

Labor Supply and Directed Technical Change: Evidence from the Termination of the Bracero Program in 1964*

Shmuel San

August 31, 2021

Abstract

This paper studies the impact of labor supply on the creation of new technology, exploiting a large exogenous shock to the US agricultural labor supply caused by the termination of the Bracero agreements between the US and Mexico at the end of 1964. Using a text-search algorithm allocating patents to crops, I show a negative labor-supply shock induced a sharp increase in innovation in technologies related to more affected crops. The effect is stronger for technology related to labor-intensive production tasks. Farm-value dynamics indicate that, despite the positive technology reaction, the policy change was undesirable for farm owners.

Keywords— Directed Technical Change, Labor Supply, Induced Innovation, Automation, Immigration Restrictions, Bracero

JEL Classifications— J08, F22, O31, O33

*Department of Economics, The Hebrew University of Jerusalem, Jerusalem 91905, Israel. Email: muly.san@mail.huji.ac.il. For their helpful comments, I thank Michael Clemens, James Fenske, Christopher Flinn, Raphael Franck, Walker Hanlon, Richard Hornbeck, Alfred Galichon, Ethan Lewis, Petra Moser, Suresh Naidu, Tzachi Raz, Martin Rotemberg, Yannay Spitzer, Kevin Thom, Sharon Traiberman, Miguel Uribe, and Joseph Zeira as well as participants at the EHS 2021 meeting, EHA Israel 2020 meeting, NBER 2020 Summer Institute (DAE poster session), SOLE 2019 meeting, AASLE 2018 meeting, Empirics and Methods in Economics Conference (Northwestern) 2018, Warwick Ph.D. Conference 2018, and seminars at NYU and the Hebrew University of Jerusalem. Yoav Sivor provided excellent research assistance. I thank the Economic History Association and the Institute for Humane Studies for financial support.

1 INTRODUCTION

Whether labor abundance/scarcity encourages or discourages technical advance is one of the oldest debates in economics (Malthus 1959; Ricardo 1951). Intuitively, when a factor such as labor becomes more expensive, it spurs invention directed to economizing the use of that factor (Hicks 1932; Zeira 1998). On the other hand, a low number of workers reduces the number of potential users of new technologies (Kremer 1993; Acemoglu 1998). Acemoglu (2010) theoretically shows that the direction of the effect depends on whether technical advance reduces or increases the marginal product of labor. Whether labor scarcity may induce technical advance in practice, however, is an open empirical question (Acemoglu 2010, p. 1071).

Technical advance includes the creation and adoption of new technologies. A few studies test the effect of labor supply on the adoption of technology (Lewis 2011; Hornbeck and Naidu 2014; Clemens et al. 2018). However, they do not address the creation of new technologies.

In this paper, I offer an empirical test for the effect of labor supply on the creation of new technologies. To do so, I utilize a large exogenous shock to the labor supply in the US agricultural sector caused by the termination of the Bracero agreements between the United States and Mexico in 1964. The Bracero agreements were a set of three bilateral agreements between the United States and Mexico to regulate bilateral flows of temporary low-skill labor, spanning 1942–1964 (Clemens et al. 2018). Varying substantially between crops, Bracero workers accounted for about 11% of the total seasonal farm workforce in 1964. The exclusion of those workers from the labor force generated a sharp decline in the labor supply in a very short period.

The first objective of this paper is to document the pattern of the creation of new technologies caused by this shock to the labor supply. Using a text-search algorithm to allocate patents to crops, I show that the Bracero exclusion induced a sharp increase in innovation in technologies related to crops with a higher share of Bracero workers relative to crops with a lower share. Innovators reacted fast, introducing new technologies right after the termination of the program. Innovation in technologies related to high-exposed crops remained high more than 15 years after the end of the program. Thus, the patent data reveal substantial directed technical change towards technologies related to crops with labor scarcity.

To further ensure the robustness of the results, I instrument the share of Bracero workers by the average distance from Mexico and the average historical Mexican population in the counties producing each crop. The IV strategy yields quantitatively similar results.

An alternative robustness check employs patent data to measure the technological similarity between crops. Utilizing the latter, I calculate the synthetic exposure of the crops to the Bracero program, predicted by the exposure of technologically similar crops, and show that the actual exposure to the Bracero program is not correlated with the synthetic exposure measure. Furthermore, I re-estimated the difference-in-differences regressions controlling for the synthetic exposure, isolating the part in the exposure to the labor-supply shock that is not predicted by the technical features of the crops. Once again, the estimated effects are similar.

The second objective of this paper is to study the heterogeneous effect of labor supply on different types of technology. Using detailed data on labor requirements by production task and crop, I compare the impact of a labor-supply shock on the creation of technologies related to more labor-intensive tasks relative to less labor-intensive tasks. Triple difference estimates confirm that the effect is stronger in technologies related to more labor-intensive production tasks. Assuming that such technologies tend to be more labor-saving, these results suggest that labor scarcity encourages the creation of labor-saving technology more than labor-augmenting technology.

In the last part of the paper, I study the impact of the termination of the Bracero program on the profitability of farm businesses. Using information from the US agricultural census, I show that more exposed counties experienced a greater decline in farm values after the shock relative to less-exposed counties. These results, however, are valid only for states that participated in the Bracero program. Taken together, these results show that a negative shock to the supply of low-skilled labor, implied by immigration restrictions, is harmful to farm owners even in the medium and long run, despite the positive technology reaction.

As mentioned earlier, the effect of labor scarcity on the invention of new technology is theoretically ambiguous. To illustrate the two opposite forces, section 2 summarizes the main results of two simple models, with different assumptions about the exact way technology is shaping production.¹ These models have contradicting predictions about the effect of labor supply on innovation activity. If technology increases the quantity of production for every level of labor, as the standard macroeconomics literature suggests, then an increase in labor supply encourages technological progress. On the other hand, if new technology replaces workers, an increase in labor supply discourages technological progress. Taking together, the sign of the effect is an open theoretical question that must be answered empirically. This theoretical framework also motivates the examination of the heterogeneous impact of labor scarcity on the creation of labor-saving technologies and labor-augmenting technologies.

In recent years, there has been an increasing interest in the joint dynamics of artificial intelligence (AI) technology and the labor market (Aghion et al. 2017; Acemoglu and Restrepo 2018). Whether an increase in the available labor supply encourages or discourages technological progress is crucial for the rate of development of AI. If greater labor supply discourages the development of automated technologies, as suggested by the results of this paper, an initial positive shock to AI technologies will increase the available supply of labor (or reduce wages) and hence discourages the development of further AI technology. In other words, the discouraging effect of labor abundance on the creation of labor-saving technologies limits the long-run growth rate of AI technology on the one hand, and unemployment due to automation on the other hand (Nakamura and Zeira 2018; Acemoglu and Restrepo 2018).

The effect of labor supply on technological progress is a fundamental question in economic history. The famous Habakkuk hypothesis claims that US labor scarcity in the 19th century induced

¹ The models themselves are presented in Appendix B.

rapid technological progress relative to Britain (Habakkuk 1962). Similarly, Allen (2009) claims that high wages in 18th century Britain were a preponderant reason for the Industrial Revolution occurring there as opposed to elsewhere.² This paper contributes to the economic history literature by providing causal evidence for the impact of labor supply on the creation of new technology.

The claim that the termination of the Bracero program increased the pace of labor-saving technological innovation is not new. For example, Runsten and LeVein (1981) argued that

"For many years, California agriculture has relied upon abundant supplies of cheap foreign labor, coming mainly from Mexico. As the rural labor market maintained segmented from the rest of the economy, this allowed the mechanization of these specialty crops to be postponed. In 1964, when the use of Mexican labor became constrained by the end of the Bracero Program, a strong inducement was given to introduce mechanical harvesting techniques."

This claim, however, has never been rigorously tested.

Clemens et al. (2018) exploit the termination of the Bracero program to study the effect of labor scarcity on the labor market. They use state-level variation in the exposure to the Bracero program to show that the program's termination did not affect local wages or employment. They also provide supporting evidence for the positive effect of labor scarcity on the adoption of already-existing technologies. The current paper complements their findings by offering an explicit mechanism for their results: the positive innovation response to labor scarcity can dampen the wage response.³

This article contributes to the literature on the effect of factor supplies on technological progress. Newell et al. (1999) and Popp (2002) demonstrated that increased energy prices re-direct innovation to more energy-efficient technology. Hanlon (2015) found that the scarcity of US cotton exported to England during the US Civil War induced the development of new technologies that augmented Indian cotton. Closer to the current topic, Lewis (2011) and Hornbeck and Naidu (2014) have shown that areas with a lower relative supply of low-skilled labor adopted more advanced technology.⁴

The results of this paper are also relevant for understanding the impact of immigration on technological change. Most of the literature has focused on high-skilled immigration, which affects the supply side of innovation (Hunt and Gauthier-Loiselle 2010; Kerr and Lincoln 2010; Borjas and Doran 2012; Moser et al. 2014; Moser and San 2020). In contrast, this paper studies the technological response to limits on low-skilled foreign labor.

Two recent working papers by Doran and Yoon (2019) and Andersson et al. (2019) study the effects of mass migration waves on technological innovation using geographical variation in receiving and sending communities, respectively. The current research complements those papers in two ways. First, my identification strategy takes advantage of an exogenous policy shock that affected only

² See also Hayami and Ruttan (1970) and Alesina et al. (2018) for similar arguments.

³ Methodologically, unlike Clemens et al. (2018) that use state-level variation in the exposure to the Bracero program, I use crop-level variation that better captures changes in innovation.

⁴ See also Lafortune et al. (2015), Lew and Cater (2018), and Abramitzky et al. (2019) for similar results.

low-skilled workers. Moreover, it does so in the context of a guest worker program that enables to isolate the impact of a pure labor-supply shock from other changes usually accompanied by conventional immigration, coming from the role of the immigrants as citizens, consumers, and potential inventors.⁵ Second, I use patents issued in the United States, the technological leader of the time. Thus, these patents are more likely to reflect frontier inventions. This is aided by the fact that my variation is at the crop level, as opposed to spatial variation. Given that technological invention is to a large extent globally applicable, the variation in this paper is better able to capture groundbreaking inventions, as opposed to local adjustments of already-existing technologies.

Finally, this paper also contributes to the literature on the impact of immigration on firms' outcomes. Several papers have found a positive impact of increasing high-skilled immigration on the productivity, size, and profits of firms (Kerr and Lincoln 2010; Ghosh et al. 2014; Doran et al. 2014; Peri et al. 2015; Beerli et al. 2021). Brunello et al. (2020) find that low-skilled immigration also has a small but positive effect on firms' profits. The current paper strengthens the latter results by suggesting that, despite the positive technology reaction to low-skilled immigration declining, the farmers' profits (embedded at the farmland values) go down.

2 THEORETICAL FRAMEWORK

Theoretically, the impact of labor supply on technological progress depends on how one assumes technology is shaping production. In Appendix B, I build two simple models to illustrate this point. I show that if technology increases the quantity of output for every level of labor, as the standard macroeconomics literature suggests, then an increase in labor supply encourages technological progress. On the other hand, if new technology replaces workers, an increase in labor supply discourages technological progress. Taking together, the sign of the effect is an open theoretical question that must be answered empirically.

Importantly, how technology is introduced in each model offers a conceptual framework for distinguishing between different types of technological improvements. For each technological invention, one should consider whether it is more similar to one type of technology or the other. Consider, for example, the invention of the tomato harvester. For a given amount of workers (and land), it does not change the production quantity. However, it can economize the workers needed to produce the same amount of output; therefore, it better aligns with the labor-saving technology type. On the other hand, a new fertilizer, for instance, increases the production for a given amount of land and labor; hence, it better aligns with the labor-augmenting technology type. Using that classification of technologies, the theory suggests different effects for a change in labor supply on the two types of inventions. Section 6 further explores this.

⁵ See Moser and San (2020) for the effect of the 1920s quota acts on high-skilled scientists and inventors.

3 HISTORICAL BACKGROUND AND DATA

3.1 THE BRACERO PROGRAM

Existing from 1942 until 1964, the Bracero program allowed over four million Mexican agricultural workers to migrate legally, making it the most extensive guest worker program in the history of the United States (Kosack 2016).

The wartime Bracero program started on August 4, 1942, when the US government concluded with Mexico an agreement to use Mexican agricultural labor on US farms. From 1942 to 1947, more than 200,000 agricultural workers entered the United States from Mexico. The program's post-war era began in 1948 when Braceros contracted directly with US employers. Approximately 200,000 Mexican legal workers entered the United States between 1948 and 1950. In August 1951, Congress approved Public Law 78, which served as the statutory basis for Bracero contracting until it expired in December 1964. By June 1952, the Bracero system became a permanent component of US farm labor. During the period 1952-1959, on average, 335,000 Mexican workers were annually employed on US farms (Craig 1971).

Opposition to the Bracero program solidified in the early 1960s when more interest groups had joined the fight against imported Mexican labor. In March 1962, the US government required farmers to offer Braceros at least the statewide average wage that in some Bracero-using regions was considerably higher than the area wage. The program finally terminated at the end of 1964. The principal policy goal of excluding Braceros was to improve labor-market conditions for US farm-workers by reducing the size of the workforce (Craig 1971; Clemens et al. 2018).⁶

3.2 DATA

The measure of the exposure to the shock in the primary analysis (section 5) is the share of foreign seasonal labor just before the termination of the Bracero program. The outcome variable is the innovation activity by year and crop, measured by the number of agricultural patents for each year and crop, possibly scaled by forward citations. In section 6, I utilize additional information on labor requirements by crop and technological type to examine the differential effects across different types of technology. Finally, in section 7, I use data on farm values from the census of agriculture to examine the impact of the program on the farmers. Below, I describe the exposure and outcome measures used in the primary analysis. Appendix C provides additional details on these data, as well as the other data sources that I use in the paper.

EXPOSURE TO THE BRACERO PROGRAM: SHARE OF FOREIGN SEASONAL LABOR

Exact data on Bracero workers by the crop is unavailable. However, during the period 1948-1964, 94.5% of the foreign workers admitted for temporary employment in US agriculture were

⁶ See the timeline of the main events in Table A1.

Mexican (Bureau of Employment Security 1966, Tables 1 and 3). Hence, in this paper, I use the share of foreign seasonal workers in the total seasonal employment for each crop in 1964 as a proxy for the share of Bracero workers (Bureau of Employment Security 1966, Table 5).⁷ For most of the statistical analysis of this paper, I use the total number of person-hours worked annually by foreign and local seasonal workers.⁸ The sample for the primary analysis consists of 16 crops that used 4,000 or more person-months of foreign labor in 1964. This measure has a significant variation in the data, ranging from 55% in lettuce to 2% in tobacco (Table 1). For a robustness check, I also use a binary version of the exposure measure, where crops above the median of foreign share are defined as exposed crops (Table A5).

[Table 1 about here]

INNOVATION MEASURE: ALLOCATING PATENTS TO CROPS

My first measure of technology innovation is the number of USPTO patents by crop and year. To account for the quality of the innovation, I also use the number of patents weighted by the number of their forward citations. Because most of the Braceros were allocated to harvesting tasks, I focus on technological innovations related to harvesting and mowing (CPC class A01D) in the primary analysis in section 5. In section 6, I compared patenting in this class with patenting in other technological classes using information on the labor requirements by task and crop (see Table 5).

I allocated patents to crops by searching the text of patents for the crop names. I collected the full text of the patent from Google Patents and looked for the crop names in the title, abstract, claims, and description sections of the patent document. If more than one crop appears in the text of one patent, I allocate the patent to the crop which appears first.⁹ I determined the invention's date by the application date of the patent.¹⁰

Between 1948 and 1985, US inventors were awarded 2,563 patents related to harvesting and mowing, which mentioned at least one of the 16 crops (Table 1). I also collected data from the Google Patents website on the number of citations for each patent.

A basic prediction of the endogenous growth literature is that innovation activity is increasing with scale; the larger is the value of a market, the higher are the incentives to invent new technologies

⁷ This measure is highly correlated with another potential measure of the share of Bracero workers, which is the 1964 to 1965 change in the share of foreign workers (the Pearson correlation coefficient is 0.72).

⁸ For robustness, I also use the share of foreign workers in the total seasonal employment at the date of peak foreign employment from (Bureau of Employment Security 1966, Table 21).

⁹ I estimate robustness checks using other procedures, such as assigning the patents to all the crops or splitting its weight between the crops (see Table A6).

¹⁰ The application date is missing for 34 patents in the sample, for which I estimated the application date by subtracting the median lag between the application date and issue date in the sample (2.6 years) from the issue date.

relevant to that market (Romer 1990; Aghion and Howitt 1992). I use this prediction to validate the text search algorithm that assigns inventions (measured by patents) to the different markets (here, crops). Indeed, the data show a strong positive correlation between the average number of patents by crop and the average value of production by a crop for the period 1948-1985 (Figure 1).¹¹

[Figure 1 about here]

4 AGRICULTURAL INNOVATION AND PATENTING IN THE US

Before turning to the primary analysis, this section provides descriptive information on innovation and patenting in the US. I focus on patents in the primary sample, namely patents of harvesting and mowing technologies related to one of the 16 crops described above, and compare them to the entire set of USPTO patents in the same year.

I start by examining the identity of the inventors. In particular, I check whether the inventors are individual inventors that assign the patents to themselves or firms and other institutions that employ the inventors. Figure A1 shows a monotonic decline in the share of patents that belong to individual inventors for the patents in our sample over the years, from more than 60% in 1948 to less than 40% in 1985. A similar decline in the share of patents by individual inventors occurred to other US patents, from 40% in 1948 to less than 20% in 1985. However, the levels of patents by individuals are consistently higher for the sample of agricultural innovation than other innovations. Finally, the data show no differential pattern between patents issued before and after the end of the Bracero program for both patents inside and outside the sample.

Next, Figure A2 shows the average number of inventors per patent by year. In this case, the data show a monotonic *increase* for both patents in the sample and other patents. At the beginning of the analyzed period, the average number of inventors per patent was around 1.2 for both groups. At the end of the period, the number of inventors per patent is 1.6 for patents in the sample and 2.0 for other patents. Once again, there is no differential pattern between patents issued before or after 1964 in both groups.

Another interesting question is how patent applications in the sample are expected to be reviewed by the US Patent Office. In particular, I check the time between the application and publication of the successful patents. Figure A3 shows that patents in the sample are reviewed faster than other patents. Over 1948-1985, the average time between application and publication was 2.6 years for patents in the sample, compared to 3.0 for other patents. Over the years, the pattern of this variable seems to be similar for patents inside and outside the sample, and without any significant change in 1964.

¹¹ The value of each crop was collected from various publications of the US Department of Agriculture. See Appendix C for additional details and see Table A2 for additional descriptive statistics of the crops in the sample.

Finally, I check some institutional features of the market for agricultural innovation studied in this article compared to the overall market for innovation. To do so, I calculated the number of unique assignees, patents per assignees, as well as other concentration (or market power) measures separately for each group of patents (sample/other patents) and years (1948-1964/1965-1985). Table A3 reports the results.

The number of assignees for patents in the sample is 637 and 800 before and after 1964. The growth rate in the number of assignees for other patents is higher (from 210,531 to 317,974). The number of patents per assignee in the sample is 1.7 for both periods. This figure is much higher in the broader set of patents, with 3.7 and 4.4 patents per assignee in 1948-1964 and 1965-1985, respectively.

To further examine the concentration in the market for invention, I calculate the Herfindahl-Hirschman Index, which is the sum of squares of the share percentage of patents issued by each assignee (ranges between 0 and 10,000). This measure shows a decrease in the market concentration between the two periods for both patents inside and outside the sample. The share of patents owned by the top 1, 3, 5, 10, 30, and 50 assignees also decreased between the two periods for both groups of patents. However, the share of patents owned by top 1, 3, 5, 10, 30, and 50 *percent* of assignees slightly decreased between the periods for patents in the sample, but *increased* for patents outside the sample.

Overall, the information presented in this section suggests no apparent difference in the type of inventors or the way the US patent office reviewed the patents before and after the termination of the bracero program. The patterns of market concentration are less clear. Some of the measures suggest no change (patents assignee) or even a decrease (share of patents by top *percent* of assignees) in the concentration between the periods for patents in the sample, compared to an increase in those measures for other patents. However, other measures (Herfindahl-Hirschman Index and share of patents by top *number* of assignees) suggest a decrease in market concentration for both groups.

5 EFFECTS OF LABOR SCARCITY ON INVENTION IN THE UNITED STATES

My empirical strategy compares changes in invention across crops that were differentially affected by the termination of the Bracero program. Figure 2 illustrates the main results. The relative number of patents for crops with low exposure to the Bracero program reveals almost no change before and after 1965.¹² However, for crops in the medium and high exposure groups, there is a noticeable jump around the end of the Bracero program. The rest of this section explores this finding more rigorously.

¹² To account for the scale differences between crops, I calculated the annual number of patents by crop and year relative to the pre-period (1948-1964) average of that crop. See Figure A4 for a similar graph by crop.

[Figure 2 about here]

The dependent variables, citation-weighted or unweighted patent counts by crop and year, are skewed and nonnegative. For example, 26.6% of the crop/year observations in the data correspond to years of no patent output; the figure climbs to 76.1% if one focuses on crop/year observations with no more than five patents.

To address this count nature of the data, I estimate the model using the Poisson Quasi Maximum-likelihood Estimator, first suggested by Hausman et al. (1984). This estimator is fully robust to distributional misspecification, and it also maintains certain efficiency properties even when the distribution is not Poisson (Wooldridge 1999, 2010).¹³ I compute QML "robust" standard errors, which are consistent even if the underlying data-generating process is not Poisson. These standard errors are robust to arbitrary patterns of serial correlation (Wooldridge 1997; Bertanha and Moser 2016).¹⁴

The years of analysis for most of the specifications are 1948-1985.¹⁵ I chose 1948 for two main reasons. First, it is the first year that the Bracero workers were employed directly by the farmers and not by the US government. Second, choosing a year in the middle of the period avoids an additional (positive) shock to the labor supply at the program's beginning.¹⁶

5.1 BASELINE SPECIFICATION

My estimating equation relates crop i 's output in year t to characteristics of i :

$$\ln [\mathbb{E}(Innovation_{it}|X_{it})] = \beta \cdot ForeignShare_i \cdot post_t + \gamma_i + \delta_t \quad (1)$$

where $Innovation_{it}$ is a measure of innovation output at crop i at year t , $ForeignShare_i$ is the share of foreign workers in the total number of seasonal workers in crop i one year before the termination of the Bracero program, $post_t$ denotes an indicator variable that switches to one after 1965, the γ_i 's correspond to crop fixed effects, the δ_t 's stand for a full set of calendar year indicator variables, and X_{it} denotes all the independent variables on the right-hand side of the equation.

Table 2 presents the main results. Column (1) examines the determinants of the 16 crops' patent count. I find a significant increase in the yearly number of patents produced after 1965 in crops

¹³ Except for the conditional mean, the distribution of the outcome variable given the dependent variables and the coefficients is entirely unrestricted. In particular, there can be overdispersion or underdispersion in the latent variable model (Wooldridge 1997).

¹⁴ Due to the small number of crops (16 in most specifications), I did not cluster the standard errors in the crop-level regressions. However, results are robust for clustering the standard errors at the crop level (Table A4).

¹⁵ The results are robust to the choice of the start and end years (Table A8).

¹⁶ Unfortunately, I cannot estimate the effect of this positive shock due to the lack of data on the share of Mexican workers by crop in the first years of the Bracero program.

that were more exposed to the Bracero program. A one percentage point rise in the share of foreign workers before the policy change increases the innovation activity by 3.3 percent (significant at one percent). Compared with an average of 4.06 annual patents per crop in 1948-1964, an increase of one standard deviation in the labor-supply shock increases the number of patents by 70.7%, which amounts to 2.87 additional annual patents per crop at the average pre-period level.¹⁷

[Table 2 about here]

Column (2) provides the results for citation-weighted patents, a measure that takes into account the quality of the innovation. The effect is somewhat smaller: A one percentage point rise in the share of foreign workers increases the quality-adjusted innovation measure by 2.3 percent.

A potential challenge to the difference-in-differences estimation is that pre-treatment trends may drive the difference between the patenting of crops with different degrees of exposure. To address this concern and to check the persistence of the effect, I explored the dynamics of the effects uncovered in Table 2 by estimating a specification in which the treatment effect interacts with a set of indicator variables corresponding to a particular calendar year and then graphing the effects and the 95% confidence interval around them.¹⁸

Following the end of the Bracero program, the treatment effect rises monotonically, peaking three to four years after Bracero exclusion, and remaining at the same level (see Figure 3). Two aspects of this result are noteworthy. First, the fact that R&D responded so quickly to the negative labor-supply shock suggests that part of the new patented technologies were ready, or at least close to being ready, at that time. Potentially, the expected labor shortage provided only a "nudge" to the inventors of these technologies to make them operational or even just to issue a patent for them. However, the data show no evidence of recovery—the effect of Bracero exclusion persisted for at least 15-20 years. This result suggests that labor scarcity not only, or mainly, induced the patenting of (almost) existing technologies but mostly induced the invention of new technologies. Second, the event study coefficients fluctuate around zero and are not significantly different from zero for periods before 1965, showing no evidence for a pre-treatment trend.

[Figure 3 about here]

5.2 ROBUSTNESS CHECKS

Appendix A provides additional evidence testing the robustness of the results. The first set of robustness checks evaluate the sensitivity of the results for the definition of the treatment. My

¹⁷ The percentage change in innovation for an increase of a standard deviation σ in the share of foreign workers is $\exp(\beta \cdot \sigma) - 1$. For $\beta = 3.258$ and $\sigma = 0.164$, the percentage change is $\exp(3.258 \cdot 0.164) - 1 = 0.707$.

¹⁸ The small size of the sample (16 crops) does not allow the estimation of many coefficients simultaneously with satisfactory precision. Therefore, I estimate biennial coefficients. A similar picture was obtained when estimating annual coefficients but with larger confidence intervals.

preferred exposure definition is a continuous variable measuring the exposure to the shock in 1964, one year before the end of the program. This measure carries more information than a binary treatment variable, which is more common in difference-in-differences studies. The results, however, are robust to the use of a dummy variable, where crops above the median of foreign share in 1964 get the value one. I estimate the model separately for the two versions of the outcome variable, the number of patents and citations-weighted patents. The Poisson estimate of the effect is an increase of 92.5 log points (significant at 1 percent). This estimate suggests that, after 1964, American inventors produced 152 percent additional patents in high exposed crops relative to low-exposed crops. The estimated effect using the quality-adjusted measure of invention is a 60.3 log points (or 83 percent) increase in invention (Table A5, columns 3-4).

The primary measure of exposure in this paper is the share of foreign seasonal workers in the total number of person-hours annually worked in 1964. However, the results are robust to using the share of foreign workers in the total seasonal employment at the date of peak foreign employment. The Poisson estimators using this measure of exposure to the Bracero program are somewhat smaller but still significant at one percent (Table A5, columns 5-6).

The process of Bracero exclusion began in 1962 when the US government raised the required wage rate for Bracero workers and was completed at the end of 1964 (Craig 1971). In the specifications above, I picked the post-year to be 1965. Defining the post-year to be 1962, however, the results are virtually unchanged (Table A5, columns 7-8).

In the baseline specification, I use the share of foreign seasonal workers in 1964 as a proxy for workers under the Bracero program. Another (probably tighter) measure of the share of Bracero workers is the 1964 to 1965 change in the share of foreign workers. Using this measure to define the exposure to the Bracero program, the estimates of the effects are somewhat smaller, but still statistically significant at least at five percent (Table A5, columns 9-10).

I also checked the sensitivity of the results to the algorithm assigning patents to crops. I compared five different alternatives. 1) The baseline algorithm assigns the patent to the first crop that appears in the text of the patent. The four alternative algorithms are: 2) assigning the patent to the crop that was mentioned more times than any other crop, 3) assigning a patent to each crop mentioned in the text, 4) assigning equal weight to each crop mentioned so that the sum of the weights is one, and 5) assigning weights proportional to the number of times each crop is mentioned. All algorithms yield similar results, both for the innovation measure based on patent counts and the innovation measure based on the number of citations (Table A6).

The next set of robustness checks alternate the crops included in the analysis. Restricted by the data availability, the baseline sample of this paper contains crops that used 4,000 or more person-months of foreign labor in 1964. This choice implies that these crops tend to be labor-intensive. Moreover, prominent crops (e.g., wheat, corn) are not part of the original sample. To examine the validity of the results for a broader range of crops, I extended the sample by the ten field crops

with the largest acreage in the 1964 agricultural census.¹⁹ Unfortunately, I did not find exact data on the share of foreign workers in crops with less than 4,000 person-months of foreign labor in 1964, including these field crops. However, the foreign share of the category "Hay and Grain" is 1.2 percent. In what follows, I assumed that each of the ten field crops has the common foreign share of 1.2 percent. Columns 3-4 of Table A7 show the results of the difference-in-difference specification for the aforementioned extended sample. The effect of labor scarcity on innovation is positive and significant for both innovation measures. The magnitude of the effect is comparable to the original sample of 16 crops, although a bit smaller. The effect of a one percentage point rise in the share of foreign workers is 2.8 percent for a simple patents-count innovation measure and 1.5 percent for the quality-adjusted measure.

I also extended the sample to include ten additional crops of which information on the share of foreign workers in 1962 in California is available.²⁰ Columns 5-6 of Table A7 report the results for a sample containing the 16 original crops and the ten "California crops". The results are virtually identical to the baseline results, with estimates of a 3.1 and 2.3 percent increase in the number of patents and citations, respectively. Finally, columns 7-8 of Table A7 show the results for all crops together. The estimated effects are 2.8 and 1.5 percent (significant at one percent).

The years of analysis in the baseline specification are 1948-1985. To check the sensitivity of the results for this choice, which is somewhat arbitrary, I estimated the baseline specification for different periods. Table A8 indicates that the results are not sensitive to that decision.

Although the preferred statistical model for count data is the Poisson model, I checked the sensitivity of the results for three alternative models. Columns 3-4 of Table A9 report the results of a Negative Binomial model. The estimators are positive with a similar magnitude (effect of 2.2 and 2.0 percent for patent and citation counts, respectively, significant at 1 percent). Next, I estimated a zero-inflated Poisson regression. This model assumes that two different processes generated the outcome variable. The first process is governed by a binary distribution that generates extra zeros. If the first process yields zero, the outcome is simply zero. However, if the binary process yields one, the outcome is sampled from a Poisson distribution. I assume that the excess zero counts (the first process) come from a logit model. Maximum likelihood estimates of this model yield results that are very similar to the baseline Poisson model (2.9 and 1.8 percent increase, and significant at one percent, Table A9, columns 5-6). Finally, I estimated an OLS model, where the outcome variable is the natural log of the count of patents and citations (observations with zero patents/citations are dropped from the regression). An estimate of the effect using the patents measure shows an effect of a 1.5 percent increase, smaller from the baseline Poisson estimate (significant at 1 percent). Using the quality-adjusted inventions measure, I find that the estimated effect is an increase of 2.0 percent, similar to the baseline estimate (significant at one percent, Table A9, columns 7-8).

As I discussed earlier, the data show no evidence for pre-treatment trends. To further address

¹⁹ These crops are: barley, corn, flax-seed, oats, peanuts, rice, rye, sorghum, soybeans, and wheat.

²⁰ These crops are apricots, cherries, olives, peaches, pears, plums, prunes, lemons, almonds, and walnuts.

this concern, I estimated the baseline specification with crop-specific linear pre-trends. The estimates are greater (4.9 and 4.5 percent for patents and citations, respectively) and statistically significant at one percent (Table A10, columns 3-4).

5.3 THE DECISION ABOUT THE BRACERO WORKERS AND INSTRUMENTAL VARIABLES ESTIMATION

What explains the variation in the share of Mexican workers between the crops? My identifying assumption is that controlling for crop and year fixed effects, changes in patenting would have been comparable for crops with a high and low share of Mexican workers if the US government had not terminated the Bracero program. This assumption is violated if the percentage of Mexican workers is correlated with factors that generate unparalleled innovation activity trends for unrelated reasons. For example, suppose the share of Mexican workers is higher in crops with higher labor requirements per acre, and there is convergence in the invention dynamics such that crops with higher labor requirements close the gap by having more labor-saving inventions in later years. In that case, the estimated effect is not the causal effect of the Bracero exclusion.

Using data on the value of production, seasonal labor, and acreage of the crops, Table A11 reports the correlation between the foreign share of seasonal labor in 1964 and various measures of labor productivity, yield, and market size. The data suggest no significant correlation between the foreign share and any of these measures (p-value is always greater than 10 percent). While the data show no correlation between the exposure measure and any of the observable characteristics of the crops, the rest of this section uses instrumental variables to address the possibility that other (unobservable) characteristics might violate the parallel trend assumption.

Two logical instruments are the distance from Mexico and the historical share of the population of Mexican origin. The data show that, other things equal, seasonal Mexican workers tended to work in places closer to the US-Mexico border and in places that attracted older waves of immigration. The exclusion restriction requires that those variables must not affect the technological progress differentially in the pre and post periods, except for its impact through the channel of the Bracero program termination. Indeed, it is unlikely that the proximity to Mexico or the historical share of Mexicans impacts the technological progress differentially before and after 1965 other than its effect through the Bracero program.²¹

To construct the instrument, I use county-level information on the distance from Mexico, and the share of the Mexican population in 1940, taken from the US census of population. As my innovation measures are at the crop level, I need to transform the instruments from the spatial dimension into the crop dimension. To do so, I use the US agricultural census from 1964 for information on the

²¹ The exclusion restriction does not require that the cross-sectional variation of the instruments would not affect the technological progress itself. The instruments remain valid even if the cross-sectional variation is related to these instruments. For instance, if the crops closer to Mexico tend to have faster technological progress, this would be picked up by the crop fixed effect, and the exclusion restriction would still hold.

crops produced in each county. More precisely, the average distance from Mexico of a crop i is measured by $d_i = \sum_c d_c w_{ic}$ where d_c is the minimal distance between the Mexican border and the center of the county c , and w_{ic} is the percent of the acreage of crop i in county c in the total acreage of crop i in 1964. The crop-average Mexican population is calculated similarly.

To implement the IV for count-data, I use a model first introduced by Mullahy (1997). It is widely used in the empirical literature and has better asymptotic properties than the additive errors models.²² The model takes the form:

$$Innovation_{it} = \exp[\beta \cdot ForeignShare_i \cdot post_t + \gamma_i + \delta_t] \cdot \epsilon_{it} \quad (2)$$

where ϵ_{it} is a unit-mean error term. The treatment variable $ForeignShare_i \cdot post_t$ is instrumented by $z_i \cdot post_t$, where z_i is either the average distance from Mexico, or the 1940 average percentage of Mexicans in the population of the counties growing the crops (or both). The GMM estimators of the model are presented in Table 3.²³ In the first and fourth columns, the instrument used is the average distance from Mexico. The estimates for a one percentage point rise in the share of foreign workers are 4.9 and 5.3 percent for the patents and citations measures, respectively (significant at one percent). In the second and fifth columns, I use the average share of the Mexican population similarly. The estimates of the effect are a 3.0 percent increase in patents (marginally significant, p-value = 0.067) and a 4.3 percent increase in citations (significant at five percent). Finally, the third and sixth columns report the estimates where both instruments are used. The estimates are 4.5 and 5.0 percent, respectively, both significant at 1 percent. Overall, the IV estimates of the effect are somewhat higher than the baseline estimates. This indicates that, if anything, the simple Poisson estimates of the effect are biased toward zero.

[Table 3 about here]

5.4 BUILDING SYNTHETIC EXPOSURE USING A TECHNOLOGY-BASED SIMILARITY MATRIX

An additional threat for the identification strategy comes from a potential technical similarity between groups of crops. If exposure to the Bracero program is not randomly distributed across the groups, differential technical progress between the groups might confound the results.

²² See Cameron and Trivedi (2013) for a review on count-data instrumental variables estimation. Similar results were obtained using additive-errors and control-function models.

²³ See Table A12 for the corresponding "first-stage" estimates. The Poisson-GMM estimation method does not rely on an actual first stage but on the moment conditions implied by the exclusion restriction. Anyway, I report the OLS coefficients of a linear first stage. Those are the exact coefficients obtained in the first stage when using the control function method to estimate the IV-Poisson (instead of GMM), which provides very similar results.

To address this concern, I checked the correlation between the exposure to the Bracero program and the technical features of the crops. To do so, I build a synthetic exposure measure which is the predicted exposure according to the exposure of crops that are similar to the original crop regarding technical properties. The synthetic exposure enables me to check whether the technical features of a crop predict its actual exposure to the program.

I measured the technical similarity between crops by the number of patents that mention both crops. If many technological innovations are relevant for two crops simultaneously, those crops have a lot in common regarding technical properties. Specifically, I build a similarity matrix where the off-diagonal entry (i, i') is the number of patents in the sample that mention crops i and i' somewhere in the text, and the diagonal entries are set to zero. Then, each row in the matrix is normalized to sum to one. Table A13 shows the similarity matrix. The results indicate, for example, that citrus is most similar to apples and that asparagus is a combination of celery, lettuce, and tomatoes. Using this similarity matrix, I constructed the synthetic exposure as follows:

$$ForeignShare_i^{syn} = \sum_{i' \neq i} w_{i,i'} ForeignShare_{i'} \quad (3)$$

where $ForeignShare_{i'}$ is the foreign seasonal workers shares of crop i' . The data show no correlation between the actual foreign shares and the synthetic ones (Figure 4). This indicates that the crop's exposure to foreign labor is orthogonal to the technological features of the crops measured by the patent-based similarity measure described above.

[Figure 4 about here]

Furthermore, I re-estimated the Poisson regressions controlling for the synthetic exposure. By doing so, I isolated the part in the exposure to the labor-supply shock that is not predicted by the technical features of the crops. The difference-in-differences specification takes the form:

$$\begin{aligned} \ln [\mathbb{E}(Innovation_{it}|X_{it})] &= \beta \cdot ForeignShare_i \cdot post_t \\ &+ \alpha \cdot ForeignShare_i^{syn} \cdot post_t + \gamma_i + \delta_t \end{aligned} \quad (4)$$

The results for the two invention measures are reported in Table 4. The estimated effect is 3.6 and 2.5 percent for the patents and citations measures, respectively (significant at 1 percent). These estimates are close to the baseline results.

[Table 4 about here]

Overall, the results in this section provide additional support to the claim that the allocation of Mexican workers between crops was not systematically correlated with features of the crops that

affect future technological innovation. Therefore, the labor-supply shocks can be treated as if they are randomly assigned to the crops.

6 EFFECTS BY TYPE OF TECHNOLOGY

The main prediction of the theoretical model presented above is that a negative shock of the labor supply should increase labor-saving technological progress more than labor-augmenting technologies. An ideal way to check this prediction is to identify labor-savings and labor-augmenting technologies from the text of the patent. However, in practice, this task is not easy to perform, the more so through an automatic algorithm. For example, a patent for "Grape Harvester" was granted in 1973 (US patent number 3,766,724). This innovation improves the performance of a mechanical grape harvester that replaces manual laborers; therefore, it should be classified as a labor-saving technology. In the text of the patent, however, none of the words "labor", "work", "job", "employment", "task", "save", or "replace" appear.

To bypass this problem, I use information on the labor intensity of different tasks as a proxy for the probability of a technological innovation related to these tasks to be labor saving. The underlying assumption is, *ceteris paribus*, the incentive to develop new labor-saving technology for a particular task is higher the higher that task's labor intensity is. To conduct this, I used the technological classification of the patents together with data on labor requirements per task and crop.

In particular, I collected data on labor requirements from the State of California's "Report and Recommendations of the Agricultural Labor Commission" (State of California 1963). This data includes information on California's 25 most valuable crops in 1960. For each of these crops, the report lists in detail all the tasks needed to produce the crop, together with estimates on person-hours and labor cost required to produce an acre of that crop. Among the 25 crops included in this data set, 18 have information on the exposure to the Bracero program (either at the country level in 1964 or for the state of California in 1962).

To get a link between agricultural tasks and technology, I manually classified each task into one of the six agricultural patent subclasses.²⁴ To do so, I chose the technology subclass most similar to the task's description using the subclass's detailed definitions.²⁵ For example, the production of tomatoes in 1960 required 12.5 person-hours of "thinning" per acre, at the cost of 13.12 dollars. I classified this task into CPC subclass A01B ("Soil Working In Agriculture"), which contains the group "thinning machines" (A01B 41).²⁶

For each class-crop pair, I calculated the share of labor requirements for this technological class

²⁴ See the definition of these subclasses in Table A14.

²⁵ The definitions of the CPC classification can be found at <https://www.uspto.gov/web/patents/classification/cpc/html/cpc.html>.

²⁶ Table A15 reports the classification of all tasks.

over the total labor requirements of that crop. I used two versions of these labor-intensity measures, one using person-hours and the second using monetary cost. The second measure takes into account potential differences in skills or efficiency units of the labor inputs. Among the six subclasses, only three have a significant percentage of labor: Soil Working (A01B), Harvesting (A01D), and Cultivating (A01G). The average share of person-hours labor inputs for these crops is 15 percent, 50 percent, and 26 percent, respectively (Table A14).²⁷ As a robustness check, I also estimated a specification where the labor-intensity measure equals one for Harvesting, which is the most labor-intensive category on average, and zero for the other two categories. This specification does not require information on the actual labor intensity of each class-crop, thus allowing estimation with all crops which I have data on their exposure to the Bracero program.

Using those measures, I estimated the following continuous triple-difference specification:

$$\ln [\mathbb{E}(\text{Innovation}_{ijt} | X_{ijt})] = \beta \cdot \text{ForeignShare}_i \cdot \text{Intensity}_{ij} \cdot \text{post}_t + \gamma_{ij} + \delta_{it} + \epsilon_{jt} \quad (5)$$

where Innovation_{ijt} is the number of US patents/citations in crop i , technological class j , and year t . ForeignShare_i is the foreign percentage of seasonal workers in crop i in 1964. Intensity_{ij} is a measure of labor inputs required to perform task j in crop i . post_t indicates years after 1964. γ_{ij} , δ_{it} , and ϵ_{jt} are crop-task, crop-year, and task-year fixed effects, respectively.

The Poisson quasi-maximum likelihood estimates of equation 5 imply a substantial higher effect of Bracero exclusion after 1964 in the more labor-intensive tasks relative to the less labor-intensive tasks. This result is true for the three different labor-intensity measures. The first measure is the percentage of person-hours required for tasks in class j over the total person-hours required for producing crop i . Using that measure, the effect of a one percentage point rise in the share of foreign workers on patents after 1965 is higher by 3.2 percent in technological classes required 100 percent of the labor compared with technological classes required no labor. (Table 5, column 1, significant at one percent). The effect is slightly smaller, 2.2 percent, when using the citations measure of innovation (Table 5, column 2, significant at five percent).

[Table 5 about here]

In the theoretical model described above, I assumed workers are homogeneous, and therefore there is only one wage level. In reality, however, some tasks can only be performed by higher-skilled workers and cost more per hour of work. The second measure of labor intensity takes this into account by weighting the hours required to perform a task with the wage rate paid for that task. The measure is the share of labor cost for a class of tasks in the total labor cost of a crop. Using this labor-intensity measure, I find the results are virtually the same as the results of the first measure. The estimates for β are now 3.1 and 2.1 percent, respectively (5, columns 3-4).

²⁷ Results are robust for including all six technology subclasses in the analysis.

The first two labor-intensity measures require exact information about the labor requirements of each task and crop. This data is available only for 18 out of 26 crops with information on the exposure to the Bracero program. The third measure of labor intensity equals one for harvesting tasks, the most labor-intensive class on average, and zero for the other two classes. Using this measure I can estimate equation 5 with all 26 crops.²⁸ The triple-difference estimates are now slightly smaller, 2.5 and 1.8 percent, respectively (Table 5, columns 5-6).

To further investigate the triple-difference results, I estimate difference-in-differences specifications for each technology subclass separately. The results in Figure 5 show that the point estimates are monotonically increasing with the labor intensity of the task. In the Soil Working subclass, which is the least labor-intensive subclass (accounts for 14 percent of the labor requirements per crop on average), a one percentage point rise in the share of foreign workers *decreases* the number of patents by 0.53 percent (the 95 percent confidence interval is [-1.96,0.90]). The estimated effect in the Cultivating subclass (26 percent of the labor requirements) is an increase of 0.87 percent, and the confidence interval is [-0.05,1.79]. Finally, the estimate for the Harvesting subclass (50 percent of the labor requirements) is 3.14, and the confidence interval is [2.26,4.01].

[Figure 5 about here]

Overall, the results presented in this section indicate that the effect of labor scarcity on technological progress is greater in more labor-intensive tasks. Under the assumption that labor-saving technologies are more likely to be developed for labor-intensive tasks, the results suggest that labor scarcity encourages the invention of labor-saving technologies more than other technologies, in accordance with the theory.

Without making this assumption, a more modest interpretation of the results of this section can be offered. One can think of the additional information about the task's labor intensity as a measure of the shock intensity. The extra dimension allows adding of fixed effects for year-crop, year-class, and crop-class pairs. These fixed effects address many potential threats to the baseline difference-in-differences specification, such as time-varying crop-specific demand shocks that, for some reason, are correlated with the Bracero shock. Thus, the results offer an additional robustness check for the effect of labor supply on technical innovation.

7 THE IMPACT ON FARM OWNERS

How the Bracero exclusion impacted the farm owners that employed the Bracero workers? In principle, the profits of farm owners should decline after the Bracero exclusion due to their loss of low-cost labor. Under standard assumptions (e.g., without externalities), this is true even if we

²⁸ These crops include the 16 crops in the main data set and additional ten crops with information on the exposure in the state of California. The results are similar when restricting the sample to the 16 original crops (estimates of 2.7 and 1.7 percent, respectively, significant at five percent).

consider the positive technological reaction, as revealed by the fact that farmers chose to hire the Bracero workers when possible and fought against the program's cancellation. However, if there are sufficiently large externalities such as knowledge spillovers, the policy change may cause a "big push" dynamic that makes farmers better off in the medium and long run by "being forced" to adopt the newly available labor-saving technology (Hornbeck and Naidu 2014).

To check whether farmers won or lost from the policy change, I used the land value to measure the profits of the farm owners. The value of a farm would increase if, following the end of the program, it became more profitable to be a farmer in farms that were more exposed to the program.

In particular, I use the US census of agriculture for the years 1950-1982 to build a panel data of land-value per acre by county and year. Additionally, using the same data sets and the exposure measures by crop, I construct a measure of the exposure of a county c to the Bracero program in the following way:

$$Exposure_c = \sum_i ForeignShare_i \cdot AcreageShare_{ic} \quad (6)$$

where $ForeignShare_i$ is the foreign percentage of seasonal workers in crop i and $AcreageShare_{ic}$ is the share of crop i in the total acreage of county c in the 1964 census. The regression equation is:

$$\ln(Value_{ct}) = \sum_{\tau=1950}^{1982} \beta_{\tau} \cdot \mathbb{I}(t = \tau) \cdot Exposure_c + \gamma_c + \delta_t + \epsilon_{ct} \quad (7)$$

where γ_c and δ_t are county and year fixed effects, respectively. I ran separate regressions for Bracero and non-Bracero states.²⁹ Figure 6 shows a permanent decrease in farm values of counties that are relatively more exposed to the shock. These results are valid only for states that participated in the Bracero program.

[Figure 6 about here]

The results of this section show that despite the positive technology reaction, farmers who employed Bracero workers were adversely affected by the termination of the program, even in the medium and long run. In other words, innovation was not enough to offset the lower labor supply for the affected farms.³⁰ This fact comports with historical documentation about farmers' opposition to the program's termination.

²⁹ Following Clemens et al. (2018), I defined Bracero states as having some Braceros in 1955 and non-Bracero states as having zero Braceros in 1955.

³⁰ Hornbeck and Naidu (2014) also find a negative impact on farmland values per acre after a negative shock to the supply of low-skilled workers. Similar results (although not always statistically significant) also obtained by Lafortune et al. (2015) and Abramitzky et al. (2019).

Moreover, in a recent study, Clemens et al. (2018) show that although the termination of the program aimed to increase the wages and employment rate of local US workers, both employment and wages were not affected. Taken together, one can conclude that Bracero exclusion made capital worse off while making labor no better off.³¹

Finally, the fact that the negative impact starts only after the termination of the Bracero program suggests that the policy change was also unexpected, justifying the identification strategy of this paper.

8 CONCLUSION

This study provides evidence that the supply reduction of seasonal Mexican workers in the United States after the termination of the Bracero program caused the invention of new harvesting machines. I demonstrated that US inventors focused their efforts on developing new technologies that supported the production of crops that were affected by the labor-supply shock. Moreover, I showed that more inventions related to production tasks that required intensive labor input were invented, probably because those technologies tended to be more labor-saving. Finally, I show that the immigration restrictions were harmful to farm businesses, despite the positive technological response.

The termination of the Bracero agreement caused a massive negative shock to agricultural labor supply with high variation between the different crops. This shock provides a rare opportunity to study the effect of labor supply on the creation of new technologies. The fact that this study focused on innovations in the United States, technological leaders of the time, and the use of between crop variation helped capture this type of technological progress.

I developed a new method to classify patents into crops, using the entire text of the patent. While the vast majority of studies use only the count of patents and citations of the patents to measure technology, the patent text provides a new rich world of information that needs to be explored. The current study takes a small step in this direction, but there is much more to be done. One concrete example is the identification of labor-saving innovations. This study attempts to indirectly measure it using the information on the labor-intensiveness of different tasks. However, a direct measure based on the terminology used in the patent could be more effective.

This study focused only on one industry, agriculture. Despite the importance of this industry in economic development and economic history, the direction and magnitude of the effect in different industries would also be of much interest. Moreover, there are reasons to believe that technological progress in agriculture tends to be more labor-saving than in other sectors (Acemoglu 2002, 2010).

³¹ It is behind the scope of this paper to study the impact of the Bracero exclusion on aggregate welfare. The increased innovation could have had off-farm benefits such as higher productivity, lower output prices, or spillovers to other innovations that offset (or even reverse) the harm to the landowning farmers. Nevertheless, the results suggest that the Bracero exclusion caused a Pareto loss to both domestic workers and farmers collectively.

Thus, the finding that labor scarcity encourages innovation in this industry is consistent with the theory. Future research on the heterogeneity of the effect by industry and the factors that can explain this heterogeneity is needed.

REFERENCES

- Abramitzky, R., P. Ager, L. P. Boustan, E. Cohen, and C. W. Hansen (2019). The effects of immigration on the economy: Lessons from the 1920s border closure. *Unpublished*.
- Acemoglu, D. (1998). Why do new technologies complement skills? Directed technical change and wage inequality. *The Quarterly Journal of Economics* 113(4), 1055–1089.
- Acemoglu, D. (2002). Directed technical change. *The Review of Economic Studies* 69(4), 781–809.
- Acemoglu, D. (2010). When does labor scarcity encourage innovation? *Journal of Political Economy* 118(6), 1037–1078.
- Acemoglu, D. and P. Restrepo (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108(6), 1488–1542.
- Aghion, P. and P. Howitt (1992). A Model of Growth Through Creative Destruction. *Econometrica*, 323–351.
- Aghion, P., B. F. Jones, and C. I. Jones (2017). Artificial intelligence and economic growth. *Unpublished*.
- Alesina, A., M. Battisti, and J. Zeira (2018). Technology and labor regulations: theory and evidence. *Journal of Economic Growth* 23(1), 41–78.
- Allen, R. C. (2009). *The British industrial revolution in global perspective*. Cambridge University Press Cambridge.
- Andersson, D., M. Karadja, and E. Prawitz (2019). Mass Migration and Technological Change. *Unpublished*.
- Beerli, A., J. Ruffner, M. Siegenthaler, and G. Peri (2021). The abolition of immigration restrictions and the performance of firms and workers: evidence from Switzerland. *American Economic Review* 111(3), 976–1012.
- Bertanha, M. and P. Moser (2016). Spatial errors in count data regressions. *Journal of Econometric Methods* 5(1), 49–69.
- Borjas, G. J. and K. B. Doran (2012). The collapse of the Soviet Union and the productivity of American mathematicians. *The Quarterly Journal of Economics* 127(3), 1143–1203.

- Brunello, G., E. Lodigiani, and L. Rocco (2020). Does low-skilled immigration increase profits? Evidence from Italian local labour markets. *Regional Science and Urban Economics* 85, 103582. Publisher: Elsevier.
- Bureau of Employment Security (1966, March). Farm labor developments. Technical report, Department of Labor, Manpower Administration, Washington D.C.
- Cameron, A. C. and P. K. Trivedi (2013). *Regression analysis of count data*, Volume 53. Cambridge university press.
- Clemens, M. A., E. G. Lewis, and H. M. Postel (2018). Immigration restrictions as active labor market policy: Evidence from the mexican bracero exclusion. *American Economic Review* 108(6), 1468–87.
- Craig, R. B. (1971). *The Bracero program: Interest groups and foreign policy*. University of Texas Press.
- Doran, K., A. Gelber, and A. Isen (2014). The effects of high-skilled immigration policy on firms: Evidence from H-1B visa lotteries. *Unpublished*.
- Doran, K. and C. Yoon (2019). Immigration and Invention: Evidence from the Quota Acts. *Unpublished*.
- Ghosh, A., A. M. Mayda, and F. Ortega (2014). The impact of skilled foreign workers on firms: an investigation of publicly traded US firms. *Unpublished*. Publisher: IZA Discussion Paper.
- Habakkuk, H. J. (1962). *American and British technology in the nineteenth century: The search for labour-saving inventions*. London: Cambridge University Press.
- Haines, M., P. Fishback, and P. Rhode (2019). United States agricultural data, 1840-2012 (ICPSR 35206).
- Hanlon, W. W. (2015). Necessity is the mother of invention: Input supplies and Directed Technical Change. *Econometrica* 83(1), 67–100.
- Hausman, J., B. H. Hall, and Z. Griliches (1984). Econometric Models for Count Data with an Application to the Patents-R&D Relationship. *Econometrica* 52(4), 909–38.
- Hayami, Y. and V. W. Ruttan (1970). Factor prices and technical change in agricultural development: The United States and Japan, 1880-1960. *Journal of Political Economy* 78(5), 1115–1141.
- Hicks, J. (1932). *The theory of wages*. London: Macmillan.
- Hornbeck, R. and S. Naidu (2014). When the levee breaks: black migration and economic development in the American South. *American Economic Review* 104(3), 963–990.

- Hunt, J. and M. Gauthier-Loiselle (2010). How much does immigration boost innovation? *American Economic Journal: Macroeconomics* 2(2), 31–56.
- Kerr, W. R. and W. F. Lincoln (2010). The supply side of innovation: H-1B visa reforms and US ethnic invention. *Journal of Labor Economics* 28(3), 473–508.
- Kosack, E. (2016). Guest Worker Programs and Human Capital Investment: The Bracero Program in Mexico, 1942-1964. *Unpublished*.
- Kremer, M. (1993). Population growth and technological change: One million BC to 1990. *The Quarterly Journal of Economics* 108(3), 681–716.
- Lafortune, J., J. Tessada, and C. González-Velosa (2015). More hands, more power? Estimating the impact of immigration on output and technology choices using early 20th century US agriculture. *Journal of International Economics* 97(2), 339–358.
- Lew, B. and B. Cater (2018). Farm mechanization on an otherwise ‘featureless’ plain: tractors on the Northern Great Plains and immigration policy of the 1920s. *Cliometrica* 12(2), 181–218. Publisher: Springer.
- Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics* 126(2), 1029–1069.
- Malthus, T. R. (1959). *Population: the first essay*, Volume 31. University of Michigan Press.
- Moser, P. and S. San (2020). Immigration, Science, and Invention: Evidence from the Quota Acts. *Unpublished*.
- Moser, P., A. Voena, and F. Waldinger (2014). German Jewish émigrés and US invention. *American Economic Review* 104(10), 3222–55.
- Mullahy, J. (1997). Instrumental-variable estimation of count data models: Applications to models of cigarette smoking behavior. *Review of Economics and Statistics* 79(4), 586–593.
- Nakamura, H. and J. Zeira (2018). Automation and Unemployment: Help is on the Way. *Unpublished*.
- Newell, R. G., A. B. Jaffe, and R. N. Stavins (1999). The induced innovation hypothesis and energy-saving technological change. *The Quarterly Journal of Economics* 114(3), 941–975.
- of California, U. (1963). *Seasonal labor in California agriculture*. California: Division of Agricultural Sciences, University of California.

- Peri, G., K. Shih, and C. Sparber (2015). STEM workers, H-1B visas, and productivity in US cities. *Journal of Labor Economics* 33(S1), S225–S255. Publisher: University of Chicago Press Chicago, IL.
- Popp, D. (2002). Induced innovation and energy prices. *American Economic Review* 92(1), 160–180.
- Ricardo, D. (1951). *Works and Correspondence*, edited by P. Sraffa. Cambridge University Press, Cambridge.
- Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy* 98(5, Part 2), S71–S102.
- Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyer, J. Pacas, and M. Sobek (2018). Ipums usa: Version 8.0 [dataset]. *Minneapolis, MN: IPUMS 10*, D010.
- Runsten, D. and P. LeVeen (1981). *Mechanization and Mexican labor in California agriculture*. La Jolla, Calif. :: Program in United States-Mexican Studies, University of California, San Diego,.
- State of California (1963). Report and recommendations of the Agricultural Labor Commission. Technical report, Agricultural Labor Commission, Sacramento, Ca.
- Wooldridge, J. M. (1997). Quasi-likelihood methods for count data. *Handbook of Applied Econometrics* 2, 352–406.
- Wooldridge, J. M. (1999). Distribution-free estimation of some nonlinear panel data models. *Journal of Econometrics* 90(1), 77–97.
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.
- Zeira, J. (1998). Workers, machines, and economic growth. *The Quarterly Journal of Economics* 113(4), 1091–1117.

Table 1: Number of Harvesting and Mowing Patents between 1948-1985, the Number of Patents Before and After 1965, and the Share of Seasonal Foreign Workers in 1964 by Crop

Crop	Number of Patents 1948-1985	Before 1948-1964	After 1965-1985	Ratio After/Before	Foreign Share of Seasonal Work
Apples	74	27	47	1.7	3.8
Asparagus	51	12	39	3.3	19
Beans	277	127	150	1.2	2.4
Celery	20	8	12	1.5	32.4
Citrus	201	42	159	3.8	21.6
Cotton	772	494	278	.6	3.7
Cucumbers	39	11	28	2.5	27.4
Grapes	137	13	124	9.5	3.3
Lettuce	25	9	16	1.8	55.3
Melons	18	6	12	2	28.4
Potatoes	213	139	74	.5	3.6
Strawberries	37	12	25	2.1	13.8
Sugarbeets	117	63	54	.9	19.9
Sugarcane	155	38	117	3.1	46.9
Tobacco	206	76	130	1.7	1.9
Tomatoes	121	26	95	3.7	26.2
Sum/Median	2463	1103	1360	1.9	19.4

Notes: This table summarizes the outcome measure (number of patents) and the treatment variable (foreign share of seasonal workers) for each crop in the main sample. The last row presents the sum for the second, third, and fourth columns and the median for the last two columns.

Table 2: Effects of Bracero Exclusion on Invention: Baseline Estimates

	(1) Patents	(2) Citations
Foreign share \times post	3.258*** (0.474)	2.271*** (0.497)
Average response	2.87	10.79
N (crops \times years)	608	608
Mean patents/citations before 1965	4.06	23.90
Treatment mean	0.19	0.19
Treatment sd	0.16	0.16
Year FE	Yes	Yes
Crop FE	Yes	Yes

Notes: Difference-in-differences regressions with continuous treatment compare changes in patenting per year in more exposed crops with changes in less exposed crops: $\ln [\mathbb{E}(Innovation_{it}|X_{it})] = \beta \cdot ForeignShare_i \cdot post_t + \gamma_i + \delta_t$ where $Innovation_{it}$ is the number of US patents/citations in crop i and year t , $ForeignShare_i$ is the foreign percentage of seasonal workers in crop i in 1964, $post_t$ indicates years after 1964, and γ_i and δ_t are crop and year fixed effects, respectively. The table reports the Poisson quasi-maximum likelihood estimators of the percentage change in innovations resulting from an increase of one percentage point in the exposure to foreign labor. The average response is the estimated change in the number of patents/citations per year for a one standard deviation increase in the exposure at the average number of patents/citations per crop and year before 1965. All specifications include crop and year fixed effects. Robust standard errors are shown in parentheses.

Table 3: Effects of Bracero Exclusion on Invention: Instrumental Variables

	Patents			Citations		
	(1)	(2)	(3)	(4)	(5)	(6)
Foreign share \times post	4.849*** (1.565)	2.968* (1.622)	4.466*** (1.499)	5.272*** (1.587)	4.272** (1.742)	4.968*** (1.538)
Instruments	Distance	Population	Both	Distance	Population	Both
N (crops \times years)	608	608	608	608	608	608
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Crop FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Difference-in-differences regressions with instrumental variables: $Innovation_{it} = \exp[\beta \cdot ForeignShare_i \cdot post_t + \gamma_i + \delta_t] \cdot \epsilon_{it}$ where ϵ_{it} is a unit-mean error term. The treatment variable $ForeignShare_i \cdot post_t$ is instrumented by $z_i \cdot post_t$, where z_i is either the average distance from Mexico, or the average percentage of the Mexican population in 1940 of the counties growing the crops (or both). The dependent variable is the number of patents in columns 1-3 and the number of forward citations in columns 4-6. Estimators presented are based on Mullahy (1997) count-data IV model with multiplicative errors. The results reported in columns 1 and 4 use the average distance from Mexico as an instrument variable. Columns 2 and 5 show the results using the Mexican population IV, and in column 3 and 6 both instruments are used. The average distance from Mexico of a crop i is measured by $d_i = \sum_c d_c w_{ic}$ where d_c is the minimal distance between the Mexican border and the centroid of county c , and w_{ic} is the percent of acreage of crop i in county c out of the total acreage of crop i . The average Mexican population of a crop is calculated in a similar way using data from the 1940 US population census. All specifications include crop and year fixed effects. Robust standard errors are shown in parentheses.

Table 4: Effects of Bracero Exclusion on Invention: Continuous Difference in Differences with Synthetic Treatment

	(1)	(2)
	Patents	Citations
Foreign share \times post	3.588*** (0.517)	2.474*** (0.557)
Synthetic foreign share \times post	2.392** (0.981)	1.377 (1.188)
N (crops \times years)	608	608
Mean patents/citations before 1965	4.06	23.90
Year FE	Yes	Yes
Crop FE	Yes	Yes

Notes: Poisson quasi-maximum likelihood estimators of Difference-in-differences model with continuous treatment and "synthetic" continuous treatment: $\ln [\mathbb{E}(Innovation_{it}|X_{it})] = \beta \cdot ForeignShare_i \cdot post_t + \alpha \cdot ForeignShare_i^{syn} \cdot post_t + \gamma_i + \delta_t$. A "synthetic" foreign share of a crop i is the weighted average of foreign shares of all other crops, where the weights are a measure of the similarity between the crops, measured by the number of patents in the sample that mentions both crops. The weights are normalized to sum to one. All specifications include crop and year fixed effects. Robust standard errors are shown in parentheses.

Table 5: Effects of Bracero Exclusion on Invention in Labor Intensive Tasks: Triple-difference Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	Patents	Citations	Patents	Citations	Patents	Citations
Foreign percentage \times labor-class \times post	3.224*** (0.964)	2.271** (1.052)				
Foreign percentage \times cost-class \times post			3.133*** (0.953)	2.161** (1.024)		
Foreign percentage \times class \times post					2.459*** (0.550)	1.775*** (0.628)
N (crops \times classes \times years)	1,447	1,447	1,447	1,447	2,096	2,096
Mean patents/citations before 1965	2.19	14.14	2.19	14.14	1.89	12.72
Crop-Class FE	Yes	Yes	Yes	Yes	Yes	Yes
Crop-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Class-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Triple-difference regressions with continuous treatment comparing the effect of Bracero exclusion on patenting in labor-intensive tasks with the effect in less labor-intensive tasks: $\ln[\mathbb{E}(Innovation_{ijt}|X_{ijt})] = \beta \cdot ForeignShare_i \cdot Intensity_{ij} \cdot post_t + \gamma_{ij} + \delta_{it} + \epsilon_{jt}$. $Innovation_{ijt}$ is the number of US patents/citations in crop i , technological class j , and year t . $ForeignShare_i$ is the foreign percentage of seasonal workers in crop i in 1964. $Intensity_{ij}$ is a measure of labor inputs required to to perform task j in crop i . $post_t$ indicates years after 1964. γ_{ij} , δ_{it} , and ϵ_{jt} are crop-task, crop-year, and task-year fixed effects, respectively. The table reports the Poisson quasi-maximum likelihood estimators of β . In the four first columns, I use information on the labor requirement by crop and task to measure the relative labor intensity of a crop-class pair. In columns (1) and (2), $Intensity_{ij}$ is the percentage of hours of labor required for tasks in class j for producing crop i , while in columns (3) and (4) it is the relative labor cost. In columns (5) and (6), $Intensity_{ij}$ equals one for harvesting and mowing tasks (the most labor intensive class on average) and zero for other classes. Robust standard errors are shown in parentheses.




Figure 1: Correlation between Invention and Market Values

Notes: This figure shows the correlation between the innovation activity related to a crop and the crop's market value. Innovation is measured by the log of the average number of US patents related to harvesting technologies in 1948-1985. The text-search algorithm for allocating patents to crops is described in the text. Log average value of production in the years 1948-1985 in 1980 dollars. Data in market values exist for all crops in the sample except apples. The coefficients (and robust standard errors) of the fitted line are -9.89 (1.92) and 0.81 (0.14) for the intercept and slope, respectively. The Pearson correlation coefficient is 0.80

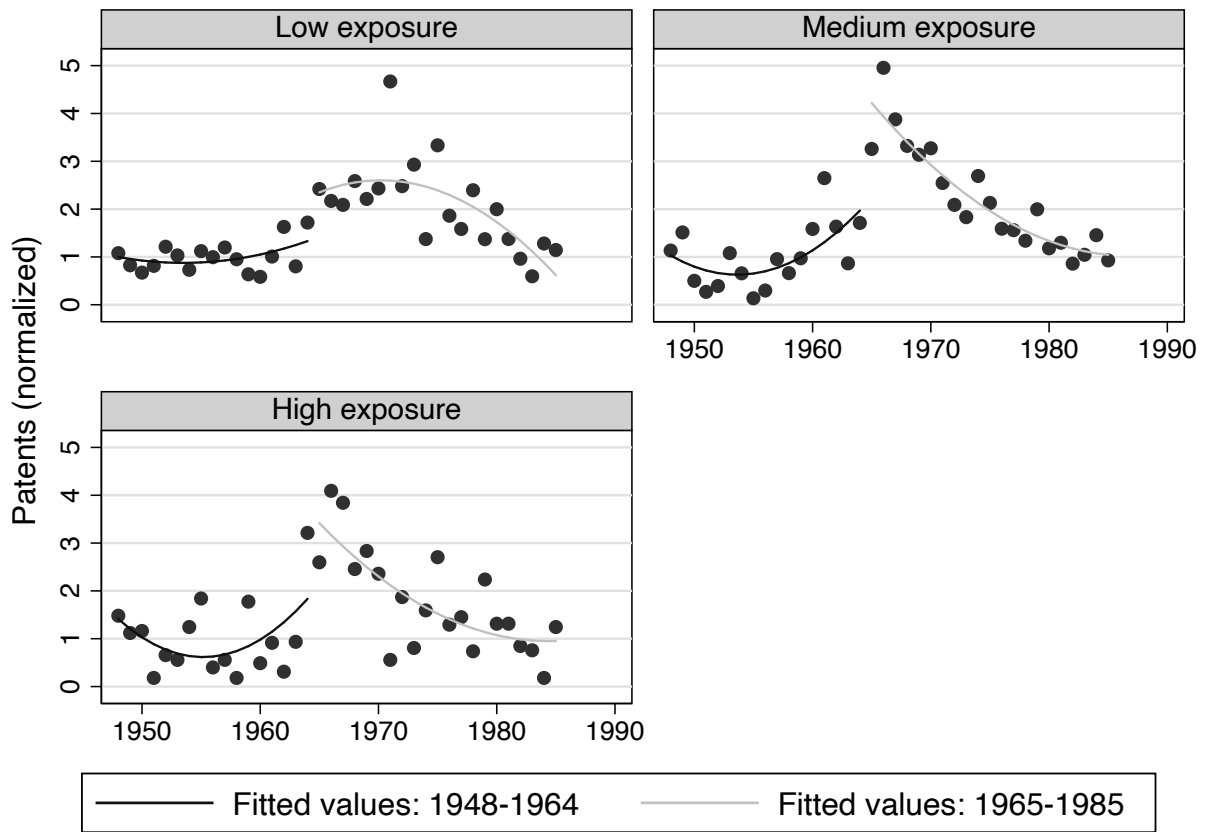


Figure 2: Invention over Time for Crops with Low, Medium and High Exposure to the Bracero Program

Notes: Low exposure: six crops with at most 3.8 percent of foreign workers. Medium exposure: five crops with between 3.8-26.2 percent foreigners. High exposure: five crops with at least 26.2 percent foreigners. The normalized patents measure is the average normalized number of patents for the crops in the exposure group. Each crop-year observation is divided by the crop's pre-period (1948-1964) average number of patents per year.




Figure 3: Effects of Bracero Exclusion on Invention: Event Study

Notes: Event study regression with continuous treatment comparing patenting per year in more exposed crops with patenting in less exposed crops: $\ln [\mathbb{E}(Innovation_{it}|X_{it})] = \beta_t \cdot ForeignShare_i + \gamma_i + \delta_t$ where $Innovation_{it}$ is the number of US patents in crop i and year t , $ForeignShare_i$ is the foreign percentage of seasonal workers in crop i in 1964 (in percentage points), β_t is the bi-annual indicator variable and γ_i and δ_t are crop and year fixed effects, respectively. The graph plots the Poisson quasi-maximum likelihood estimators of β_t and the 95 percent confidence interval (using robust standard errors) of these coefficients.

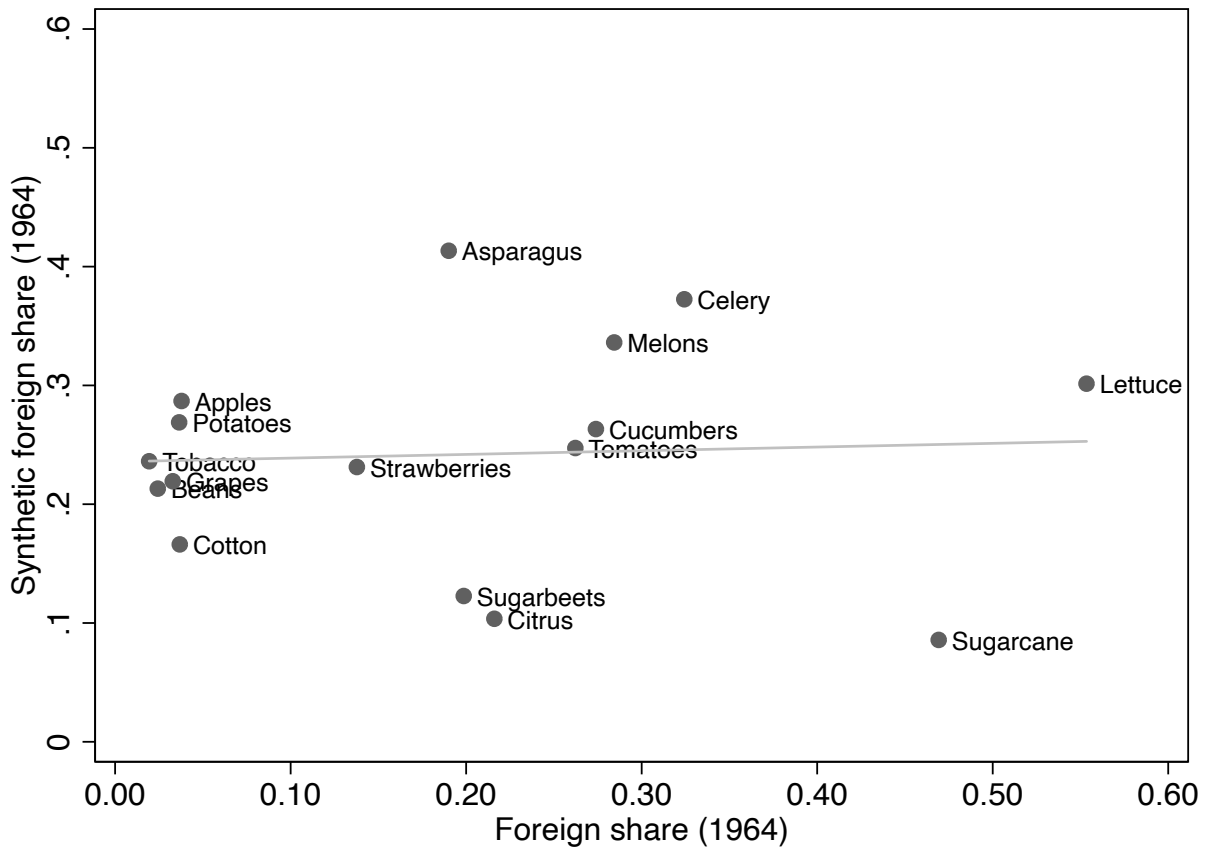


Figure 4: Correlation between Actual and Synthetic Share of Foreign Seasonal Workers in 1964

Notes: This figure show the correlation between the share of foreign seasonal workers in 1964 and the synthetic share by crop. A "synthetic" foreign share of a crop i is the weighted average of foreign shares of all other crops, where the weights are a measure of the similarity between the crops, measured by the number of patents in the sample that mentions both crops. The weights are normalized to sum to one. The coefficients (and robust standard errors) of the fitted line are 0.17 (0.13) and 0.10 (0.47) for the intercept and slope, respectively. The Pearson correlation coefficient is 0.06

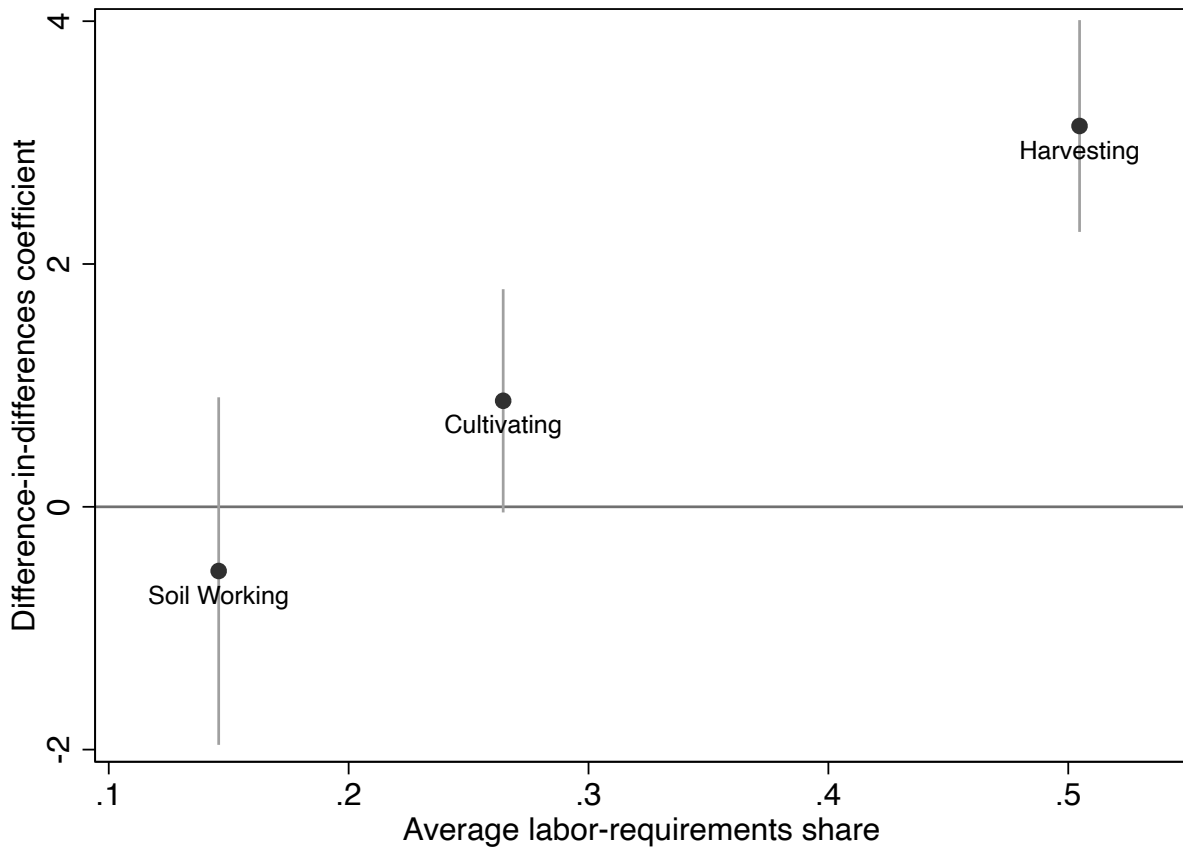


Figure 5: Difference-in-Differences coefficients by Technological Class

Notes: This graph shows the difference-in-differences estimates of the effect of exposure to the Bracero program on innovation (equation 1) for three agricultural CPC patent subclasses separately. The y-axis shows the Poisson quasi-maximum likelihood estimates and their 95 percent confidence interval (using robust standard errors). The x-axis shows the share of labor requirements for the corresponding technological subclass over the total labor requirements of a crop, averaged across the crops.




Figure 6: Effects of Bracero Exclusion on Farm Values

Notes: Event study regression with continuous treatment comparing farm values per agricultural-census year in more exposed counties with farm values in less exposed counties: $\ln(Value_{ct}) = \sum_{\tau=1950}^{1987} \beta_{\tau} \cdot \mathbb{I}(t = \tau) \cdot Exposure_c + \gamma_c + \delta_t + \epsilon_{ct}$. $Value_{ct}$ is the value of an acre of agricultural land in county c in census year t . $Exposure_c$ is a measure of the exposure of county c to the Bracero program, calculated by $Exposure_c = \sum_i ForeignShare_i \cdot AcreageShare_{ic}$ where $ForeignShare_i$ is the foreign percentage of seasonal workers in crop i and $AcreageShare_{ic}$ is the share of crop i in the total acreage of county c in the 1964 census. The exposure is normalized to have a mean of zero and a unit standard deviation. β_t is a census specific indicator variable and γ_c and δ_t are county and year fixed effects, respectively. The graph plots the OLS estimators of β_t and the 95 percent confidence interval of these coefficients. Standard errors are clustered at the county level.

APPENDICES (FOR ONLINE PUBLICATION)

A APPENDIX TABLES AND FIGURES

Table A1: Timeline of Events

Date	Event
August 1942	Wartime program started
January 1948	Postwar era: Braceros contracted directly with US employers
August 1951	Congress approved Public Law 78, which served as the statutory basis for the program until its end
March 1962	US government required farmers to offer Braceros at least the statewide average wage
December 1964	Termination of the program

Notes: The table is based on Craig (1971).

Table A2: Summary Statistics for Crops in the Sample, United States, 1948-1985.

Crop	total labor	domestic labor	foreign labor	foreign share	acreage	production	value
Lettuce	122,500	54,600	67,800	0.553	220,351	44,483	498,007
Sugarcane	105,700	56,100	49,600	0.469	588,511	491,619	698,747
Celery	44,400	30,000	14,400	0.324	34,434	15,389	172,452
Melons	64,700	46,300	18,400	0.284	426,346	41,421	313,727
Cucumbers	105,500	76,600	28,900	0.274	178,723	14,914	158,070
Tomatoes	345,100	254,600	90,500	0.262	474,035	131,244	887,708
Citrus	319,800	250,800	69,100	0.216	.	225,213	1,456,449
Sugarbeets	160,600	128,700	31,900	0.199	1,093,495	442,956	703,143
Asparagus	60,500	49,000	11,500	0.190	122,811	2,936	118,202
Strawberries	308,500	266,100	42,500	0.138	72,702	5,551	265,991
Apples	132,000	127,000	5,000	0.038	.	.	.
Cotton	1,769,400	1,704,200	65,200	0.037	14,420,034	61,122	4,898,662
Potatoes	246,600	237,700	9,000	0.036	1,422,176	292,770	1,613,069
Grapes	179,600	173,700	5,900	0.033	592,258	82,798	701,520
Beans	263,100	256,700	6,400	0.024	1,481,324	18,670	415,374
Tobacco	767,200	752,300	14,900	0.019	1,124,646	19,458	3,137,989

Notes: Seasonal hired labor, by crop and origin of worker, United States, 1964 and average acreage, production and value by crop, United States, 1948-1985. Data from Farm Labor Developments and USDA annual statistical bulletins (see Appendix C for details). Seasonal labor in person-months, acreage in acres, production in 1000 Cwt (100,000 pounds), and value of production in 1980 dollars. Crops listed in descending order of foreign seasonal hired labor relative to the total.

Table A3: Measures of Market Concentration

	Sample		Other patents	
	1948-1964	1965-1985	1948-1964	1965-1985
N. assignees	637	800	210531	317974
Av. patents per assignee	1.7	1.7	3.7	4.4
Herfindahl-Hirschman Index	125.1	96.3	21	11.3
Share patents by top 1 assignee	7.3	6.3	2.9	1.3
Share patents by top 3 assignees	16	14	5.8	2.7
Share patents by top 5 assignees	21.1	18.5	7.6	3.9
Share patents by top 10 assignees	26.4	23.5	10.8	6.8
Share patents by top 30 assignees	35.7	31.9	17.3	14.3
Share patents by top 50 assignees	40.6	36.6	21	19.1
Share patents by top 1 percent	22.4	21.9	51.5	59
Share patents by top 3 percent	32.1	30.1	59.8	66.4
Share patents by top 5 percent	36	34.4	63.9	69.9
Share patents by top 10 percent	43	41.5	70	75
Share patents by top 30 percent	59.6	58.8	81.3	84.2
Share patents by top 50 percent	71.1	70.6	86.6	88.7

Notes: This table reports various measures of the concentration of patents issued by assignees. The Herfindahl-Hirschman Index is the sum of squares of the share percentage of patents issued by each assignee (ranges between 0 and 10,000). The sample consists of USPTO patents in the CPC "harvesting and mowing" subclass that mention one of the 16 crops (see text for details). Other patents include all other USPTO patents. All measures are calculated separately for each group of patents (sample/other patents) and years (1948-1964/1965-1985).

Table A4: Effects of Bracero Exclusion on Invention: Alternative Standard Errors

	Analytical clustered SEs		Bootstrap clustered SEs	
	(1) Patents	(2) Citations	(3) Patents	(4) Citations
Foreign share \times post	3.258*** (1.257)	2.271*** (0.825)	3.258** (1.487)	2.271** (0.975)
N (crops \times years)	608	608	608	608

Notes: This table shows alternative standard errors for the baseline specifications reported in Table 2. The first two columns show the analytical clustered standard errors. The last two columns show bootstrap clustered standard errors based on 1,000 replications.

Table A5: Effects of Bracero Exclusion on Agricultural Invention: Alternative Definitions of the Treatment

	Baseline		Binary		Peak season		Post=1962		Change 64-65	
	(1) Patents	(2) Citations	(3) Patents	(4) Citations	(5) Patents	(6) Citations	(7) Patents	(8) Citations	(9) Patents	(10) Citations
Foreign share \times post65	3.258*** (0.474)	2.271*** (0.497)								
Binary exposure \times post65			0.925*** (0.146)	0.603*** (0.163)						
Peak season \times post65					2.718*** (0.402)	1.848*** (0.426)				
Foreign share \times post62							3.324*** (0.509)	2.539*** (0.540)		
Foreign share change 64-65 \times post65									2.858*** (0.758)	1.620** (0.759)
N (crops \times years)	608	608	608	608	608	608	608	608	608	608
Mean patents/citations before 1965	4.06	23.90	4.06	23.90	4.06	23.90	4.06	23.90	4.06	23.90
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crop FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table checks the sensitivity of the results to the definition of the treatment. The first two columns repeat the baseline specification, where the continuous treatment is the share of foreign workers in the total seasonal employment and the "post" year is 1965, the first year after the abrogation of the *bracero* program. The following two columns use a binary treatment: crop is in the treatment group if the foreign percentage is above the median. In columns (5) and (6), the treatment is defined according to the foreign share at the date of peak foreign employment of each crop. Columns (7) and (8) use the baseline (continuous) measure of the crop' exposure to the Bracero exclusion, but change the "post" year to be 1962, when the US administration started to restrict the program. The last two columns use the 1964 to 1965 change in the share of foreign workers to define the treatment. All specifications include crop and year fixed effects. Robust standard errors are shown in parentheses.

Table A6: Effects of Bracero Exclusion on Agricultural Invention, Robustness to the Text-search Algorithm

	First crop		Maximal crop		All crops		Equal weights		Proportional weights	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Patents	Citations	Patents	Citations	Patents	Citations	Patents	Citations	Patents	Citations
Foreign share \times post	3.258*** (0.474)	2.271*** (0.497)	3.223*** (0.467)	2.163*** (0.500)	3.028*** (0.449)	2.220*** (0.498)	3.046*** (0.443)	2.128*** (0.464)	3.182*** (0.457)	2.178*** (0.483)
N (crops \times years)	608	608	608	608	608	608	608	608	608	608
Mean patents/citations before 1965	4.06	23.90	4.06	23.90	4.37	26.74	4.06	23.90	4.06	23.90
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crop FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table checks the sensitivity of the results to how the text-search algorithm allocates patents to crops. The first two columns repeat the baseline algorithm, where the patent is allocated to the first crop mentioned in the text of the patent. The next two columns allocate the patent to the crop with the maximum mentions in the text. Columns (5) and (6) assign one patent to each one of the crops mentioned in the text. Columns (7) and (8) assign equal weights to each one of the crops mentioned such that the sum of the weights is one. Finally, the last two columns assign weights proportional to the number of times each crop is mentioned. All specifications include crop and year fixed effects. Robust standard errors are shown in parentheses.

Table A7: Effects of Bracero Exclusion on Agricultural Invention, Robustness to the Sample of Crops

	Baseline crops		Baseline + Field		Baseline + California		All crops	
	(1) Patents	(2) Citations	(3) Patents	(4) Citations	(5) Patents	(6) Citations	(7) Patents	(8) Citations
Foreign share \times post	3.258*** (0.474)	2.271*** (0.497)	2.848*** (0.414)	1.481*** (0.442)	3.137*** (0.445)	2.329*** (0.470)	2.765*** (0.399)	1.545*** (0.423)
N (crops \times years)	608	608	988	988	988	988	1,368	1,368
Mean patents/citations before 1965	4.06	23.90	3.59	21.53	2.65	16.17	2.70	16.50
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crop FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table checks the sensitivity of the results to the crops comprising the sample. The first two columns repeat the baseline results, where the sample includes the sixteen crops for which there exist data on the foreign percentage of the total US seasonal labor. Crops are included in the data if they employed 4,000 or more person-months of foreign labor in 1964. In the next two columns, the sample is extended to include the ten greatest field crops (in terms of acreage, according to the 1964 agricultural census), and the foreign exposure of those crops is assumed to be equal to the foreign percentage of the group "Hay and Grain". The sample in columns (5) and (6) includes the baseline sixteen crops and additional ten crops for which data on the percentage of foreign workers in 1962 in California is available. The last two columns include all thirty-six crops together. All specifications include crop and year fixed effects. Robust standard errors are shown in parentheses.

Table A8: Effects of Bracero Exclusion on Agricultural Invention, Changing the Period of the Sample

		Total Patents									
Last Year:	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990
First Year:											
1943	2.943*** (0.460)	2.924*** (0.451)	2.894*** (0.445)	2.939*** (0.442)	2.884*** (0.439)	2.886*** (0.434)	2.826*** (0.431)	2.761*** (0.430)	2.712*** (0.428)	2.702*** (0.422)	2.717*** (0.422)
1944	2.965*** (0.464)	2.946*** (0.455)	2.916*** (0.449)	2.961*** (0.446)	2.906*** (0.443)	2.909*** (0.438)	2.849*** (0.435)	2.784*** (0.434)	2.735*** (0.432)	2.725*** (0.427)	2.740*** (0.426)
1945	3.137*** (0.472)	3.116*** (0.463)	3.084*** (0.457)	3.129*** (0.454)	3.074*** (0.450)	3.076*** (0.446)	3.016*** (0.443)	2.951*** (0.442)	2.902*** (0.440)	2.890*** (0.434)	2.906*** (0.434)
1946	3.145*** (0.480)	3.124*** (0.470)	3.091*** (0.464)	3.136*** (0.461)	3.080*** (0.458)	3.083*** (0.454)	3.022*** (0.451)	2.958*** (0.449)	2.908*** (0.447)	2.896*** (0.441)	2.912*** (0.441)
1947	3.278*** (0.492)	3.255*** (0.482)	3.221*** (0.475)	3.265*** (0.473)	3.209*** (0.469)	3.212*** (0.465)	3.150*** (0.462)	3.085*** (0.460)	3.035*** (0.458)	3.022*** (0.452)	3.038*** (0.452)
1948	3.326*** (0.501)	3.301*** (0.491)	3.267*** (0.484)	3.311*** (0.482)	3.255*** (0.478)	3.258*** (0.474)	3.196*** (0.471)	3.132*** (0.469)	3.082*** (0.467)	3.068*** (0.461)	3.084*** (0.461)
1949	3.317*** (0.515)	3.292*** (0.504)	3.257*** (0.497)	3.301*** (0.495)	3.245*** (0.491)	3.248*** (0.487)	3.187*** (0.484)	3.122*** (0.482)	3.072*** (0.480)	3.058*** (0.473)	3.074*** (0.474)
1950	3.313*** (0.529)	3.287*** (0.518)	3.252*** (0.511)	3.296*** (0.509)	3.240*** (0.505)	3.243*** (0.501)	3.181*** (0.498)	3.116*** (0.496)	3.066*** (0.494)	3.052*** (0.488)	3.068*** (0.488)
1951	3.397*** (0.543)	3.369*** (0.532)	3.332*** (0.525)	3.376*** (0.522)	3.319*** (0.518)	3.321*** (0.515)	3.258*** (0.512)	3.193*** (0.510)	3.142*** (0.508)	3.127*** (0.501)	3.143*** (0.501)
1952	3.270*** (0.543)	3.242*** (0.532)	3.205*** (0.525)	3.249*** (0.523)	3.192*** (0.519)	3.194*** (0.516)	3.132*** (0.513)	3.066*** (0.511)	3.015*** (0.508)	3.000*** (0.502)	3.016*** (0.502)
1953	3.001*** (0.531)	2.975*** (0.520)	2.939*** (0.513)	2.983*** (0.512)	2.926*** (0.508)	2.929*** (0.504)	2.866*** (0.502)	2.801*** (0.499)	2.750*** (0.497)	2.737*** (0.491)	2.752*** (0.492)

Notes: This table checks the sensitivity of the results to the period of the analysis. Every cell in the table reports the Poisson quasi-maximum likelihood estimator of β in the equation $\ln[\mathbb{E}(Innovation_{it}|X_{it})] = \beta \cdot ForeignShare_i \cdot post_t + \gamma_i + \delta_t$ where the analysis sample begins at one of the years 1943-1953 and end in one of the years 1980-1990. $Innovation_{it}$ is the number of US patents in crop i and year t , and the other variables are as explained above. Robust standard errors are shown in parentheses.

Table A9: Effects of Bracero Exclusion on Agricultural Invention, Robustness to the Econometric Model

	Poisson		Negative binomial		Zero-inflated Poisson		OLS	
	(1) Patents	(2) Citations	(3) Patents	(4) Citations	(5) Patents	(6) Citations	(7) Log(patents)	(8) Log(citations)
Foreign share \times post	3.258*** (0.474)	2.271*** (0.497)	2.230*** (0.449)	1.978*** (0.596)	2.938*** (0.463)	1.808*** (0.424)	1.453*** (0.365)	1.985*** (0.550)
N (crops \times years)	608	608	608	608	608	608	446	446
Mean patents/citations before 1965	4.06	23.90	4.06	23.90	4.06	23.90	6.20	36.52
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crop FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table checks the sensitivity of the results to the econometric model in use. The dependent variable is the number of patents in odd columns and the number of citations in even columns. The first two columns repeat the baseline results of the Poisson quasi-maximum likelihood model. Columns (3) and (4) report the results of the negative binomial model. The next two columns show the estimates of a zero-inflated Poisson model where the equation that determines the observed count is zero is logit with the same covariates as the main estimation equation. The last two columns are the results of an OLS model, where the dependent variable is the natural log of patents or citations, and crop-year pairs with zero patents/citations are not included in the regression. All specifications include crop and year fixed effects. Robust standard errors are shown in parentheses.

Table A10: Effects of Bracero Exclusion on Invention: Controlling for Linear Pretrends

	Baseline		Linear pretrends	
	(1) Patents	(2) Citations	(3) Patents	(4) Citations
Foreign share \times post	3.258*** (0.474)	2.271*** (0.497)	4.856*** (0.944)	4.547*** (1.052)
N (crops \times years)	608	608	608	608
Mean patents/citations before 1965	4.06	23.90	4.06	23.90
Year FE	Yes	Yes	Yes	Yes
Crop FE	Yes	Yes	Yes	Yes
Crop-specific linear pre-trends	No	No	Yes	Yes

Notes: Poisson quasi-maximum likelihood estimators of the Difference-in-differences model with continuous treatment. The regressions reported in columns 3-4 include crop-specific linear pre-trends: $\ln [\mathbb{E}(Innovation_{it}|X_{it})] = \beta \cdot ForeignShare_i \cdot post_t + \eta_i \cdot t \cdot (1 - post_t) + \gamma_i + \delta_t$. Robust standard errors are shown in parentheses.

Table A11: Correlation Matrix

	Foreign share	Value/ labor	Acreage/ labor	Value/ acreage	Value	Labor	Acreage
Foreign share	1.000						
Value/ labor	-0.102 (0.729)	1.000					
Acreage/ labor	-0.304 (0.291)	0.459 (0.115)	1.000				
Value/ acreage	0.303 (0.314)	0.087 (0.777)	-0.628 (0.022)	1.000			
Value	-0.450 (0.106)	0.378 (0.182)	0.573 (0.041)	-0.149 (0.626)	1.000		
Labor	-0.415 (0.110)	0.193 (0.509)	0.526 (0.053)	-0.176 (0.566)	0.976 (0.000)	1.000	
Acreage	-0.344 (0.228)	0.199 (0.516)	0.688 (0.007)	-0.286 (0.343)	0.908 (0.000)	0.934 (0.000)	1.000

Notes: Pairwise correlation between the variables. Observations are crops ($N = 16$). %Foreign is the share of foreign seasonal workers in the total seasonal labor in 1964. Seasonal labor in 1964 in person-months units. Average acreage in 1948-1964 in acres. Average value of production in 1948-1964 in 1980 dollars. P-values in parentheses.

Table A12: First Stage: Regressions of the Exposure to the Bracero Program on Instruments

	Foreign percentage \times post		
	(1)	(2)	(3)
Log Av. distance from Mexico \times post	-0.099*** (0.012)		-0.016 (0.023)
Av. Mexican population percentage \times post		1.375*** (0.141)	1.246*** (0.281)
N (crops \times years)	608	608	608
Year FE	Yes	Yes	Yes
Crop FE	Yes	Yes	Yes
F-statistic on instrument(s)	66.67	119.74	60.03

Notes: This table reports the coefficients on the instruments from OLS regressions of the treatment variable $ForeignShare_i \cdot post_t$ on the instrument $z_i \cdot post_t$, where z_i is either the log average distance from Mexico, or the average percentage of the Mexican population in 1940 of the counties growing the crops (or both). The average distance from Mexico of a crop i is measured by $d_i = \sum_c d_c w_{ic}$ where d_c is the minimal distance between the Mexican border and the centroid of county c , and w_{ic} is the percent of acreage of crop i in county c out of the total acreage of crop i . The average Mexican population of a crop is calculated in a similar way using data from the 1940 US population census. All specifications include crop and year fixed effects. Robust standard errors are shown in parentheses. Last row reports the Cragg-Donald Wald F-statistic.

Table A13: Similarity Matrix

Crop	Apples	Aspara	Beans	Celery	Citrus	Cotton	Cucumb	Grapes	Lettuc	Melons	Potato	Strawb	Sugarb	Sugarc	Tobacc	Tomato
Apples	0	1	2	0	58	3	1	11	1	1	7	3	2	0	2	8
Asparagus	8	0	0	17	0	0	0	8	25	0	8	0	8	0	0	25
Beans	3	0	0	3	3	25	7	14	3	2	7	3	3	2	7	17
Celery	0	10	10	0	0	0	10	0	25	0	10	0	0	0	10	25
Citrus	63	0	2	0	0	5	4	9	1	1	7	2	1	0	0	6
Cotton	7	0	25	0	8	0	3	3	2	5	10	8	7	2	12	7
Cucumbers	1	0	6	3	6	3	0	6	6	7	15	8	1	0	7	31
Grapes	24	2	15	0	19	4	7	0	4	2	7	4	0	0	2	11
Lettuce	3	9	6	14	3	3	11	6	0	6	9	6	0	0	0	26
Melons	4	0	4	0	4	12	19	4	8	0	4	0	0	0	4	38
Potatoes	7	1	3	2	6	5	9	3	2	1	0	7	24	1	5	26
Strawberries	10	0	5	0	5	13	15	5	5	0	20	0	3	0	5	15
Sugarbeets	4	2	4	0	2	9	2	0	0	0	64	2	0	4	2	2
Sugarcane	0	0	14	0	0	14	0	0	0	0	14	0	29	0	29	0
Tobacco	5	0	10	5	0	17	12	2	0	2	15	5	2	5	0	20
Tomatoes	7	2	8	4	5	3	17	5	7	8	24	5	1	0	6	0

Notes: The similarity between two crops is measured by the number of patents in the sample that mention both crops. The weights are normalized such that each row sums to one hundred.

Table A14: Plant-Agricultural Subclasses in the CPC Classification System: Definition of the Subclass, Number of Crop-Specific Patents and Labor Requirements

Subclass	Definition	Patents			Labor share	
		1948-64	1965-85	Total	mean	sd
B	Soil Working In Agriculture Or Forestry; Parts, Details, Or Accessories Of Agricultural Machines Or Implements, In General	204	195	399	0.15	0.14
C	Planting; Sowing; Fertilising	192	288	480	0.04	0.07
D	Harvesting; Mowing	981	936	1,917	0.50	0.25
F	Processing Of Harvested Produce; Hay Or Straw Presses; Devices For Storing Agricultural Or Horticultural Produce	50	77	127	0.03	0.06
G	Horticulture; Cultivation Of Vegetables, Flowers, Rice, Fruit, Vines, Hops Or Seaweed; Forestry; Watering	198	581	779	0.26	0.15
N	Preservation Of Bodies Of Humans Or Animals Or Plants Or Parts Thereof; Biocides, E.G. As Disinfectants, As Pesticides, As Herbicides Pest Repellants Or Attractants; Plant Growth Attractants; Plant Growth Regulators	3	38	41	0.02	0.01

Notes: The table shows the definition and summary statistics for the six subclasses of the A01 class (Agriculture) in the Cooperative Patent Classification (CPC), which are related to plants. Columns (3)-(5) show the number of US patents belonging to each subclass that mention one of the crops in the extended sample (Baseline + California) in 1948-1964, 1965-1985, and 1948-1985, respectively. The sixth column reports the share of hours of labor related to each subclass required to produce an acre of a crop, averaged over eighteen crops for which there exist information on both the seasonal foreign labor share and labor requirements in California in 1960. The last column reports the standard deviation of those averages.

Table A15: Task Classification

Task	Class	Task	Class	Task	Class
balancing	D	harvest	D	planting cover crops	C
banking out	D	harvest labor	D	plowing	B
bird control	G	harvesting	D	post-harvest cleanup	B
blight cutting	G	hauling	D	post-harvest mite control	N
brush disposal	B	hauling out	D	pre-irrigation	G
brush removal	B	hauling seed	D	preparing ground for harvest	B
brush shredding	B	hauling to market	D	propping	G
checking	G	hauling to mill	D	pruning	G
chiseling	B	hoeing	B	racking	D
chopping fern	B	housing	D	raking	D
combining	D	hulling	F	removing	C
cover crop	G	hulling boxes	F	renovation	B
cultivating	G	insecticide application	N	ridging	B
cultivation	G	inspection	D	rodent control	N
cultural labor	G	irrigating	G	rolling	B
cutting	D	irrigation	G	rolling beds	B
cutting (seed)	C	irrigation preparation	G	scratching	G
defoliate air application	N	knocking	D	seeding	C
digging	D	land planing	B	shaking	D
digging around trees	G	land preparation	B	shping beds	B
disease control	N	leveling	B	sledding operation	D
disease prevention	N	installing flumes	G	splitting ridges	B
disking	B	list	B	spraying	N
distributing	D	listing beds	B	subsoiling	B
ditching	B	loading	D	subsoilingiling	B
draining	G	lrrigating	G	supervision	D
dusting	N	maintaining checks	G	survey for leveling	B
dusting or spraying	N	manure application	C	surveying	G
fertilizing	C	mechanical harvesting	D	thinning	B
field hauling	D	mowing	D	threshing	F
floating	G	mulching beds	G	tillage	B
flooding	G	packing	D	turning vines	G
frost protection	G	pest	N	tying	G
furrowing	B	pest control	N	weed control	N
furrowing out	B	picking	D	weeding	B
girdling	G	planting	C		
harrowing	B	planting cover crop	C		

Notes: The table lists all task descriptions that appear in the labor requirement data (State of California 1963), and the technological subclass (of the CPC Agriculture class A01) I assigned the task to (see Table A14 for the definition the subclasses). If a task description combines multiple tasks (e.g. "cultivating and fertilizing"), I classify it according to (and report in this table only) the first task.




Figure A1: Share of Patents by Individual Inventors by Year

Notes: This figure shows the share of patents by individual inventors for the years 1948-1985. Patents are classified as invented by individual inventors if one of the inventors is also the patent's assignee. The sample consists of USPTO patents in the CPC "harvesting and mowing" subclass that mention one of the 16 crops (see text for details). Other patents include all other USPTO patents.




Figure A2: Average Number of Inventors per Patent by Year

Notes: This figure shows the average number of inventors per patent for the years 1948-1985. The sample consists of USPTO patents in the CPC "harvesting and mowing" subclass that mention one of the 16 crops (see text for details). Other patents include all other USPTO patents.




Figure A3: Years Between Patent Application and Publication by Year

Notes: This figure shows the average number of years between patent application and publication for the years 1948-1985. The sample consists of USPTO patents in the CPC "harvesting and mowing" subclass that mention one of the 16 crops (see text for details). Other patents include all other USPTO patents.




Figure A4: Invention over Time by Crop

Notes: Low exposure: six crops with at most 3.8 percent of foreign workers. Medium exposure: five crops with between 3.8-26.2 percent foreigners. High exposure: five crops with at least 26.2 percent foreigners. The normalized patents measure is the number of patents divided by the crop's pre-period (1948-1964) average number of patents per year. Observations with (normalized) patent values higher than ten are replaced with ten.

B MODELS

This appendix constructs a theoretical framework to capture the opposing effects of labor supply on technological change. I explicitly introduce two types of technologies. The first is labor-augmenting machines that increase the production for every level of labor (Acemoglu 1998). The second type is labor-saving technologies in the spirit of Zeira (1998).³² I show that an increase in the labor supply encourages the creation of new labor-augmenting technologies but discourages the creation of new labor-saving technologies. Summing up the effects on the two different types of technology, the overall effect of labor supply on technological progress is theoretically unclear.

B.1 LABOR-AUGMENTING IMPROVEMENTS

Following Acemoglu (1998, 2002), the production function of a representative competitive firm is:

$$Y = AL^\beta \tag{B1}$$

where:

$$A = \int_0^1 q_A(a)x_A(a)^\alpha da \quad , \quad q_A(a) \in \{0, 1\} \quad , \quad \alpha, \beta > 0 \quad , \quad \alpha + \beta < 1 \tag{B2}$$

The output produced from labor input L , assumed to be supplied inelastically, and labor-augmenting machines $\{x_A(a)\}$. The technology level is determined by the set of technologies available, $\{q_A(a)\}$.

Technology products are supplied by technology monopolists. Each monopolist sets a rental price $p_A(a)$ for the technology it supplies to the market. Following Alesina et al. (2018), I assume that the invention cost of technology a increases with a :

$$K_A(a) = g_A(a) \quad , \quad g'_A(a) > 0 \tag{B3}$$

For simplicity, I also assume that the cost function is continuous and satisfies the Inada conditions $g_A(0) \rightarrow 0$ and $g_A(1) \rightarrow \infty$. After the invention of the machine, the inventor has full rights on that technology. The marginal cost of producing one machine unit is ψ_A .

³² In a recent paper, Acemoglu and Restrepo (2018) model the invention of new tasks, besides labor-saving inventions. Note, however, that new task inventions in their model are similar to labor-augmenting inventions, as they effectively increase the productivity of manual labor. More precisely, the invention of a new task is equivalent to a positive labor-augmenting shock and a negative labor-saving shock, as the share of automated tasks decreases. One can think about the invention of new tasks as an example of labor-augmenting technology improvements.

Definition of equilibrium: Given labor supply, L , marginal cost of machines production ψ_A , entry cost function $g_A(a)$, and the Cobb-Douglas parameters α and β , an equilibrium is defined by the wage rate w , machine prices $\{p_A(a)\}$, machine quantities $\{x_A(a)\}$, and the set of technologies available $\{q_A(a)\}$, such that:

1. Given the prices and the set of technologies available, $\{x_A(a)\}$ and L maximize the producer's profits.
2. For each task $a \in [0, 1]$ such that $q_A(a) = 1$, the machine price $p_A(a)$ maximizes the monopolist's gross profits.
3. Free entry condition: for each task $a \in [0, 1]$, the monopolist chooses $q_A(a) = 1$ if and only if her net profits are positive.

The competitive producer chooses the quantity of machines of each type $\{x_A(a)\}$, and the quantity of the labor input L in order to maximize profits:

$$\max_{\{x_A(a)\}, L} \left(\int_0^1 q_A(a) x_A(a)^\alpha da \right) L^\beta - wL - \int_0^1 q_A(a) p_A(a) x_A(a) da \quad (\text{B4})$$

where the wage rate w , the prices of the machines $\{p_A(a)\}$, and the set of machines available $\{q_A(a)\}$, are given. The price of the final good is normalized to 1.

From the producer's first order conditions, the demand for machines is:

$$x_A(a) = \left(\frac{p_A(a)}{\alpha} \right)^{-\frac{1}{1-\alpha}} L^{\frac{\beta}{1-\alpha}} \quad (\text{B5})$$

and the demand for labor is:

$$L = w^{-\frac{1}{1-\beta}} (\beta A)^{\frac{1}{1-\beta}} \quad (\text{B6})$$

The inventor is confronted by a two-stage problem: (1) whether to enter the market and pay the fixed cost required to develop the new technology; and (2) to choose the optimal monopolistic price. Starting with the second problem, given that the new technology is available, the inventor maximizes gross profits (not including the entry cost):

$$\max_{p_A(a)} \Pi_A(a) = (p_A(a) - \psi_A) x_A(a) \quad (\text{B7})$$

where the demand function is given in (B5). The optimal monopolistic price is:

$$p_A(a) = \frac{\psi_A}{\alpha} \quad (\text{B8})$$

and the corresponding gross profits are:

$$\Pi_A = \psi_A \left(\frac{1-\alpha}{\alpha} \right) \left(\frac{\psi_A}{\alpha^2} \right)^{\frac{-1}{1-\alpha}} L^{\frac{\beta}{1-\alpha}} \quad (\text{B9})$$

which are independent of a : $\Pi_A(a) = \Pi_A$.

Returning to the inventor's first stage problem: the optimal gross profits are equal across all machine types, but the fixed cost is an increasing function of a .

the assumptions above about the entry cost function $g_A(a)$ guarantee the existence of a unique internal threshold $\bar{a} \in (0, 1)$, such that the inventor has zero net profits:

$$g_A(\bar{a}) = \Pi_A(L) \quad (\text{B10})$$

To see it, note that because $\beta \in (0, 1)$, the gross profits Π_A are positive for each combination of the parameters ψ_A , β , and L . Second, because $g_A(0) \rightarrow 0$, there exist $\epsilon > 0$ small enough such that $g_A(\epsilon) < \Pi_A$. Similarly, because $g_A(1) \rightarrow \infty$, there exist $\delta < 1$ close enough to 1 such that $g_A(\delta) > \Pi_A$. Finally, the monotonicity and continuity of $g_A(a)$ guarantee the existences of a unique $\bar{a} \in (0, 1)$ such that $g_A(\bar{a}) = \Pi_A$.

Next, notice that an increase in the labor supply, L , increases the profits of each technology monopolist, and therefore the technology level, \bar{a} :

$$\frac{\partial \bar{a}}{\partial L} > 0 \quad (\text{B11})$$

This is a pure Acemoglu (1998)'s market size effect: a larger market for the technology, namely more workers who use it, leads to more innovation.

Next, I examine what happens to the other equilibrium outcomes of the model, namely the number of machines (x_A), the TFP (A), the output (Y), and the wage rate (w), when the labor supply is increasing. From equation B9, we have $\Pi_A = C_1 L^{\frac{\beta}{1-\alpha}}$ where C_1 is a constant term that depends only on the parameters ψ_A and β . Because, $g_A(a)$ is monotonically increasing in a , we obtain: $\frac{\partial \bar{a}}{\partial L} > 0$.

From equation B5, the amount of machines is

$$x_A(a) = \left(\frac{\psi_A}{\alpha^2} \right)^{-\frac{1}{1-\alpha}} L^{\frac{\beta}{1-\alpha}} \equiv C_2 L^{\frac{\beta}{1-\alpha}} \quad (\text{B12})$$

which increases with L . Substituting into equation B2, the TFP is

$$A = C_2 \bar{a} L^{\frac{\alpha\beta}{1-\alpha}} \quad (\text{B13})$$

which is also an increasing function of L . So is the output $Y = A \cdot L^\beta$. The wage rate is now:

$$w = C_3 \bar{a} L^{-\frac{1-\alpha-\beta}{1-\alpha}}. \quad (\text{B14})$$

Note that with an exogenous technology level $\bar{a}(L) = \bar{a}$, the wage rate is a decreasing function of L . This results from the decreasing return to scale production function, together with the constant price of machines. However, when the technology level \bar{a} is endogenous, it increases when L increases. This effect dampens the wage response to a change in the labor supply and might even change the sign of the effect.

B.2 LABOR-SAVING IMPROVEMENTS

What happens when technology, rather than augmenting the production of each unit of labor, replaces human labor? In this section, I present a simple model with labor-saving technology progress in the spirit of Zeira (1998). The producer's technology is now:

$$Y = A \int_0^1 (e(l) + q_L(l)x_L(l))^\beta dl \quad (\text{B15})$$

where $\beta \in (0, 1)$ and $q_L(l) \in \{0, 1\}$.³³ Each task, l , can be done by manual labor $e(l)$. If a machine of this type exists, i.e., $q_L(l) = 1$, the task can also be done by labor-replacing machine $x_L(l)$. Manual labor, L , is again assumed to be supplied inelastically. The cost of inventing technology l is:

$$K_L(l) = g_L(l) \quad , \quad g'_L(l) > 0, \quad (\text{B16})$$

where $g_L(l)$ is continuous and satisfies the Inada conditions $g_L(0) \rightarrow 0$ and $g_L(1) \rightarrow \infty$. The

³³ Qualitative similar results obtained when $A = \int_0^1 q_A(a)x_A(a)^\alpha da$, $\alpha + \beta < 1$, and $q_A(a) \in \{0, 1\}$ is given exogenously. For simplicity, I show the results of a model with a fixed A .

marginal cost of producing one machine unit is ψ_L .

Definition of equilibrium: Given the labor supply, L , marginal cost of machines production ψ_L , entry cost function $g_L(l)$, and the Cobb-Douglas parameter β , an equilibrium is defined by the wage rate w , machine prices $\{p_L(l)\}$, manual labor demand $\{e(l)\}$, machine quantities $\{x_L(l)\}$, and the set of technologies available $\{q_L(l)\}$, such that:

1. Given the prices and the set of technologies available, $\{x_L(l)\}$ and $\{e(l)\}$ maximize the producer's profits.
2. For each task $l \in [0, 1]$ such that $q_L(l) = 1$, the machine price $p_L(l)$ maximizes the monopolist's gross profits.
3. Free entry condition: for each task $l \in [0, 1]$, the monopolist chooses $q_L(l) = 1$ if and only if her net profits are positive.
4. The labor market clears: $\int_0^1 e(l)dl = L$

The producer chooses the quantity of machines of each type $\{x_L(l)\}$, and the quantity labor for each task $\{e(l)\}$ to maximize profits:

$$\begin{aligned} \max_{\{x_L(l)\}, \{e(l)\}} A \left[\int_0^1 q_L(l)x_L(l)^\beta dl + \int_0^1 e(l)^\beta dl \right] \\ - \int_0^1 p_L(l)x_L(l)dl - w \int_0^1 e(l)dl \end{aligned} \quad (\text{B17})$$

where the uniform wage rate w , the prices of the machines $\{p_L(l)\}$, and the set of machines available $\{q_L(l)\}$, are given. Because of the perfect substitution between manual labor and machines, if a machine of type l is available, the producer will use only the cheaper factor. Moreover, if no one buys the machine in equilibrium, the machine will not be invented as it is costly to invent it. For simplicity, I assume that if the producer is indifferent between hiring manual labor or machines, she will choose to employ only machines. Taking together, there is a threshold \bar{l} such that tasks $l \leq \bar{l}$ are produced by machines, and tasks $l > \bar{l}$ are produced by labor. From the first-order conditions, the demand for machines is:

$$x_L(l) = p_L(l)^{-\frac{1}{1-\beta}} (\beta A)^{\frac{1}{1-\beta}}, \quad (\text{B18})$$

and the demand for labor is:

$$e(l) = w^{-\frac{1}{1-\beta}} (\beta A)^{\frac{1}{1-\beta}}. \quad (\text{B19})$$

Given that a machine l exists, the monopolist inventor sets the price to maximize gross profits:

$$\Pi_L(l) = (p_L(l) - \psi_L) p_L(l)^{-\frac{1}{1-\beta}} (\beta A)^{\frac{1}{1-\beta}}. \quad (\text{B20})$$

Taking the first order condition with respect to the price, the optimal price is $p_A(a) = p_A = \frac{\psi_A}{1-\beta}$. However, where $w < \frac{\psi_A}{1-\beta}$, because machines and labor are perfect substitutes, if the monopolistic charged this price, the producers would choose to produce with labor and to pay a lower price. Therefore, the monopolistic price is:

$$p_L(l) = \min\left(\frac{\psi_L}{\beta}, w\right). \quad (\text{B21})$$

Additionally, if $w < \psi_L$, the producers would lose from the production of every machine (even without the entry cost), and therefore will choose not to produce (even if the technology already exists).

We can distinguish between three cases: 1) $w \geq \frac{\psi_L}{\beta}$, 2) $\psi_L \leq w \leq \frac{\psi_L}{\beta}$, and 3) $w \leq \psi_L$.

Case 1: $w \geq \frac{\psi_L}{\beta}$. In this range, the price of the machines does not depend on l or L . The gross profits $\Pi_L^* = \left(\frac{\psi_L}{\beta} - \psi_L\right) \left(\frac{\psi_L}{\beta}\right)^{-\frac{1}{1-\beta}} (\beta A)^{\frac{1}{1-\beta}}$ are fixed and positive. Because of the continuity and monotonicity of $g_L(l)$ and the Inada conditions, a unique equilibrium technology level exists that satisfies $\bar{l} \in (0, 1)$. This technology level is independent of L : $\frac{\partial \bar{l}}{\partial L} = 0$.

Case 2: $\psi_L \leq w \leq \frac{\psi_L}{\beta}$. In this case $p_L = w$. The technological level \bar{l} is determined such that the marginal inventor has zero net profits:

$$g_L(\bar{l}) = \Pi_L = (w - \psi_L) \cdot (\beta A)^{\frac{1}{1-\beta}} w^{-\frac{1}{1-\beta}}. \quad (\text{B22})$$

An increase in the wage rate increases the maximal price the inventors can charge for the machines, and because this price is below the optimal unrestricted price, $w < \frac{\psi_L}{\beta}$, this increases the profits of all inventors, hence the technology level. Now, because $g_L(l)$ is continuous, monotonically increasing in l , and satisfies the Inada conditions, for each level of $\Pi_L \geq 0$ a unique $\bar{l} \in [0, 1)$ exists such that $g_L(\bar{l}) = \Pi_L$ and $\frac{\partial \bar{l}}{\partial \Pi_L} > 0$. Hence, we can write $w = h(\bar{l})$ such that $h'(\bar{l}) > 0$. Using the simplifying assumption that if a producer is indifferent between manual labor and machines she will use machines only, together with the uniform wage rate and the decreasing return to scale in each task ($\beta < 1$), we obtain $e(l) = \frac{L}{1-\bar{l}}$. Substituting into equation B19, we obtain:

$$L = \left(\frac{h(\bar{l})}{\beta A}\right)^{-\frac{1}{1-\beta}} (1 - \bar{l}) \quad (\text{B23})$$

Differentiating L with respect to \bar{l} , we get:

$$\frac{\partial L}{\partial \bar{l}} = - \left(\frac{h(\bar{l})}{\beta A} \right)^{-\frac{1}{1-\beta}} - \frac{h'(\bar{l})}{\beta A(1-\beta)} \left(\frac{h(\bar{l})}{\beta A} \right)^{\frac{\beta-2}{1-\beta}}. \quad (\text{B24})$$

Because $h(\bar{l}) > 0$ and $h'(\bar{l}) > 0$, both terms are negative, so we have $\frac{\partial L}{\partial \bar{l}} < 0$ and hence $\frac{\partial \bar{l}}{\partial L} < 0$. Finally, $\frac{\partial w}{\partial L} = \frac{\partial h(\bar{l})}{\partial L} = \frac{\partial h(\bar{l})}{\partial \bar{l}} \frac{\partial \bar{l}}{\partial L} < 0$

Case 3: $w < \psi_L$. In this case, the gross profits of the inventor are negative for every positive amount of production. Hence, there is no reason to pay the fixed cost of the invention, and no machine is invented ($\bar{l} = 0$). Changes in the labor supply in this range do not affect the technology level $\frac{\partial \bar{l}}{\partial L} = 0$.

In what follows, I show that we can divide the parameters space into three disjoint groups that correspond to the three cases. First, note that in all three cases $\frac{\partial w}{\partial L} < 0$. I have already demonstrated this for case 2. For cases 1 and 3, it can be seen from the producer's F.O.C:

$$w = \beta A \left(\frac{L}{1-\bar{l}} \right)^{\beta-1}, \quad (\text{B25})$$

and the fact that $\frac{\partial \bar{l}}{\partial L} = 0$.

Claim: For each set of parameters L, ψ_L, β , and entry cost function $g_L(l)$, there exist a unique equilibrium $(w, p_L, x_L, e, \bar{l})$. Moreover, let $\bar{L} = \left(\frac{\psi_L}{\beta A} \right)^{-\frac{1}{1-\beta}}$ and $\underline{L} = \left(\frac{\psi_L}{\beta^2 A} \right)^{-\frac{1}{1-\beta}} (1 - g_L^{-1}(\Pi_L^*))$, then:

1. $w \geq \frac{\psi_L}{\beta} \iff L \leq \underline{L}$.
2. $\psi_L \leq w \leq \frac{\psi_L}{\beta} \iff \underline{L} \geq L \geq \bar{L}$.
3. $w \leq \psi_L \iff L \geq \bar{L}$.

Proof: Assume $w \geq \frac{\psi_L}{\beta}$ and let $L = \underline{L}$. Because we are at the range of case 1, $\Pi_L = \Pi_L^*$ and therefore $\bar{l} = g_L^{-1}(\Pi_L^*)$. From equation B25 we obtain $w = \frac{\psi_L}{\beta}$, so we can verify that indeed $w \geq \frac{\psi_L}{\beta}$. Now, assume that $L < \underline{L}$. Because w is strictly decreasing in L , we have $w > w(\underline{L}) = \frac{\psi_L}{\beta}$. On the other hand, for $L < \underline{L}$ we have $w < \frac{\psi_L}{\beta}$. Therefore, a solution of the type $w \geq \frac{\psi_L}{\beta}$ exists only if $L \leq \underline{L}$, and exists and is unique if $L \leq \underline{L}$.

Now, consider a solution of the type $w \leq \psi_L$. In this case $\bar{l} = 0$ and $w = \beta A L^{\beta-1}$. For $L = \bar{L}$, we have $w = \psi_L$. Because w is strictly decreasing in L , a solution of this type exists if only if $L \geq \bar{L}$ and is unique if $L \geq \bar{L}$.

Finally, consider the case of $\psi_L \leq w \leq \frac{\psi_L}{\beta}$. If $L = \bar{L}$, $w = \psi_L$ and $\bar{l} = 0$ solve the system of equations B22 and B25. Similarly, If $L = \underline{L}$, the unique solution is $w = \frac{\psi_L}{\beta}$ and $\bar{l} = g_L^{-1}(\Pi_L^*)$.

Again, because w is strictly decreasing in L , a solution of this type exists if only if $\underline{L} \geq L \geq \bar{L}$ and is unique if $\underline{L} \geq L \geq \bar{L}$.

Notice that for each of the three ranges of w , a solution exists only if L satisfies the corresponding parametric condition; therefore, there cannot be more than one equilibrium for a given set of parameters. \square

Intuitively, when labor supply is very high, wages are low, so the marginal cost of producing each machine is higher than its potential price, bound above by the wage rate; therefore, no technology is invented. On the other hand, when the labor supply is very low, the optimal monopolistic price of the machines is lower than the wage rate. In this range, the gross profits of the inventors, hence the technological level, are constant and do not react to shifts in the labor supply.

When labor supply is not very high or very low, an exogenous increase in the labor supply decreases the wage rate and therefore decreases the technology level, \bar{l} :

$$\frac{\partial \bar{l}}{\partial L} < 0. \quad (\text{B26})$$

This is Hicks (1932)'s substitution effect. When the labor supply is higher, wages are lower; therefore, the potential saving from each new machine is lower. In this case, fewer labor-saving technologies will be developed.

I have already shown that the technology level \bar{l} does not depend on the labor supply when $L \leq \underline{L}$ and $L \geq \bar{L}$, and is decreasing in L otherwise. I have also shown that w is decreasing in L in all regions. Next, I turn to explore what happens to the other elements of the model.

Case 1: In this region, the technology level and the monopolistic price of the machines are constant in L ; therefore the quantity of machines $x_L(l)$ is also constant for each $l \leq \bar{l}$ (denote it by $x_L(l) = x^*$). The output in this case is $Y = A \left[\bar{l}x^{*\beta} + (1 - \bar{l}) \left(\frac{L}{1 - \bar{l}} \right)^\beta \right]$ which increases in L .

Case 2: In this case we have $w = p_L(l)$ and $e(l) = x(l)$. As $w = \beta A e(l)^{\beta-1}$ and $\frac{\partial w}{\partial L} < 0$, we know that $\frac{\partial e(l)}{\partial L} = \frac{\partial x(l)}{\partial L} > 0$. The output is $Y = A e(l)^\beta$, increases in L .

Case 3: In this case we have $w = \beta A L^{\beta-1}$ and $e(l) = L$. There are no machines. The output is $Y = A L^\beta$, increases in L .

C DATA APPENDIX

C.1 US PATENTS

I collected patent data from Google Patents for successful applications for patents between 1940-1990. The data include the full text of the patent (title, abstract, claims, and description) as well as patent identification numbers, number of citations, CPC classification, and application and publication (issue) years. I used the application, rather than publication year, to define the timing of the invention because the application date is closer to the date of the invention. Publication

dates are typically delayed by several years. I proxy the application date for patents with missing application dates by subtracting the median lag between application and publication dates (2.6 years) from the publication date.

C.2 CROP-LEVEL INFORMATION

Total and foreign seasonal labor by crop for the years 1964-1965 was collected from (Bureau of Employment Security 1966, Table 5, p. 11). Total and foreign seasonal labor by crop at the date of peak foreign employment was collected from (Bureau of Employment Security 1966, Table 21, p. 48). Information for the total and Mexican seasonal labor in California for the additional ten crops was collected from (of California 1963, Table 5, p. 17)

To construct the labor-intensity measures, I collected data on labor requirements per acre (in terms of hours and cost of labor) by task and crop in California in 1960 from the State of California's "Report and Recommendations of the Agricultural Labor Commission" (State of California 1963). This data includes information on person-hours and labor cost per acre for the various tasks required for the production process for California's 25 most valuable crops in 1960. I manually classified each task into one of the six agricultural CPC subclasses. See the definition of these subclasses in Table A14. I classify it according to the first task if a task description combines multiple tasks (e.g., "cultivating and fertilizing"). Table A15 reports the classification of all tasks.

Production, acreage, and value by crop for the years 1940-1990 were collected from various publications of the Department of Agriculture (see Table C1).

Table C1: Sources of data on acreage, production, and value by crop

Crop	Acreage	Production	Value (current prices)
Sugarbeets	NASS (40-90)	NASS (40-90)	NASS (40-90)
Sugarcane	NASS (40-90)	NASS (40-90)	NASS (78-90)
Bean	NASS (40-90)	NASS (40-90)	NASS (40-90)
Cotton	NASS (40-90)	NASS (40-90)	NASS (40-90)
Tobacco	NASS (40-90)	NASS (40-90)	NASS (40-90)
Grapes	NASS (47-90)	NASS (44-90)	NASS (44-90)
Potatoes	NASS (40-90)	NASS (40-90)	NASS (40-90)
Tomatoes	RE (40-59), ERS (60-90)	RE (40-59), ERS (60-90)	RE (40-59), ERS (60-90)
Lettuce	RE (40-49), ERS (50-90)	RE (40-49), ERS (50-90)	RE (40-49), ERS (50-90)
Asparagus	RE (40-49), ERS (50-81,83-90)	RE (40-49), ERS (50-81,83-90)	RE (40-49), ERS (50-81,83-90)
Straebwrrries	RE (40-59, 64-69), AS (60-63), ERS (70-90)	RE (40-59, 64-69), AS (60-63), ERS (70-90)	RE (40-59, 64-69), AS (60-63), ERS (70-90)
Celery	RE (40-59, 64-81), AS (60-63)	RE (40-59, 64-81), AS (60-63)	RE (40-59, 64-81), AS (60-63)
Cucumbers	RE (40-59, 64-70, 74-81)	RE (40-59, 64-70, 74-81)	RE (40-59, 64-70, 74-81)
Mellons	RE (40-81)	RE (40-81)	RE (40-81)
Citrus	No Data	CF (40-81)	CF (40-81)
Apples	No Data	No Data	No Data

Notes: This table uses the following abbreviations for the Department of Agriculture's publications. NASS: National Agricultural Statistics Service. RE: Revised Estimates. ERS: Economic Research Service. AS: Annual Summary. CF: Citrus Fruits.

C.3 COUNTY-LEVEL INFORMATION

Share of Mexicans by county is calculated from the complete count 1940 US population census (Ruggles et al. 2018). Distance from Mexico is the minimal distance between the Mexico border and the centroid of the county. To calculate the weight of each county in the production of each

crop, I use information on the total acreage by crop and county from the 1964 census of agriculture (Haines et al. 2019). Farm values per acre were collected from the census of agriculture for the years 1950, 1954, 1959, 1964, 1969, 1974, 1978, and 1982 (Haines et al. 2019). Finally, the list of Bracero and non-Bracero states used in section 7 are from (Clemens et al. 2018, Figure 2, p. 1476).