phi} q(B)$ and therefore increase TFP and output per low-skill worker. □

B Two-Sector Extension: Proof of Proposition A1

B.1 Setup and equilibrium objects

Consider industries $k \in \{H, L\}$ with common factor prices (w_H, w_L) and common implementation technology. Industries differ only in

$$\phi_H > \phi_L, \qquad \theta_L > \theta_H.$$

Thus the H-industry is more innovation-intensive and less low-skill-intensive.

Industry output is

$$y_k = \left(S_k^{\phi_k} q(B_k)\right)^{\kappa} L_k^{\theta_k} M_k^{1-\theta_k}, \tag{19}$$

with

$$B_k = \frac{M_k}{L_k}, \qquad q(B) = \frac{B^\chi}{1+B^\chi}, \qquad \chi > 0. \tag{20}$$

Factor markets clear:

$$S_H + M_H + S_L + M_L = \bar{H}, \qquad L_H + L_L = \bar{L}.$$

Firm FOCs can be written as

$$w_H S_k = y_k\, \kappa\phi_k, \qquad w_H M_k = y_k\, \widetilde{A}_k^H(B_k), \qquad w_L L_k = y_k\, A_k^L(B_k), \tag{21}$$

where

$$\widetilde{A}_k^H(B) \equiv (1-\theta_k) + \kappa\chi\big(1-q(B)\big), \tag{22}$$

$$A_k^L(B) \equiv \theta_k - \kappa\chi\big(1-q(B)\big), \tag{23}$$

$$A_k^H(B) \equiv \kappa\phi_k + \widetilde{A}_k^H(B). \tag{24}$$

The within-production elasticities satisfy

$$\widetilde{A}_k^H(B) + A_k^L(B) = 1.$$

Combining manager and low-skill FOCs:

$$\omega \equiv \frac{w_H}{w_L} = \Omega_k(B_k) \equiv \frac{1}{B_k}\frac{\widetilde{A}_k^H(B_k)}{A_k^L(B_k)}. \tag{25}$$

Define the pass-through objects

$$K_k \equiv K_k(B_k) \equiv \frac{X(B_k)}{\widetilde{A}_k^H(B_k)A_k^L(B_k)}, \qquad -\frac{d\ln B_k}{d\ln\omega} \equiv \frac{1}{1+K_k}, \qquad X(B) \equiv \kappa\chi^2\frac{B^\chi}{(1+B^\chi)^2} > 0. \tag{26}$$

Hence $dB_k/d\omega < 0$ and $dq(B_k)/d\omega < 0$.

Demand closure for price movements. To pin down movements in ω, we use a small-open-economy closure with fixed revenue weights $\lambda_H, \lambda_L > 0$ and $\lambda_H + \lambda_L = 1$, so

$$y_k = \lambda_k\, y_{\text{tot}}.$$

We take y_{tot} as fixed and exogenous under this closure, so $d \ln y_k = 0$ for $k \in \{H, L\}$. We also take w_L as numeraire. Define implemented varieties and productivity:

$$N_k \equiv S_k^{\phi_k} q(B_k), \qquad TFP_k \propto N_k^{\kappa}.$$

B.2 Proposition and proof

Lemma 1 (Bandwidth ordering from factor intensity)**.** *In an interior equilibrium with common factor prices and $\theta_L > \theta_H$, one has*

$$B_L^* < B_H^*.$$

Proof. From (25),

$$\Omega_k(B) = \frac{1}{B} \frac{\widetilde{A}_k^H(B)}{A_k^L(B)} = \frac{1}{B} \frac{(1-\theta_k) + \kappa\chi(1-q(B))}{\theta_k - \kappa\chi(1-q(B))}.$$

Holding B fixed and differentiating with respect to θ_k:

$$\frac{\partial \Omega_k}{\partial \theta_k}(B) = -\frac{1}{B\,[A_k^L(B)]^2} < 0$$

on the interior admissible set $A_k^L(B) > 0$. Therefore $\theta_L > \theta_H$ implies $\Omega_L(B) < \Omega_H(B)$ pointwise in B.

By (26), $\Omega_k'(B) < 0$, so each Ω_k is strictly decreasing and invertible. At equilibrium, both sectors face the same skill premium ω^*, hence

$$B_L^* = \Omega_L^{-1}(\omega^*) < \Omega_H^{-1}(\omega^*) = B_H^*.$$

□

Proposition A1 (Two-sector implications of a low-skill inflow)**.** *Assume an interior equilibrium with $\chi > 0$. By Lemma 1, $\theta_L > \theta_H$ implies $B_L^* < B_H^*$. Define $K_k^* \equiv K_k(B_k^*)$ and*

$$\textit{(A1)} \quad A_L^L(B_L^*) > A_H^L(B_H^*), \tag{27}$$

$$\textit{(A2)} \quad \frac{\widetilde{A}_L^H(B_L^*)}{1+K_L^*} - \frac{\widetilde{A}_H^H(B_H^*)}{1+K_H^*} > \kappa(\phi_H - \phi_L), \tag{28}$$

$$\textit{(A3)} \quad \chi\left[\frac{1-q(B_L^*)}{1+K_L^*} - \frac{1-q(B_H^*)}{1+K_H^*}\right] > (\phi_H - \phi_L). \tag{29}$$

For an increase in $\bar{L}$ holding $\bar{H}$ fixed:

1. $d\omega/d\bar{L} > 0$ *and* $d\ln q(B_k)/d\ln \bar{L} < 0$ *for* $k \in \{H, L\}$*;*

2. *Equation* (27) *implies*
$$\frac{d\ln L_L}{d\ln \bar{L}} > \frac{d\ln L_H}{d\ln \bar{L}};$$

3. $d\ln(y_k/L_k)/d\ln \bar{L} < 0$ *in both industries, and the decline is larger in* L *if Condition* (28) *holds;*

4. $d\ln TFP_k/d\ln \bar{L} < 0$ *in both industries, and the decline is larger in* L *if Condition* (29) *holds.*

Proof. Step 1: an increase in $\bar{L}$ *implies an increase in* ω. Using the closure $y_k = \lambda_k y_{\text{tot}}$ and aggregating (21):

$$w_H \bar{H} = y_{\text{tot}} \sum_{k\in\{H,L\}} \lambda_k A_k^H(B_k), \qquad w_L \bar{L} = y_{\text{tot}} \sum_{k\in\{H,L\}} \lambda_k A_k^L(B_k). \tag{30}$$

Dividing:
$$\omega \frac{\bar{H}}{\bar{L}} = \frac{\lambda_H A_H^H(B_H(\omega)) + \lambda_L A_L^H(B_L(\omega))}{\lambda_H A_H^L(B_H(\omega)) + \lambda_L A_L^L(B_L(\omega))} \equiv \Psi(\omega). \tag{31}$$

From (25), $dB_k/d\omega < 0$. Since $q'(B) > 0$:

$$\frac{dq(B_k)}{d\omega} < 0, \qquad \frac{dA_k^H}{d\omega} = \kappa\chi\left(-\frac{dq(B_k)}{d\omega}\right) > 0, \qquad \frac{dA_k^L}{d\omega} = \kappa\chi\frac{dq(B_k)}{d\omega} < 0.$$

Hence $\Psi(\omega)$ is increasing in ω. Holding ω fixed, an increase in $\bar{L}$ lowers the left-hand side $\omega(\bar{H}/\bar{L})$. Under the maintained unique-equilibrium crossing in (31), the rebalanced equilibrium therefore requires a higher skill premium. Thus $d\omega/d\bar{L} > 0$. From (25), this implies $dB_k/d\bar{L} < 0$ and $dq(B_k)/d\bar{L} < 0$ in both industries.

Step 2: labor absorption ordering. From (21) and the closure,

$$L_k = \frac{\lambda_k y_{\text{tot}}}{w_L} A_k^L(B_k).$$

Taking log differences across sectors eliminates common terms:

$$d\ln L_L - d\ln L_H = d\ln A_L^L - d\ln A_H^L.$$

Define
$$E_k \equiv -\frac{d\ln A_k^L(B_k)}{d\ln \omega}.$$

Since $A_k^L(B) = \theta_k - \kappa\chi(1 - q(B))$ and $X(B) \equiv \kappa\chi B q'(B) = \kappa\chi^2 \dfrac{B^\chi}{(1+B^\chi)^2}$,

$$\frac{dA_k^L}{dB_k} = \kappa\chi q'(B_k) = \frac{X(B_k)}{B_k}, \qquad \frac{d\ln B_k}{d\ln\omega} = -\frac{1}{1+K_k},$$

so

$$E_k = \frac{X(B_k)}{A_k^L(B_k)\big(1+K_k\big)}.$$

Therefore

$$\frac{d\ln L_L}{d\ln \bar{L}} - \frac{d\ln L_H}{d\ln \bar{L}} = E_\omega\,(E_H - E_L), \qquad E_\omega \equiv \frac{d\ln\omega}{d\ln\bar{L}} > 0.$$

Evaluated at the equilibrium, condition (27) is exactly $E_H^* > E_L^*$.[31] Hence:

$$\frac{d\ln L_L}{d\ln \bar{L}} > \frac{d\ln L_H}{d\ln \bar{L}}.$$

Step 3: output per worker. Using (19) and $B_k = M_k/L_k$,

$$\ln\left(\frac{y_k}{L_k}\right) = \kappa\phi_k \ln S_k + \kappa \ln q(B_k) + (1-\theta_k)\ln B_k.$$

Under the closure, $d\ln y_k/d\ln\bar{L} = 0$, and from the scientist FOC $S_k = (y_k\kappa\phi_k)/w_H$, so $d\ln S_k/d\ln\bar{L} = -E_\omega$. Hence

$$\frac{d\ln(y_k/L_k)}{d\ln\bar{L}} = E_\omega\left[-\kappa\phi_k + (1-\theta_k)\frac{d\ln B_k}{d\ln\omega} + \kappa\frac{q'(B_k)}{q(B_k)}\frac{dB_k}{d\ln\omega}\right].$$

Because $d\ln B_k/d\ln\omega < 0$, $dB_k/d\ln\omega < 0$, and $q'(B_k) > 0$, every term in brackets is negative, so $d\ln(y_k/L_k)/d\ln\bar{L} < 0$ for both sectors.

For the cross-sector comparison:

$$\frac{d\ln(y_L/L_L)}{d\ln\bar{L}} - \frac{d\ln(y_H/L_H)}{d\ln\bar{L}} = E_\omega\left[\kappa(\phi_H - \phi_L) + \frac{\widetilde{A}_H^H(B_H^*)}{1+K_H^*} - \frac{\widetilde{A}_L^H(B_L^*)}{1+K_L^*}\right].$$

Condition (28) makes the bracket negative, implying a larger decline in y_L/L_L.

[31] To see this, note that E_k can be rewritten as $E_k = \frac{1}{\frac{A_k^L(B_k)}{X(B_k)} + \frac{1}{\widetilde{A}_k^H(B_k)}}$. Since $\widetilde{A}_k^H = 1 - A_k^L$, E_k is strictly decreasing in A_k^L for fixed X and strictly increasing in X for fixed A_k^L. By Lemma 1, $B_L^* < B_H^*$. A direct calculation gives $X'(B) = \kappa\chi^3 B^{\chi-1}(1-B^\chi)/(1+B^\chi)^3$, so $X'(B) > 0$ whenever $B < 1$. Because $B = M/L < 1$ empirically (managers are scarcer than production workers; in our calibration $B \approx 0.15$), both equilibrium bandwidths satisfy $B_L^* < B_H^* < 1$, and therefore $X(B_L^*) \le X(B_H^*)$. Combined with Condition (27), both channels reinforce: the L-industry has a larger A^L *and* a smaller X, so $E_H^* > E_L^*$.

Step 4: TFP.

$$d\ln TFP_k = \kappa(\phi_k\, d\ln S_k + d\ln q(B_k))\,.$$

Using $d\ln S_k/d\ln \bar{L} = -E_\omega$ and $d\ln q(B_k)/d\ln \bar{L} = \chi(1-q(B_k))\, d\ln B_k/d\ln \bar{L} = -\chi(1-q(B_k))E_\omega/(1+K_k^*)$:

$$\frac{d\ln TFP_k}{d\ln \bar{L}} = -\kappa E_\omega\left[\phi_k + \chi\frac{1-q(B_k^*)}{1+K_k^*}\right] < 0.$$

Thus TFP declines in both sectors. Moreover,

$$\frac{d\ln TFP_L}{d\ln \bar{L}} - \frac{d\ln TFP_H}{d\ln \bar{L}} = -\kappa E_\omega\left[\phi_L - \phi_H + \chi\left(\frac{1-q(B_L^*)}{1+K_L^*} - \frac{1-q(B_H^*)}{1+K_H^*}\right)\right].$$

Condition (29) makes the bracket positive, so the difference is negative: TFP falls more in industry L. □

Interpretation of A1–A3. (A1) is the exact labor-absorption condition comparing the ω-elasticity of effective low-skill demand across industries. (A2) controls whether output-per-worker falls more in L despite H's higher ϕ. (A3) controls whether the implementation bottleneck channel dominates enough for TFP to fall more in L.

C Innovation Intensity and Endogenous Factor Intensity under CES Production

This appendix shows that cross-industry differences in innovation intensity ϕ_k can generate differences in effective factor intensity under CES production, through the same implementation bottleneck used in the baseline model.

C.1 CES environment with bottlenecks

For each industry $k \in \{H, L\}$:

$$I_k = S_k^{\phi_k}, \qquad N_k = S_k^{\phi_k} q(B_k), \qquad B_k = \frac{M_k}{L_k}, \qquad q(B) = \frac{B^\chi}{1+B^\chi}.$$

Define the CES composite $X_k \equiv N_k^\kappa M_k$ and output

$$Y_k = \left(X_k^\rho + L_k^\rho\right)^{1/\rho}, \qquad \rho = \frac{\sigma-1}{\sigma} \in (0,1). \tag{32}$$

Let

$$s_{X,k} \equiv \frac{X_k^\rho}{X_k^\rho + L_k^\rho}, \qquad s_{L,k} \equiv 1 - s_{X,k},$$

and profits

$$\Pi_k = Y_k^\mu - w_H(S_k + M_k) - w_L L_k, \qquad y_k \equiv \mu Y_k^\mu.$$

C.2 FOCs and the bottleneck wedge

FOCs multiplied by input levels are

$$w_H S_k = y_k \, \kappa\phi_k \, s_{X,k}, \tag{33}$$

$$w_H M_k = y_k \, s_{X,k}\Big[1 + \kappa\chi\big(1 - q(B_k)\big)\Big], \tag{34}$$

$$w_L L_k = y_k\Big[s_{L,k} - \kappa\chi\big(1 - q(B_k)\big)s_{X,k}\Big]. \tag{35}$$

Define the bottleneck term

$$\Xi_k \equiv \kappa\chi\big(1 - q(B_k)\big). \tag{36}$$

Then

$$A_k^L(B_k, s_{X,k}) \equiv \frac{w_L L_k}{y_k} = s_{L,k} - \Xi_k s_{X,k} = 1 - (1 + \Xi_k)s_{X,k}. \tag{37}$$

Relative to frictionless CES shares, bottlenecks reduce the effective low-skill side and raise managerial load.

C.3 Local mapping to an effective Cobb-Douglas exponent

We use a local mapping around an interior equilibrium $(\omega^*, B_k^*, s_{X,k}^*)$.

Proposition A1 (Reference-point CES-Cobb-Douglas mapping)**.** *Define*

$$\Xi_k^* \equiv \kappa\chi\big(1 - q(B_k^*)\big), \qquad \theta_k \equiv (1 + \Xi_k^*)\big(1 - s_{X,k}^*\big). \tag{38}$$

Consider the reduced-form,

$$\widetilde{Y}_k \propto N_k^\kappa L_k^{\theta_k} M_k^{1-\theta_k},$$

with the same bottleneck $q(\cdot)$ and same aggregate closure as in the baseline model. Then the reduced form reproduces the low-skill expenditure object exactly at the reference equilibrium. Around that point, the first-order gap is governed by both CES-share responses and bottleneck-coefficient responses.

Proof. At the reference equilibrium, (37) gives

$$A_k^L(B_k^*, s_{X,k}^*) = 1 - (1 + \Xi_k^*)s_{X,k}^*.$$

Using (38), the reduced-form counterpart $\widetilde{A}_k^L \equiv \theta_k - \Xi_k^*$ satisfies

$$\widetilde{A}_k^L = (1 + \Xi_k^*)(1 - s_{X,k}^*) - \Xi_k^* = 1 - (1 + \Xi_k^*)s_{X,k}^* = A_k^L(B_k^*, s_{X,k}^*).$$

Hence levels match at the reference point. For local comparative statics,

$$dA_k^L = -s_{X,k}^*\, d\Xi_k - (1 + \Xi_k^*)\, ds_{X,k}.$$

The reduced-form counterpart is $\widetilde{A}_k^L = \theta_k - \Xi_k$, so $d\widetilde{A}_k^L = -d\Xi_k$. Therefore the first-order gap is

$$dA_k^L - d\widetilde{A}_k^L = (1 - s_{X,k}^*)\, d\Xi_k - (1 + \Xi_k^*)\, ds_{X,k}.$$

Thus local approximation quality depends on both channels: movements in the CES share $s_{X,k}$ and movements in bottleneck tightness Ξ_k. □

C.4 Implication for cross-industry factor intensity

The mapping

$$\theta_k = (1 + \Xi_k^*)\big(1 - s_{X,k}^*\big)$$

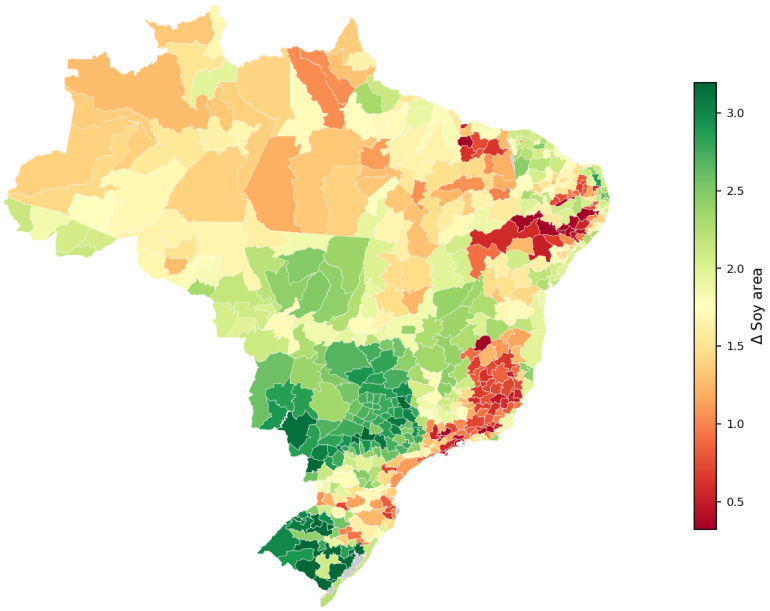
links effective low-skill intensity to both the CES share $s_{X,k}^*$ and bottleneck tightness Ξ_k^*.

Higher ϕ_k raises the return to scientists (via (33)), tends to raise $s_{X,k}^*$, and tends to relax bottlenecks (lower Ξ_k^*) through greater innovation capacity. Both forces reduce θ_k. Therefore, industries with higher innovation intensity naturally appear less low-skill intensive in the reduced form used in the main text.

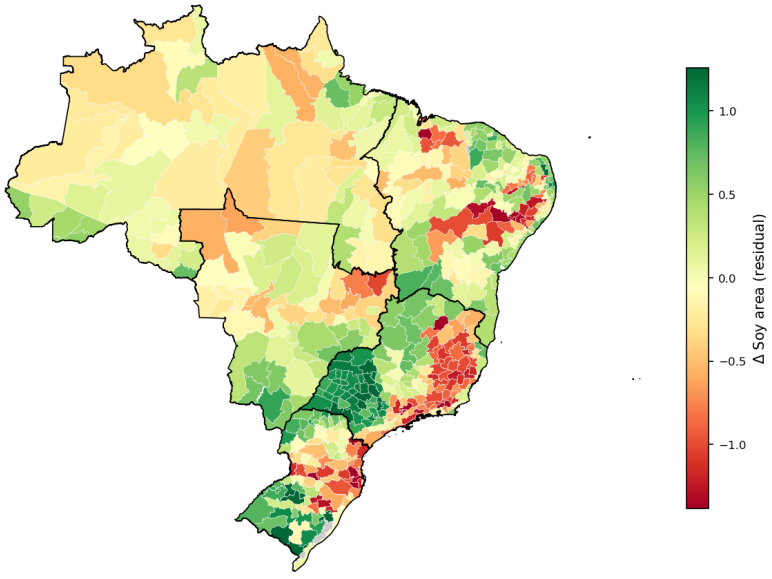
D Appendix Figures

Figure D.1: Δ in Potential Soy Yield 2000–2010

(a) Soy area change



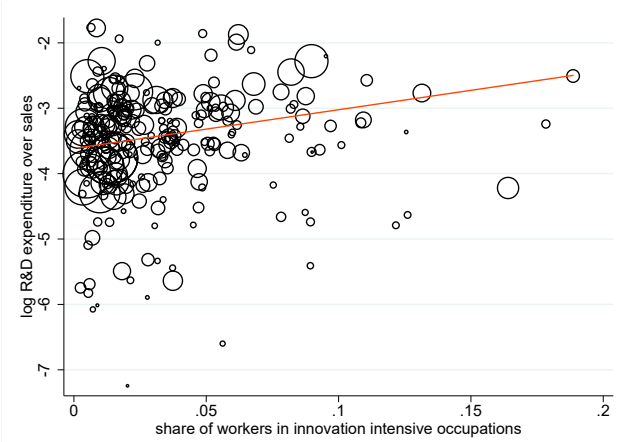
(b) Soy area change (residual, within macroregion)



Notes: Authors' calculations from FAO-GAEZ data. Panel (a) shows the change in potential soy yield between 2000 and 2010 for each microregion, computed as the difference between average potential yield under high inputs and average potential yield under low inputs. Panel (b) shows the residual variation after removing macroregion fixed effects. Black lines delimit Brazil's five macroregions (Norte, Nordeste, Sudeste, Sul, Centro-Oeste). Colors range from red (low) to green (high).

Figure D.2: Correlations between share of workers in innovation intensive occupations and industry-level measures of innovation

(a) R&D expenditure/sales



(b) Patent share



(c) Product innovation share



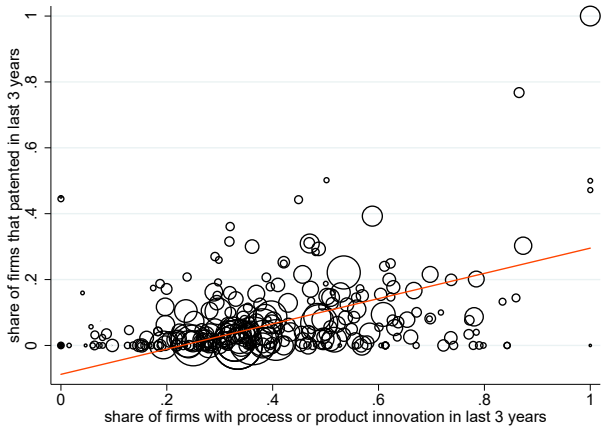
(d) Process innovation share



(e) Product or Process innovation share



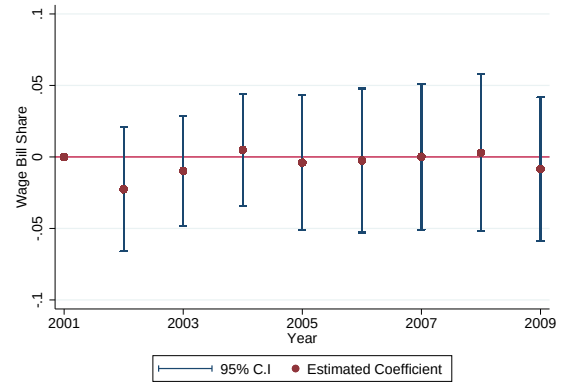
(f) Patents vs Product or Process innovation share



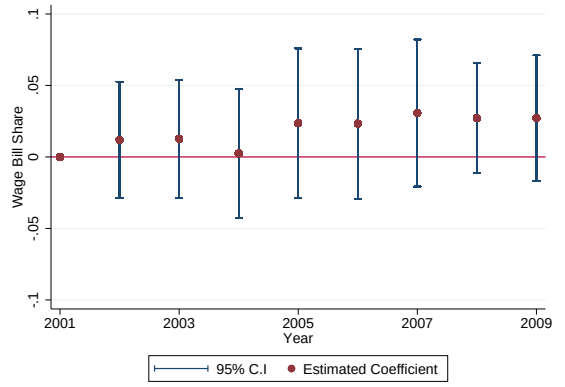
Notes: The share of workers in innovation intensive occupations in each sector is constructed using RAIS data for the year 2000 and the methodology described in Section 3.4. All measures of innovation at industry level (R&D expenditure over sales, patent share, product and process innovation share) are computed using the 2000 *Pesquisa de Inovação Tecnológica (PINTEC)*. Dot size captures size of the industry in terms of number of employees in 2000.

Figure D.3: The effect of agricultural technical change on occupational transitions of innovation workers in H industries

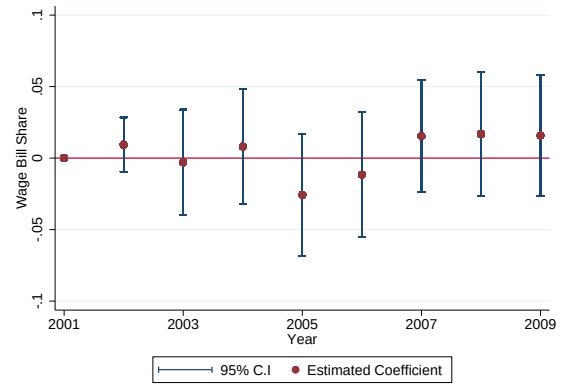
(a) **Stay in Innovation, Stay in H Industry**



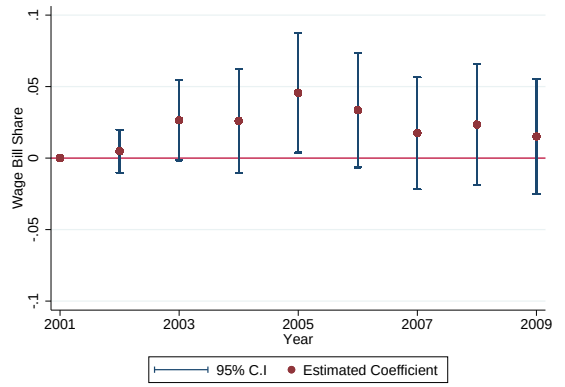
(b) **Stay in Innovation, Move to Other**



(c) **Move to Managerial, Stay in H Industry**



(d) **Move to Managerial, Move to Other**



Notes: The figure shows the point estimates and the 95% confidence intervals for the estimates of the β_j of equation (9) where the outcome variable is the wage bill share of the indicated category of skilled innovation workers initially employed in high-R&D (H) industries in microregion k located in region r at the end of year t in the manufacturing sector. Each share is computed as the wage bill of the indicated worker category divided by the total wage bill of all H-industry skilled innovation workers in that microregion and year. Workers are identified as innovators in 2000, 2001, or 2002. Standard errors are clustered at the microregion level.

E Appendix Tables

Table E.1: Effect of Agricultural Technical Change on GE Soy Adoption

Panel A	Δ GE-soy area share (1)	Δ GE-soy area share (2)	Δ non-GE soy area share (3)	Δ non-GE soy area share (4)
ΔA^{soy}	0.022***	0.020***	-0.007*	-0.008**
	[0.005]	[0.004]	[0.004]	[0.004]
Share rural population	0.034***	0.117***	-0.009	-0.057**
	[0.010]	[0.023]	[0.009]	[0.023]
Log Income per capita		-0.009		-0.002
		[0.006]		[0.007]
Literacy rate		0.162***		-0.043
		[0.034]		[0.035]
Log population density		0.005***		-0.006***
		[0.001]		[0.001]
Observations	557	557	557	557
R-squared	0.094	0.208	0.013	0.053

Panel B	Δ Soy area share (1)	Δ Soy area share (2)	Δ Maize area share (3)	Δ Maize area share (4)
ΔA^{soy}	0.022***	0.016***	-0.006	0.000
	[0.004]	[0.004]	[0.004]	[0.004]
ΔA^{maize}	-0.003**	-0.001	0.005***	0.003*
	[0.001]	[0.001]	[0.002]	[0.002]
Share rural population	0.029***	0.064***	0.020***	0.012
	[0.007]	[0.013]	[0.008]	[0.015]
Log Income per capita		-0.010*		-0.011
		[0.006]		[0.007]
Literacy rate		0.122***		-0.002
		[0.018]		[0.023]
Log population density		-0.001		0.003***
		[0.001]		[0.001]
Observations	557	557	556	556
R-squared	0.135	0.245	0.041	0.066

Notes: Changes in dependent variables are calculated over the years 1996 and 2006 (source: Agricultural Census). The unit of observation is the microregion. Robust standard errors reported in brackets. Significance levels: $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.1$.

Table E.2: Summary Statistics: By Industry

		Panel A: Decadal Variables by Industry				
		2000		Δ2000-2010		
	Source	Mean	SD	Mean	SD	Observations
Employment Shares	*Population Census*					
Agriculture		0.279	0.140	-0.050	0.055	557
Manufacturing L Industry		0.081	0.055	0.007	0.033	557
Manufacturing H Industry		0.067	0.043	-0.001	0.025	557
Services		0.573	0.118	0.044	0.057	557
Log. Employment	*Population Census*					
Agriculture		8.268	0.890	0.122	0.249	557
Manufacturing L Industry		7.076	1.569	0.358	0.400	557
Manufacturing H Industry		6.898	1.485	0.309	0.394	557
Services		9.194	1.287	0.404	0.175	557

	Panel B: Yearly Variables by Industry			
	Source	Mean	SD	Observations
Manufacturing Employment	*RAIS (1998-2009)*			
Log. Total Employment				
Manufacturing L Industry		6.725	1.929	6,064
Manufacturing H Industry		6.454	1.962	6,064
Log. Low-Skill Employment				
Manufacturing L Industry		6.392	1.916	6,064
Manufacturing H Industry		6.127	1.903	6,064
Log. High-Skill Employment				
Manufacturing L Industry		5.258	2.109	6,064
Manufacturing H Industry		4.890	2.287	6,064
Manufacturing Productivity	*PIA (1998-2009)*			
Log. Value Added per Worker				
Manufacturing L Industry		10.278	0.723	3,360
Manufacturing H Industry		10.951	0.929	3,360
Log. Value Added per Wage Bill				
Manufacturing L Industry		1.329	0.485	3,360
Manufacturing H Industry		1.622	0.606	3,360
Log. Total Factor Productivity				
Manufacturing L Industry		5.608	0.736	3,360
Manufacturing H Industry		5.720	0.706	3,360
Log. Capital	*PIA (1998-2009)*			
Manufacturing L Industry		16.799	2.173	3,360
Manufacturing H Industry		16.335	2.559	3,360

Notes: The data sources are the Population Census (2000 and 2010), RAIS (1998-2009) and PIA (1998-2009). Manufacturing industries are classified as Low-R&D or High-R&D intensive depending on whether their R&D intensity is below or above the median in 2000 (weighting industries by number of employees so that each group captures around 50 percent of total manufacturing employment). We define R&D intensity as R&D expenditure as a share of total sales at baseline and we source it from the 2000 *Pesquisa de Inovação Tecnológica* (PINTEC).

Table E.3: Classification of Manufacturing Industries by R&D Intensity

IBGE Code	Description	R&D Share of Sales	Skill Intensity
26091	Ceramic products	0.106	0.275
34001	Manufacturing and assembly of motor vehicles	0.105	0.738
23030	Production of nuclear fuels	0.100	0.830
31002	Electrical material for vehicles	0.088	0.599
27001	Steel products	0.084	0.659
35030	Construction, assembly and repair of airplanes	0.080	0.875
28002	Foundries, stamping shops, powder metallurgy and metal treatment services	0.066	0.502
33003	Machines, equipment for electronic systems for industrial automation, and control	0.064	0.848
24020	Pharmaceutical products	0.062	0.809
33001	Medical equipment	0.061	0.753
29002	Appliances	0.058	0.709
34002	Cabins, car bodies, trailers and parts for motor vehicles	0.058	0.637
20000	Wooden products	0.055	0.247
33004	Equipment, instruments and optical, photographic and cinematographic material	0.055	0.709
33002	Measuring, testing and control equipment - except for controlling industrial processes	0.054	0.725
24010	Paints, dyes, varnish, enamels and lacquers	0.053	0.656
25020	Plastic products	0.052	0.543
32000	Electronic material and communications equipment	0.052	0.757
31001	Machines, equipment and miscellaneous electric material - except for vehicles	0.051	0.678
27003	Foundries	0.051	0.462
15043	Other food products	0.049	0.426
36090	Miscellaneous products	0.048	0.576
23010	Coke plants	0.047	0.487
37000	Recycling	0.045	0.304
35090	Miscellaneous transportation equipment	0.044	0.581
21002	Corrugated cardboard, packaging, and paper and cardboard objects	0.044	0.577
17001	Processing of fibers, weaving and cloth making	0.043	0.471
28001	Metal products - except machines and equipment	0.042	0.496
24030	Soap, detergents, cleaning products and toiletries	0.042	0.658
29001	Machines and equipment - except appliances	0.041	0.605
21001	Pulp, paper and smooth cardboard, poster paper and card paper	0.040	0.602
34003	Reconditioning or restoration of engines of motor vehicles	0.038	0.556
24090	Miscellaneous chemical products	0.037	0.635
25010	Rubber products	0.036	0.567
26092	Miscellaneous products of non-metallic minerals	0.035	0.382
22000	Editing, printing and reproduction of recordings	0.035	0.702
19012	Leather objects	0.034	0.453
30000	Office machines and data-processing equipment	0.034	0.852
36010	Pieces of furniture	0.034	0.402
26010	Glass and glass products	0.031	0.576
15021	Preserves of fruit, vegetables and other vegetable products	0.029	0.484
17002	Manufacturing of textile objects based on cloth - except for garments	0.028	0.433
18001	Making of clothing articles and accessories - except on order	0.023	0.425
18002	Making clothing articles and accessories - on order	0.023	0.435
18999	Making of clothing articles and accessories - on order or not	0.023	0.690
27002	Non-ferrous metals	0.022	0.644
15030	Dairy products	0.022	0.433
19020	Footwear	0.019	0.348
15010	Slaughtering and preparation of meat and fish	0.018	0.355
35010	Construction and repair of boats	0.018	0.493
23020	Products in oil refining	0.015	0.763
33005	Chronometers, clocks and watches	0.015	0.751
23400	Alcohol production	0.014	0.350
15041	Manufacturing and refining of sugar	0.013	0.334
15042	Roasting and grinding of coffee	0.013	0.499
19011	Tanning and other preparations of leather	0.013	0.325
16000	Tobacco products	0.013	0.496
15050	Beverages	0.012	0.555
15022	Vegetable fat and oil	0.009	0.446
35020	Construction and assembly of locomotives, cars and other rolling stock	0.004	0.632
	Median	0.041	0.432

Notes: The industry codes correspond to the CNAE-Domiciliar, the industry classification used in the 2000 Population Census. Industries are sorted by their R&D intensity at baseline. We measure R&D intensity as R&D expenditure as a share of total sales at baseline and we source it from the 2000 *Pesquisa de Inovação Tecnológica* (PINTEC). We define skill intensity as the share of skilled individuals in a particular industry in Brazil at baseline and we source it from the 2000 Population Census. The correlation between these variables is approximately 0.34. Industries below the median are classified as low and the ones above the median as high.

Table E.4: Keywords Used to Identify Innovative Occupations

Panel A: Nouns or combination of nouns from task description of occupations

Portuguese	English
pesquisa e desenvolvimento	research and development
inovação	innovation
p&d	R&D
desenvolvimento de produtos	product development
desenvolvimento de processos	process development
pesquisador	researcher
novas tecnologias	new technologies
protótipos	prototypes
pesquisas tecnologicas	technological research
automaçao de processos	process automation

Panel B: Actions (verb + noun) from task description of occupations

Portuguese	English
desenvolvem produtos	develop products
desenvolvem pesquisas	develop research
desenvolvem equipamentos	develop equipment
desenvolvem processos	develop processes
desenvolvem dispositivos	develop devices
otimizam métodos	optimize methods
otimizam os meios	optimize means
aperfeiçoam sistemas	improve systems
aperfeiçoam processos	improve processes
aperfeiçoam produtos	improve products
aperfeiçoam dispositivos	improve devices
implementam dispositivos de automaçao	implement automation devices
desenvolvem, testam e supervisionam sistemas, processos e métodos produtivos	develop, test and supervise systems, processes and production methods

Panel C: Nouns or combinations of nouns from PINTEC survey

Portuguese	English
produto novo / novo produto	new product
produtos novos / novos produtos	new products
produto aprimorado	improved product
produtos aprimorados	improved produts
inovação de produto	product innovation
aperfeiçoamento de produto	product improvement
processo novo / novo processo	new process
processos novos / novos processos	new processes
processo aprimorado	improved process
processos aprimorados	improved processes
inovação de processo	process innovation
aperfeiçoamento de processo	process improvement

Notes: The Table reports the keywords used to identify innovation intensive occupations and their English translation. Keywords reported in Panels A and B are sourced from the "Brazilian Classification of Occupations", Ministry of Labor, 3rd Edition (2010). Keywords reported in Panel C are sourced from the Technical Appendix of the 2008 PINTEC Survey.

Table E.5: Industry-level Measures of Innovation

Outcome:	Measure of Industrial Innovation				
Measure:	log R&D expenditure over sales	Share of patenting firms	Share of product innovation firms	Share of process innovation firms	Share of product or process innovation firms
	(1)	(2)	(3)	(4)	(5)
Share of innovation	5.929**	1.557***	2.178***	0.885**	1.628***
	(2.502)	(0.411)	(0.406)	(0.427)	(0.473)
Observations	271	274	274	274	274
R-squared	0.062	0.166	0.172	0.037	0.098

Notes: Dependent variables are calculated from 2000 *Pesquisa de Inovação Tecnológica* (PINTEC). The unit of observation is the 4-digit CNAE industry. These regressions compute the OLS coefficient of a number of outcomes on the share of innovation workers. The share of innovation workers in each industry are computed for the year 2000 using the methodology described in Section 3.4. All regressions are weighted by number of workers in each industry in 2000. Robust standard errors reported in brackets. Significance levels: $^{***}p < 0.01$,** $p < 0.05$,* $p < 0.1$.

Table E.6: Keywords Used to Identify Managerial Occupations

Panel A: Managerial keywords in occupation names

Portuguese	English
gerente	manager
supervisor / supervisora	supervisor
coordenador / coordenadora	coordinator
chefe	chief
administrador / administradora	administrator
dirigente	director / leader
responsável	person in charge
encarregado / encarregada	person in charge
gestor / gestora	manager
presidente	president
vice-presidente	vice-president
superintendente	superintendent
assessor / assessora	advisor
diretor / diretora	director

Panel B: Exclusion terms for "diretor" (artistic/creative directors)

Portuguese	English
cinema	cinema
fotografia	photography
arte / artístico / artística	art / artistic
musical	musical
teatro / teatral	theater / theatrical
televisão / tv	television / TV
filme	film
documentário	documentary
programa	program
espetáculo	show / spectacle
produção artística	artistic production
criação	creation
audiovisual	audiovisual
cênico / cênica	scenic

Notes: The Table reports the keywords used to identify innovation intensive occupations and their English translation. Keywords reported in Panels A and B are sourced from the "Brazilian Classification of Occupations", Ministry of Labor, 3rd Edition (2010). Keywords reported in Panel C are sourced from the Technical Appendix of the 2008 PINTEC Survey.

Table E.7: Effect of agricultural technical change on sectoral employment shares

Decadal Population Census Data (2000 and 2010)

	Agriculture (1)	Manufacturing (2)	Services (3)
Panel A: Overall			
ΔA_{soy}	-0.033***	0.025***	0.008*
	[0.005]	[0.005]	[0.004]
Observations	557	557	557
R-squared	0.246	0.166	0.359
Region FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Panel B: Low Skill			
ΔA_{soy}	-0.034***	0.027***	0.007
	[0.006]	[0.005]	[0.004]
Observations	557	557	557
R-squared	0.154	0.151	0.231
Region FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Panel C: High Skill			
ΔA_{soy}	-0.011***	0.013**	-0.001
	[0.004]	[0.005]	[0.006]
Observations	557	557	557
R-squared	0.049	0.104	0.076
Region FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Notes: The dependent variables are changes in the share of employment in agriculture, manufacturing, and services. Changes in dependent variables are calculated over the years 2000 and 2010 (source: Population Censuses). Panel A uses overall employment shares, Panel B uses employment shares among low-skill workers, and Panel C uses employment shares among high-skill workers. A worker is classified as skilled if he or she has completed high school. These regressions include as controls the share of rural population, income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census, a measure of technical change in maize, and region fixed effects. Robust standard errors are reported in brackets. Significance levels: $^{***}p < 0.01$,** $p < 0.05$,* $p < 0.1$.

Table E.8: Worker-level trajectories from agriculture to manufacturing, and wages of former agricultural workers

Sample:	Employed in Agriculture pre-2003		Employed in L industry from 2003, in Agriculture or L-industry pre-2003	
Outcomes:	Move to L industry	Move to H industry	log(wage)	
				composition-adjusted
	(1)	(2)	(3)	(4)
ΔA_{soy} in pre-2003 microregion	0.050***	-0.008***		
	[0.009]	[0.002]		
1(Employed in Agriculture pre-2003)			-0.271***	-0.118***
			[0.020]	[0.017]
Observations	1,304,659	1,304,659	951,453	951,453
R-squared	0.025	0.003	0.205	0.155
Microregion Controls	Yes	Yes	Yes	Yes

Notes: Microregion Controls include the share of rural population, income per capita (in logs), population density (in logs), literacy rate, the share of individuals earning minimum wage, all observed in 1991, and a measure of technical change in maize. Standard errors clustered at micro-region level reported in brackets. Significance levels: $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.1$.

Table E.9: Effect of agricultural technical change on skill premium and wages of skilled vs unskilled labor

Population Census (2000-2010)

	Formal Workers			Informal Workers		
	Δ Wage skilled (1)	Δ Wage unskilled (2)	Δ Skill premium (3)	Δ Wage skilled (4)	Δ Wage unskilled (5)	Δ Skill premium (6)
ΔA_{soy}	0.056***	0.020***	0.036**	0.025*	0.022***	0.003
	[0.018]	[0.008]	[0.017]	[0.014]	[0.008]	[0.014]
Observations	556	556	556	557	557	557
R-squared	0.074	0.268	0.072	0.133	0.362	0.039
Region FE	Yes	Yes	Yes	Yes	Yes	Yes
Rural	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variables are the change in the Mincerian wage premium for skilled workers (columns 1 and 4), the change in the Mincerian wage premium for unskilled workers (columns 2 and 5), and the change in the skill premium, defined as the difference between the two (columns 3 and 6). Changes in dependent variables are calculated over the years 2000 and 2010 (source: Population Censuses). Microregion-level wage premia are estimated from Mincerian regressions of log hourly wages on microregion fixed effects, controlling for skill, age, race, and sector interactions, separately for each year and formality status. Columns (1)–(3) use formal workers and columns (4)–(6) use informal workers. The sample is restricted to male workers. A worker is classified as skilled if he or she has completed high school. "Rural" represents the share of rural population in the microregion. "Other Controls" include income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census, as well as a measure of technical change in maize. Robust standard errors are reported in brackets. Significance levels: $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

Table E.10: Effect of agricultural technical change on wage bill shares in innovation and managerial occupations: H industry

	Stay in Innovation		Move to Managerial	
	Stay in H	Move to other	Stay in H	Move to other
	(1)	(2)	(3)	(4)
A_{soy}	0.009	0.015	-0.002	0.024
	[0.014]	[0.013]	[0.017]	[0.016]
Observations	2,862	2,862	2,862	2,862
R-squared	0.683	0.648	0.605	0.623
Microregion FE	Yes	Yes	Yes	Yes
Region x Year FE	Yes	Yes	Yes	Yes
Rural x Year FE	Yes	Yes	Yes	Yes
Other Controls x Year FE	Yes	Yes	Yes	Yes

Notes: The dependent variables are the wage bill share of skilled innovation workers who stay in innovation occupations and stay in high-R&D (H) industries (column 1), the wage bill share of skilled innovation workers who stay in innovation occupations but move to other industries or outside manufacturing (column 2), the wage bill share of skilled innovation workers who move to managerial occupations and stay in H industries (column 3), and the wage bill share of skilled innovation workers who move to managerial occupations and move to other industries or outside manufacturing (column 4), all measured using the Yearly Social Security Data (RAIS). Each share is computed as the wage bill of the indicated worker category divided by the total wage bill of all H-industry skilled innovation workers in that microregion and year. The sample includes workers identified as innovators in 2000, 2001, or 2002. All regressions include microregion fixed effects, region × year fixed effects, and the share of rural population interacted with year fixed effects. "Other Controls × Year FE" include income per capita (in logs), population density (in logs), and literacy rate, all observed in the 1991 Population Census and interacted with year fixed effects. All regressions control for a measure of technical change in maize. Standard errors clustered at the microregion level are reported in brackets. Significance levels: $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.