



STEG WORKING PAPER

# AGRICULTURAL TECHNOLOGICAL CHANGE, FEMALE EARNINGS, AND FERTILITY: EVIDENCE FROM BRAZIL

Vivek Moorthy

NOVEMBER 2022  
STEG WP041

# Agricultural Technological Change, Female Earnings, and Fertility: Evidence from Brazil

Vivek Moorthy

College of the Holy Cross

September 2022

## Abstract

I study how bias in agricultural technological change affects labor market opportunities and fertility in a modern developing country context. Exploiting plausibly exogenous variation in the adoption of genetically engineered soy across municipalities in Brazil, I show that these technologies reduced female earnings and employment in agriculture, without leading to a reallocation of female labor into other sectors. Further, this technology adoption increased fertility due to increases in overall household earnings and substitution effects driven by the reduction in female earnings and employment. These results suggest that, contrary to historical experience, technological progress in modern developing countries may not improve female labor market opportunities or contribute to fertility decline unless substitution effects are negative and sufficiently large.

*JEL:* J13, J16, J23, J43, I15, 015, 033.

---

I am grateful to Taryn Dinkelman, Daniel Hungerman, and Bill Evans for their advice and guidance on this project. This paper also benefited from discussions and comments from Nilesh Fernando, Joe Kaboski, Lakshmi Iyer, Alejandro Estefan, Paul Shaloka, Giorgi Piriashvili, Luca Parisotto, Marieke Kleemans, Rebecca Thornton, Esther Gehrke, Elise Rudt, and audiences at the Eastern Economic Association Conference, Development Day Conference, NEUDC Conference, STEG Annual Conference, the H2D2 Conference at the University of Michigan, the College of the Holy Cross, the University of Akron, the University of Missouri, and at Notre Dame. All errors are my own. Email the author at [vmoothy@holycross.edu](mailto:vmoothy@holycross.edu).

# 1 Introduction

Technological change is the engine by which economies grow (Solow, 1956; Romer, 1990), yet economists have long recognized that the benefits are not always shared evenly across groups. When labor markets are characterized by occupational sorting across different demographic groups, complementarities between technological change and specific occupations can alter the opportunities of a particular group. These complementarities can shrink or widen existing between-group inequalities. When such occupational sorting is gender-based, these technological changes not only affect the structure of the labor market but may further alter the structure of the family. The introduction of typewriters and computers illustrate these relationships: Throughout the 20th century in the U.S. and Europe, these technologies expanded female work in the service sector (Goldin, 2006; Rotella, 1981; Beaudry and Lewis, 2014; Black and Spitz-Oener, 2010), contributing to the fertility decline that we have come to see as synonymous with development (Galor and Weil, 1996). In this same period, skill-biased technological changes rapidly destroyed manufacturing occupations primarily performed by low-skilled men (Katz and Murphy, 1992; Acemoglu and Autor, 2011). This reduced their real earnings and affected downstream outcomes on the family by increasing divorce and lowering marital prospects (Black et al., 2003; Anelli et al., 2019).

Much of the evidence found in the literature on the relationships between technological change, labor market structure, and family structure comes from developed countries where technological progress has generally expanded female labor market opportunities. Nevertheless, in developing countries, the scope for technological change is vast because they are far from technological frontier and rely heavily on agriculture, a sector where tasks and occupations are starkly divided along gender lines. As a result, innovations in agricultural technologies are likely to change gender-specific tasks. These facts beg the question: Can technological change in developing countries reduce rather than expand earnings opportunities for women? If so, what happens to downstream outcomes such as fertility?

In this paper, I study the differential impacts of agricultural technological change on men’s and women’s labor market opportunities and then exploit these gendered labor market effects to test economic models of fertility, *à la* Becker (1960). To do so, I use the legalization and adoption of genetically engineered (GE) soy technologies in Brazil to generate variation in gender-specific labor outcomes. Brazil legalized herbicide resistant soy in 2003, which allowed farmers to spray herbicides to clear fields without affecting the soy crop. This eliminated the demand for weeding and harvesting operations, which are tasks typically performed by women (Sofa and Doss, 2011; Grassi et al., 2015). At the same time, the new technology complemented occupations commonly performed by men, such as spraying chemicals and operating machinery.

To identify how this new technology affected men and women differently, I estimate difference-in-differences

models that use the legalization of GE soy as the source of variation in gender-specific labor demand across time and the spread of GE soy technologies for variation across space. Rather than using actual yields that depend on endogenous technological adoption choices, I use estimates of potential yields that are a function of the plausibly exogenous geo-climatic conditions favorable for these new technologies.<sup>12</sup> My strategy compares labor market and demographic outcomes in municipalities within states that have higher potential GE soy yields to those municipalities with lower potential GE soy yields after the legalization of GE crop technologies versus before.

Importantly for my outcomes of interest, Brazil collects high-quality data on individual-level earnings for both formal and informal work. This allows me to investigate both female and male labor market outcomes that are typically difficult to uncover in developing countries. Moreover, comprehensive administrative records on live births allow for linking these labor market outcomes to measures of fertility.

First, I found that while this new crop technology led to an overall increase in household earnings, it also came with a large reduction in female earnings opportunities. Municipalities with a one standard deviation increase in soy technological change experienced a 10% reduction in women’s earnings in agriculture. This establishes a novel result: GE soy technological change is female labor-saving. I found that women in soy-producing regions reallocated into work for another member of the household’s job or directly for household sustenance, all for no pay. I found no evidence of movements into other sectors of the economy. Thus, while this technological change in soy production was altering labor markets in agriculture, no alternative employment opportunities for women were being created. These results are striking given the context of overall rising female labor force participation in Brazil over this time period.<sup>3</sup> I also found that the labor market shock driven by the new GE soy technology increased men’s agricultural earnings. Given the occupational sorting of work by gender in agriculture, these technological changes had disparate impacts at the group level that are masked in aggregate statistics.

My second main set of results establish another striking fact: regions with higher adoptions of the GE soy technology experienced higher fertility. I found that municipalities with a one standard deviation increase in GE soy technological change experienced an increase of 2 births per 1000 women, or a 4% increase. I reject the presence of pre-trends in fertility and infant mortality driving the results and find that the effects on fertility are persistent and increasing up to 17 years following the legalization of the GE soy technology. The changes in fertility therefore represent a sustained shift and not a retiming of births. I further find no

---

1. Nunn and Qian (2011) were one of the first to use measures of potential yields in the economics literature. Bustos et al. (2016) proposed the spatial measure used here.

2. While I do confirm that these potential yields predict the actual adoption of GE soy, the timing of the agricultural surveys measuring actual adoption do not correspond with the timing of the labor market surveys. Thus I run an ‘intention-to-treat’ analysis for my main empirical specification.

3. Women in Brazil constituted 39% of the total labor force in 2000, which grew to 44% in 2019. *Source*: World Bank World Development Indicators.

changes in differential migration, infant mortality, or child labor usage.

This finding at first seems counterintuitive because fertility tends to decline as countries develop and technologies alter the structure of labor market. I present a simple model of fertility that incorporates gender-biased technological change to show that basic Beckerian principles can fully rationalize these results. Since Becker (1960), economic models of fertility have stressed that the demand for children changes with earnings opportunities through a combination of income and substitution effects. For the income effects, economists have generated strong empirical evidence that children are normal goods (e.g., Black et al., 2013; Kearney and Wilson, 2018), implying that fertility will increase as income rises, *ceteris paribus*.<sup>4</sup> Typically, however, economic models stress that the net effect of more income hinges crucially on the substitution effect, whereby increases in labor market earnings increase the opportunity cost of time. These opportunity costs make up a large component of the “price” of children and are a key driver of the observed decline in fertility as economies grow richer.<sup>5</sup>

Assuming that children are normal goods, increases in women’s earnings make children more affordable, resulting in positive income effects. However, as women bear most of the time cost of childcare, their increased earnings also raise the opportunity cost of the time they use to rear children. This creates competing negative substitution effects that typically dominate the income effects (Schultz, 1997; Kitchens and Rodgers, 2020).<sup>6</sup> Increases in men’s earnings and overall family earnings serve as positive income effects. In the case of the labor-saving shock in Brazil, it lowered the opportunity cost of female time, thereby reducing the price of children and inducing positive substitution effects. This offsets any reduction in demand for children coming from declining female earnings. By increasing men’s agricultural earnings and overall family earnings, the shock further incentivized higher fertility through positive income effects. Altogether, this demonstrates that the ways in which new technologies affect the structure of the labor market have specific implications for how family structure evolves, and it need not evolve in ways that are favorable for women or that promote fertility decline.

Finally, I find no evidence of changes in child quality as measured by education and infant health outcomes. This is consistent with the quantity-quality trade-off, whereby reductions in the price of children lower investments in the quality of children, offsetting any increases in quality that would normally arise

---

4. Moreover, the predictions of the normal goods assumption are consistently borne out across a variety of contexts (e.g., Dettling and Kearney, 2014; Brueckner and Schwandt, 2015).

5. Becker and Lewis (1973) also incorporated child quality and assume that the income elasticity of quality exceeds that of quantity. Such models still require other special assumptions on preferences, such as a high elasticity of substitution between child quantity and parental consumption, in order to generate a negative income-fertility relationship (Jones et al., 2008). The GE soy technology does not directly impact the returns to education, and so I do not directly model quality choices but instead explore this in the empirical analysis.

6. This is true particularly among poorer women, which would be the case in agricultural employment as female earnings are more directly linked to the “price” of children.

from higher overall earnings.<sup>7</sup>

Given the context of overall declining fertility in Brazil, these effects are empirically generated by the technological change *slowing* the decline in fertility. This slowing of the decline is relevant for other developing countries, particularly in sub-Saharan Africa, where such agricultural technologies are beginning to be adopted following Brazil’s example and where fertility declines have largely stalled due to declines in female economic opportunity (Kebede et al., 2019). To characterize the magnitude of the effects of the GE soy technology in Brazil, I turn to the 2017 Brazilian Agricultural Census, which measures actual GE soy adoption.<sup>8</sup> Using potential yields as an instrument for actual yields, I show that a one standard deviation increase in the percent of farmland devoted to GE soy increases fertility by 3.5 births per 1000 people. This is approximately equal to the total decline in fertility experienced by sub-Saharan Africa over this same time period. Assuming the same effect sizes, this suggests that a similarly timed adoption of GE soy in sub-Saharan Africa could have stalled the fertility decline in the region by about a decade.

I make three key contributions to the literature. First, I expand the literature on women’s economic opportunities through the development process by showing that, contrary to historical experience as documented widely in the literature, economic development from new technologies can be gender-biased in ways that are detrimental to women’s labor market opportunities. This is in contrast to work showing the positive economic impacts for women from the expansion of textile industries (Goldin and Sokoloff, 1982; Heath and Mobarak, 2015) or from agricultural growth more generally (Qian, 2008; Carranza, 2014; Carney and Carney, 2018). My results demonstrate that the nature of technological change and its interactions with the occupational sorting by gender are what determine whether women are made better or worse off in the labor market.<sup>9</sup>

Second, I contribute to the economics literature on fertility by showing that fertility decline may not always accompany modern technological change and development. Development may even *cause* increased fertility by the same economic channels. Most existing studies on women’s labor market conditions and fertility demonstrate how rising female opportunity costs lower the demand for children. For instance, Schultz (1985) and Kitchens and Rodgers (2020) show that rising earnings opportunities for women in Sweden and the U.S., respectively, led to sizable reductions in fertility. Jensen (2012) finds the same relationship in a modern developing country context.<sup>10</sup> Most evidence of the effects of negative labor market shocks for

---

7. This is because increases in the number of children make it more costly for parents to invest in child quality as any given level of quality must now be invested in a larger number of children. For example, it is less costly to invest in a high school education for just one child versus two children.

8. Note, that while the timing of this data does not align with the timing of the measured labor market outcomes, they do align with the annual fertility data.

9. My paper is also related to Alesina et al. (2013)’s work on how historical agricultural practices shape the prevailing gender norms in society. I complement this work by focusing on a modern technological change that affects these particular occupations and tasks where women are currently working in developing countries.

10. Moreover, Gollin et al. (2021) finds that the Green Revolution significantly lowered fertility rates, however they do not focus

women on fertility comes from specific contexts within developed economies (Schaller, 2016; Autor et al., 2019). My results focus on a case of declining female labor market opportunity triggered by the types of technologies that are increasingly relevant for lower income countries. Further, by identifying a setting in which an aggregate positive income shock comes with positive rather than negative substitution effects, I exploit a unique opportunity to test the relevance of the substitution effect channel in explaining fertility behavior. I complement Kitchens and Rodgers (2020)’s findings by providing strong empirical evidence of the importance of this channel in driving fertility change over the course of development. My results confirm the qualitative predictions of economic models of fertility in this setting and stress the importance of turning to these models to anticipate the demographic responses to gendered technological change moving forward.<sup>11</sup>

My third contribution is to the literature on structural change over the development process. Recent theoretical and empirical work has focused on whether there are substantively different implications between industrial productivity growth pulling labor out of agriculture versus agricultural productivity growth eliminating agricultural labor, pushing these workers into other sectors of the economy (Alvarez-Cuadrado and Poschke, 2011; Bustos et al., 2019). My findings show that when structural change occurs from productivity shocks to agriculture, it is likely to interact with the occupational sorting by gender in the sector. Relevant to this, Ngai and Olivetti (2015) explicitly relate structural transformation to the U-shaped pattern of female employment throughout the course of development.<sup>12</sup> They create a theoretical model whereby declines in agricultural employment come with reductions in female labor force participation.<sup>13</sup> I provide causal micro-evidence of this relationship in a modern developing country. Further, Ager et al. (2020) and Gehrke and Kubitzka (2021) demonstrate that structural transformation out of agriculture caused fertility decline in the U.S. historically and in modern Indonesia, respectively. I show that this need not be the case. Rather, given the gendered division of labor, these agricultural productivity drivers of structural change can potentially exacerbate gender inequalities and have different implications for developing economies in the incipient stages of their demographic transition.

The rest of the paper proceeds as follows: Section 2 presents an economic model of fertility, 3 discusses background for the division of occupations in agriculture, fertility, and the context in Brazil. Section 4 discusses the data; Section 5 details my empirical strategy. Sections 6-8 present and discuss the results of my estimations, and Section 9 concludes.

---

on the gender-bias of these technologies. These effects may be rationalized by the fact that the Green Revolution technologies increased employment in agriculture, as these crop technologies required high amounts of labor (Moscona, 2019).

11. Further, this paper is related to the recent literature that revisits models of fertility in high income countries (Doepke et al., 2022). Although still a developing country, Brazil is a low fertility setting, and my results show that determinants of fertility from classic models still have predictive power in these settings.

12. Goldin (1994) provides the canonical framework the U shape pattern of female labor force participation over the development process.

13. Afridi et al. (2022) also show a similar relationship between increased mechanization in Indian agriculture and declines in female labor usage relative to that of men. I focus here on a case of an absolute decline in female labor market earnings.

## 2 Fertility and Gendered Technological Change

The expansion of new agricultural technologies alters many factors in the economy. However, from the household’s perspective, what matters is how these technologies change earnings and how they alter the opportunity costs of children. Here, I use a simple Beckerian model to capture the intuition of how the introduction of this technology can affect fertility through its impacts on gender-specific earnings. Using a one period comparative static framework in which a male and female solve a utility-maximizing lifetime plan between children and consumption, I establish how changes in men’s and women’s earnings affect fertility under the assumptions that children are normal goods and that the time burden of childcare is female intensive. I then embed gender-biased technological change into the model and derive predictions of how this particular change should alter the demand for children. The goal of this model is to highlight how the changing constraints of the household affect fertility choice. Thus the following model focuses on the *positive* comparative statics of the soy technology, rather than the *normative* effects on well-being or the effects of different functional form assumptions on household preferences. I discuss the broader implications on welfare in Section 8, after presenting results.

### 2.0.1 A Simple Model of Fertility Choice

I base my notation off of Galor (2012), and incorporate gender-specific earnings into a simple household model of fertility. Assume a household has preferences over consumption,  $c$ , and the number of children,  $n$ , and that children are normal goods.<sup>14</sup> The household consists of a male and a female, who each have one unit of time that they can supply to the labor market. The male earns  $y^m$  if he supplies the entire unit of his time to the labor market, and the female earns  $y^w$ , which they take as given.<sup>15</sup> Assume that the cost of children enters entirely through the opportunity cost of time used in raising them,<sup>16</sup> a cost that is allowed to vary by gender. The cost of raising each child consists of the fraction of the female’s unit time endowment,  $\tau^w$ , and the fraction of the male’s unit time endowment,  $\tau^m$ , required in child rearing.<sup>17</sup> The household’s budget constraint is then:

$$(\tau^w y^w + \tau^m y^m)n + c \leq y^w + y^m \quad (1)$$

---

14. The normal goods assumption is theoretically justified (Doepke, 2015) and has received strong empirical support (e.g. Black et al., 2013).

15. Note, if income is not earned directly, these given earnings can also be considered the value of each gender’s contribution to overall household income (e.g.,  $y^w$  represents the productive contribution of female work on the farm to overall farm profits).

16. The forgone value of time in raising children makes up a majority of the costs of childcare (Becker, 1992). Moreover, breaking up the cost into opportunity costs of time and direct costs, such as clothing and food prices, does not change any of the predictions. The relevant factors studied here are the opportunity costs from changing labor demands.

17. Allowing  $\tau^m \geq 0$  allows for men to also contribute to childcare.

where the shadow price of a child is the full opportunity cost of raising it,  $\pi_n = (\tau^w y^w + \tau^m y^m)$ . This is a function of the earnings opportunities of the male and female as well as the time commitment of each required to raise children. Assume that the family is approximated by a unified optimizing consumer.<sup>18</sup> For simplicity of exposition, assume household preferences are represented by the following log-linear utility function:

$$u(n, c) = \gamma \ln n + (1 - \gamma) \ln c; \quad \gamma \in (0, 1) \quad (2)$$

The household then maximizes (2) subject to (1), yielding the following demand for children:

$$n^*(\tau^w, \tau^m, y^w, y^m) = \frac{(y^w + y^m)\gamma}{y^w \tau^w + y^m \tau^m} \quad (3)$$

where men and women's earnings enter into the numerator and the denominator, yielding two countervailing forces. First by relaxing the family budget constraint, increases in household earnings increase the demand for children. However, by increasing the shadow price of children, it lowers the demand. To untangle how changes in gender-specific incomes should impact fertility, we can look at the gender-specific comparative statics. The comparative static for women's earnings is negative when:

$$\frac{\partial n^*}{\partial y^w} < 0 \iff y^m[\tau^m - \tau^w] < 0 \quad (4)$$

which depends on the relative burden of childcare. This condition is satisfied if women bear most of the time cost of childcare:  $\tau^w > \tau^m$ . In other words, increases in women's earnings increase the opportunity cost of having children, dominating positive incomes effects, leading to reductions in fertility. By symmetry, this assumption also ensures the comparative static with men's earnings is positive, i.e.:

$$\frac{\partial n^*}{\partial y^m} > 0 \quad (5)$$

## 2.0.2 Incorporating Gendered Technological Change into the Model

I now embed technological change into the framework by parameterizing the earnings of the household as  $y^i(\alpha)$ , for  $i \in w, m$ , where  $\alpha$  represents the technological level. The change in fertility in response to

---

18. While recent work uses Nash-bargaining mechanisms to study fertility behavior (e.g. Rasul, 2008), making this distinction does not create additional empirical content when the effects of changing distribution factors of husbands and wives influence family demands in the same way as the unified model (Schultz, 1997). Thus, the unified model in this case offers a more parsimonious representation of demand behavior.

technological change is given by the total derivative of  $n^*(\tau^w, \tau^m, y^w, y^m)$  with respect to  $\alpha$ :

$$\frac{dn^*(\tau^w, \tau^m, y^w, y^m)}{d\alpha} = \frac{\partial n^*(\tau^w, \tau^m, y^w, y^m)}{\partial y^w} \frac{dy^w(\alpha)}{d\alpha} + \frac{\partial n^*(\tau^w, \tau^m, y^w, y^m)}{\partial y^m} \frac{dy^m(\alpha)}{d\alpha} \quad (6)$$

The partial derivatives of the demand function  $n^*(\tau^w, \tau^m, y^w, y^m)$  with respect to women and men's earnings are established above in equations (4) and (5), with the former being negative and the latter being positive. But the overall effect on the demand for children depends on the signs of  $\frac{dy^w(\alpha)}{d\alpha}$  and  $\frac{dy^m(\alpha)}{d\alpha}$ , or how technological change alters the earnings of each gender.

Other research of gender-specific income shocks from technological change on fertility in the U.S. has studied decreases in men's earnings (e.g. from the expansion of robots (Anelli et al., 2019)), in which case  $\frac{dy^m(\alpha)}{d\alpha} < 0$  and  $\frac{dy^w(\alpha)}{d\alpha} = 0$ , yielding fertility decline. In cases of improvements in both female and male earnings, the sign is ambiguous. Kearney and Wilson (2018) focus on the expansion of fracking technologies where earnings for both men and women increased, or  $\frac{dy^m(\alpha)}{d\alpha} > 0$  and  $\frac{dy^w(\alpha)}{d\alpha} > 0$ . To get a sense of which force would dominate, multiplying (6) by  $\alpha$  and manipulating yields the following elasticity representation:

$$\epsilon_{n,\alpha} = \epsilon_{n,y^w} \epsilon_{y^w,\alpha} + \epsilon_{n,y^m} \epsilon_{y^m,\alpha} \quad (7)$$

where  $\epsilon_{y,x}$  represents the elasticity of variable  $y$  with respect to  $x$ . Thus, the overall responsiveness of fertility to gender-biased technological change is a weighted average of the gender-specific elasticities of fertility demand, weighted by the technological change elasticities of gender-specific earnings. Given the competing income and substitution effects of changes in female earnings (when they bear most of the time cost of childcare), we may expect the elasticity of demand for children with respect to female earnings to be smaller than that of men's, i.e.,  $|\epsilon_{n,y^w}| < \epsilon_{n,y^m}$ .<sup>19</sup> However, given that shocks such as fracking are predominately biased in favor of occupations that men typically perform (Kearney and Wilson, 2018), we may expect  $\epsilon_{y^w,\alpha} < \epsilon_{y^m,\alpha}$ , leading to a net increase in fertility as positive income effects dominate negative substitution effects. Moreover, historical experiences whereby technological change induced the expansion of the service sector and/or textile industries, disproportionately improving labor market opportunities for women (e.g., Galor and Weil, 1996), have the property  $\epsilon_{y^w,\alpha} > \epsilon_{y^m,\alpha}$ , allowing for substitution effects to dominate income effects, yielding fertility decline.<sup>20</sup>

In my setting, we do not know what the signs of  $\frac{dy^w(\alpha)}{d\alpha}$  and  $\frac{dy^m(\alpha)}{d\alpha}$  will be. Imposing the economic structure based on my hypothesized effects of the soy technological change, assume that  $\frac{dy^w(\alpha)}{d\alpha} < 0$  and

19. In support of this, Schaller (2016) finds stronger relationship between men's earnings and fertility than that of women's.

20. This is similar to the intuition of Diebolt and Perrin (2013), whereby rises in female earnings opportunities eventually create substitution effects large enough to dominate the positive income effects of technological growth.

$\frac{dy^m(\alpha)}{d\alpha} > 0$ . In other words, the soy technology interacts in the production process to increase men’s earnings while decreasing female earnings. This assumption is empirically verifiable and tested in this paper. With these hypothesized restrictions imposed, the total derivative (6) in this particular case is unambiguously positive. The soy technological change is expected to *increase* fertility. The positive substitution effects from reducing female earnings and lowering the cost of children reinforce the positive income effects from increasing male earnings.

In the next section, I discuss the Brazilian context and how it maps onto this model and discuss the relevance of other channels not explicitly modeled here.

## 3 Brazilian Context

### 3.0.1 The Brazilian Fertility Context

Brazil entered its demographic transition in the mid-20th century. Nationally, fertility reached below replacement levels by 2010. This decline was driven by the stopping of births, rather than spacing or delaying births.<sup>21</sup> These fertility behaviors led the Brazilian fertility decline to be characterized by negligible changes in starting ages of fertility and high rates of adolescent fertility. (Goldani, 2009; Martine, 1996). The existing literature in Brazil suggests that the fertility behavior of those age 20 and older are more likely to adjust in response to changes in the economic environment, a feature I will explore in the empirical analysis.

A key assumption made regarding changing female opportunity and fertility is that women bear most of the time cost of childcare, or  $\tau^w > \tau^m$ . This assumption is satisfied in Brazil, where the majority of childcare is done in the home using the mother’s time (Connelly et al., 1996). Beyond this, this technological shock may influence fertility through other channels relevant to the Brazilian context. For example, it may alter child labor usage, infant mortality, bargaining power, and non-wage incomes (such as land values).<sup>22</sup>

It is plausible that if the GE soy technology also destroys occupations that children perform, this could have an offsetting effect, lowering the demand for children.<sup>23</sup> Moreover, increased economic growth may lead to reductions in infant mortality, lowering fertility. Thus, I explore child labor and infant mortality as an outcome of interest in the analysis.

Due to data constraints, I am unable to empirically distinguish bargaining power and land values. How-

---

21. Many ascribe high rates of sterilizations and Cesarean sections as a key driver of fertility decline. However, such medical procedures are most prevalent among higher income women (Silveira et al., 2019), so they cannot explain most of the decline (Martine, 1996). I further discuss descriptive statistics on sterilizations from the 1996 Demographic and Health Surveys (DHS) in Appendix Section 3.

22. Moreover, while I do not explore this possibility here, if the introduction of new agricultural technologies lowers the price of food, it could in theory lower the (pecuniary) costs of raising children, incentivizing higher fertility. However, as soy is primarily an export good, this channel may not be directly relevant.

23. For instance Rosenzweig (1977) highlights the importance of how children transition from ‘production’ goods to ‘consumption’ goods as societies move out of agriculture.

ever, these channels would likely further reinforce the opportunity cost and income effect channels in this context. Data from the 1996 Demographic Health Survey<sup>24</sup> indicates that, on average, men prefer more children than their wives, suggesting that household bargaining channels are likely present.<sup>25 26</sup> Thus, if soy technological change increases men’s relative earnings, it would likely incentivize higher fertility through bargaining channels and yield the same predictions as the more parsimonious unified household model.

Finally, land values can affect fertility through the bargaining channels mentioned above or through increased income if the land is rented. Results from Table 320 of the 2006 Agricultural Census show that over 70% of farms in Brazil are owned by men. Any increase in land values would likely increase the bargaining power of men, and if that land is rented out, it may be reflected in men’s earnings directly and act through positive income effects, increasing fertility as well.

### 3.1 Women’s Roles in Agriculture

Broadly, work in agriculture is divided along gendered lines (Boserup, 1970; Schultz, 2001). The United Nations Food and Agriculture Organization (FAO) as well as Time Use Surveys across developing economies document the common ways in which women participate in agricultural production. These include weeding and tillage operations, fertilizer application, and harvesting (Grassi et al., 2015).<sup>27</sup> Further, qualitative research in Brazil and reports from the Brazilian Agribusiness Ministry confirm that women disproportionately work in the same tasks - weeding, harvesting, and processing produce - in small family farms and larger enterprises alike. Men are more commonly involved in management, contact with agronomists, and the use of new technologies (Brumer, 2008; Brumer, 2004; Lastarria-Cornhiel, 2017).<sup>28 29</sup>

### 3.2 Legalization of GE Soy

Genetically engineered soy from Monsanto was commercially released in 1996 but was not legalized in Brazil until 2003.<sup>30</sup> While the technology was legalized in 2003, the passage of a Biosafety Bill<sup>31</sup> in 2005

---

24. This is the latest survey round of the DHS in Brazil.

25. Westoff et al. (2010) shows that wanted total fertility for Brazilian women in 1996 was 1.8 births for woman, while for men it was 2.9. Realized fertility was 2.5 births per woman.

26. Previous research has also shown that increased bargaining power for lower educated women in Brazil reduced fertility (Klawon and Tiefenthaler, 2001).

27. The pilot study which tested the methodology for the first time use survey in Brazil was only conducted in 2009.

28. Brumer (2004) notes that farms with more technical advancement saw women focusing almost solely on domestic work. These patterns of gender divisions suggest that gains in earnings from new technologies in agriculture may be disproportionately captured by men.

29. In Appendix 2, I discuss Brazilian agricultural farm structures and labor markets. There is wide regional heterogeneity in farm structures across Brazil. While large farms (greater than 100 hectares) are many in number, Brazilian agriculture is still dominated by smaller farms that are less than 10 hectares, many of which are family farms.

30. GE soy was among a new class of agricultural technologies unveiled in the mid-1990s. Innovations in the ability to manipulate plant DNA led to the rise of new biotechnologies referred to as GE seeds. The most prominent of these technologies include herbicide tolerant crops, such as the GE soy studied here and insect repellent crops.

31. Law no. 11.105.

was the turning point for GE technologies in Brazil. The bill created a formal framework for the approval, sale, and use of GE crops. I discuss more details in Appendix Section 2.

This adoption time line made Brazil an early leader in GE crop technologies. In 2010, Brazil had the second highest global area of GE crops (only behind the U.S.), accounting for approximately 17% of world GE crop production (James, 2011). From 2003-2009, GE crops generated \$3.5 billion of farm income in Brazil. The main technology adopted across Brazil was Monsanto’s Roundup Ready herbicide resistant soy seeds, which constituted approximately 70% of all GE crops grown in Brazil that year (James, 2011).<sup>32</sup> In 2018, Brazil overtook the U.S. as the largest producer of soy, with around 95% of its soy crop being GE varieties (Cattelan and Dall’Agnol, 2018).

### 3.2.1 A Primer on GE Soy

A large component of traditional soy cultivation is weed management, as weeds compete for the nutrients, water, and sunlight required for plant growth. Further, weeds inhibit harvesting by becoming entangled in machinery and preventing efficient harvests. Traditional cultivation techniques include manual weed control, where laborers with small, hand-held equipment identify and remove weeds from the field (Bentham, 2013). Prior to the planting of the soy crop, farmers used to undergo a laborious tillage process that included identifying and removing such weeds. The main innovation of the GE soy technology is its natural resistance to glyphosate, a powerful herbicide that kills nearly all crops. Farmers could now simply spray glyphosate to eliminate all weeds without affecting the soy crop, effectively obviating the need of tillage and weeding operations in soy production across Brazil.<sup>33</sup> Thus, this technology directly eliminates female dominated occupations.<sup>34</sup>

## 4 Data and Empirical Trends

I obtain most of the data for this paper from four sources. First, I obtain data on earnings, populations, and socio-economic variables from the Brazilian Population Census published by the Instituto Brasileiro de

---

32. The 2005 framework also led to adoptions of insect resistant (Bt) maize which, along with increased mechanization beginning years earlier, allowed for the growing of maize crop in two seasons within the same year. Bt maize amounted to approximately 28% of all GE crops in Brazil in both the summer and winter seasons in 2010 (James, 2011). Maize may be grown in the same regions following the harvesting of soy (Cattelan and Dall’Agnol, 2018). Thus I control for the expansion of maize technologies in my empirical analysis.

33. Herbicides, specifically glyphosate, have been in production since the 1970s and have been used as an effective weed remover in weed management. It acted by inhibiting the production of enzymes essential for protein synthesis. This affects nearly all crops, including the crops of interest. Roundup Ready GE soy was created by introducing the genes from the bacteria *Agrobacterium*, which exhibited a natural resistance to glyphosate, to the soy crops (Funke et al., 2006). This equipped these crops with immunity to the herbicide.

34. In my discussions with Brazilian agronomists, a commonly cited benefit of GE soy was the ability to spray herbicides again at the end of the crop cycle, facilitating the use of machinery to more cleanly harvest the soy crop. This may then increase the demand for male work.

Geografia e Estatística (IBGE).<sup>35</sup> I obtain data on live births to construct birth rates from the Brazilian Vital Statistics Database (SINASC). I take inter-census population projections compiled from DATASUS (Departamento de Informática do Sistema Único de Saúde).<sup>36</sup> Finally, I obtain data on potential GE soy yields from the FAO Global Agro-Ecological Zones database (FAO-GAEZ).<sup>37</sup>

The relevant geographic units for this study include Federative Units (which I refer to as ‘states’), microregions, and municipalities. There are 26 states in Brazil. There were about 5500 municipalities in 2000. Since municipality borders change over time, I collapse municipality-level data to minimally comparable areas (AMCs) as suggested by the IBGE, which provide consistent geographic boundaries over long periods of time. AMC level variables are population weighted averages of municipality-level data. There are 4,260 AMCs in total. Six AMCs are dropped due to data availability for key variables, leaving the estimation sample with 4254 AMCs across the 26 states. I refer to AMCs as municipalities throughout the paper.

## 4.1 Data from the Census

My main labor market variables come from the micro-data from the sample supplement of the 2000 and 2010 Population Census.<sup>38</sup> The survey asks about formal and informal economic activity for each individual in the household, including if and how much they earned in the reference month for that work.<sup>39 40</sup>

My main labor market variables are on the monthly earnings and the share of women working in a given sector. I average individual level data up to municipality level by sector and gender for the analysis. I use sector codes to define three broad sectors – agriculture, manufacturing, and services. I additionally collect controls derived from extracts of the 1991 Census, including the ratio of the illiteracy rate for women vs men, the log population density, log income per capita, percent of the population that is rural, and the percent of children living in a household where the per capita household income is less than half the minimum wage. I discuss more detail about the construction of these variables in Appendix Section 1.

35. Data can be obtained from <https://www.ibge.gov.br/>. For accessing Census data, I use code provided by Datazoom, developed by the Department of Economics at PUC-Rio to help researchers access IBGE’s household surveys. More information can be found at: <http://www.econ.puc-rio.br/datazoom/english/>.

36. Data can be obtained from <https://datasus.saude.gov.br/informacoes-de-saude-tabnet/>.

37. Maps for potential yields can be obtained from <https://www.gaez.iiasa.ac.at/> after registering for a free account and brought to csv files using GIS software.

38. The way in which economic activity was measured changed in the 2000 Census, making it difficult to create comparable measures in previous waves.

39. In cases where work was seasonal in nature in agriculture, it is asked how much they earned in a typical month during the duration of that seasonal work. For self-employed workers, earnings are either reported as direct earnings, or calculated as the share of total profits they received from the enterprise.

40. In 2000, about 12% of women in agricultural were workers with formal contracts in their main occupation, 20% were workers without a formal contract, about 17% of female workers were classified as ‘own account’ workers, and the remaining 50% were not remunerated in their main line of work. The corresponding numbers for male agricultural workers were 18%, 28%, 31% and 22%, respectively. Note, many workers who are not paid in the work in which they devoted the most of their time may also have other work that they performed for pay. The 50% then is likely an *overestimate* of the total number of female workers solely working in unpaid work in agriculture.

## 4.2 SINASC

DATASUS provides administrative records on vital statistics from SINASC, a data system from the Brazilian Health Ministry. SINASC compiles data from live birth certificates and provides the data at the municipality level.<sup>41</sup> It provides information regarding the mother such as age, race, as well as information on characteristics of the birth (such as birth weight). I use live births by municipality residence. When compiling data on live births, there were approximately 58 municipality codes (about 1% of the total number of municipalities) that could not be matched to AMCs given the correspondence from the IBGE, and I drop these when constructing birth rates. Live births are taken from 1997-2019. I obtain data on inter-census population counts for the denominator of birth rates from DATASUS, an administrative database from Brazil’s Ministry of Health. I provide more details regarding these counts and inter-census projections in Appendix Section 1.

## 4.3 Potential Agricultural Yields

I obtain data on potential yields for soy from the FAO-GAEZ database, which the UN’s Food and Agriculture Organization and the International Institute for Applied Systems Analysis maintain jointly. The potential yields are a function of weather and soil characteristics rather than actual realized yields. The yields are measured in tons per hectare, representing total production capacity for each crop. Importantly, potential yields are calculated under the assumption of different technology usage. Low input regimes are calculated with traditional cultivar techniques, labor intensive techniques, no applications of nutrients, no usage of chemicals for pest and disease control, and minimum conservation measures. The high-level input regime is calculated with improved high yielding varieties, full mechanization with low labor intensity, and the “optimum applications of nutrients and chemical pest, disease, and weed control.” These aim to capture potential production capacity from new technologies. I use baseline geo-climatic conditions calculated as an average of conditions from the 1961-1990 period, estimated prior to the legalization of the technologies.<sup>42</sup>

I define my treatment as the difference in potential yields between the high and low regimes. In other terms:  $\Delta \text{Pot. Soy} = \text{Potential Yields}^{high} - \text{Potential Yields}^{low}$ . This differential yield is interpreted as the change in potential tons per hectare of soy between the high and low input regime.<sup>43</sup>

---

41. SINASC offers an extensive coverage of births, even exceeding coverage from local civil registries in Brazil (Lima et al., 2006). These data are estimated to have less than 4% under-reporting (Marteleto et al., 2020).

42. Documentation for different input levels and user guides can be found at [https://www.iiasa.ac.at/web/home/research/researchPrograms/water/GAEZ\\_v4.html](https://www.iiasa.ac.at/web/home/research/researchPrograms/water/GAEZ_v4.html), more information panel (under the GAEZ V3.0 GLOBAL AGRO-ECOLOGICAL ZONES box), and User’s Guide.

43. Brazil’s soy production in 2000/2001 was estimated at about 2.7 tons per hectare. Table 1 panel E shows the mean value of the potential yield differential from high to low technology in soy is 1.8 tons per hectare. Information on yields for different crops and years can be seen from the United States Department of Agriculture Foreign Agricultural Service briefs at [https://ipad.fas.usda.gov/cropeexplorer/pecad\\_stories.aspx?regionid=br&ftype=prodbriefs](https://ipad.fas.usda.gov/cropeexplorer/pecad_stories.aspx?regionid=br&ftype=prodbriefs).

## 4.4 Summary Statistics

Table 1, Panel A shows the birth rate, defined as the number of live births per 1000 women aged 10-49.<sup>44</sup> In 2000, the average birth rate across municipalities was about 55 births per 1000 women aged 10-49, which declined by about 11 births per 1000, or a 20% decline, by 2010.<sup>45</sup>

Panels B-C of Table 1 show the trends in average monthly earnings from 2000-2010 in 2010 Brazilian reals. These panels show persistent and large gender wage gaps across the Brazilian economy.<sup>46</sup> Row 3 of Panel B shows that in the agricultural sector across municipalities, women made on average 51% of the earnings of men in agriculture in 2000.<sup>47</sup> However, Column 3 shows all sectors experienced proportionally larger increases in women's earnings than that of men's over the decade.

Figure 1 Panel A shows the average of municipality (N=4254) sector employment shares in Brazil in 2000 and 2010. Sector shares are defined as the fraction of all workers (i.e. both men and women) aged 15-55 working in that particular sector.<sup>48</sup> There is a decline in the average employment share in agriculture by about 19%, and an increase in the average manufacturing and service employment shares of about 9% and 2% , respectively, over this time period. Figure 1 also shows the employment share in light industry, a subset of manufacturing, which includes textile and leather goods manufacturing.<sup>49</sup> Light industries made approximately 3% of the employment share in 2000 and grew by about 8% over the decade. Overall, Figure 1 Panel A suggests that industries that employ women, such as services and light manufacturing, may not be expanding rapidly enough to absorb rapid reductions of employment in agriculture.

Figure 1 Panel B then examines the female employment share for each sector, which is defined as the share of total employment age 15-55 in that particular industry that are women. In terms of the three broad sectors, women make up the majority of the employment share in services. In 2000, women's employment share in light manufacturing was about 72% and grew slightly over this period (1.6%). Overall, this decade saw increases in the female labor shares across all broad categories.

---

44. This age range reflects the Health Ministry's definition of fertile age. The young age ranges are included due to the relatively high and stable levels of adolescent fertility in Brazil.

45. Note, defining fertility as the number of births per 1000 women 15-44, which is the more common General Fertility Rate used in other contexts, the average municipality birth rate was about 75 in 2000, and declined to 59 by 2010. This is about equal to the overall U.S. General Fertility Rate in 2018.

46. These persist in a context of educational parity between men and women. In 2019, Brazil was ranked 95th out of 162 countries in terms of gender inequality, and policy makers identify gender inequalities as critical problems to tackle in the Brazilian context. For example, see: <https://www.worldbank.org/en/news/feature/2017/03/08/ser-mujer-brasil>.

47. Using the 2010 exchange rate with the U.S. Dollar (which is 1.75 in August, 2010, according to FRED (series DEXBZUS)), average municipality earnings for women in Brazilian Agriculture are around \$160 per month in 2010 dollars, and for men, \$346 per month.

48. This includes those working in jobs for no remuneration.

49. This industry is particularly relevant as it heavily employs women. Historical experiences saw the expansions of these types of industries improve labor market opportunities for women. In Brazil, these industries also heavily employ women, ranging from 37% of total employment in weaving and fiber processing to 86% of workers in the manufacture of clothing items.

## 5 Empirical Strategy

### 5.1 Variation in Agricultural Technological Change

The main challenge in identifying the impact of agricultural technologies on gender-specific outcomes and fertility is the endogeneity of technological adoption. For instance, areas with higher productivity workers or different cultural norms regarding women’s work may experience differential technological adoption and labor market and fertility trends. Areas with higher fertility preferences may also have lower labor force participation of women and thus relatively higher female wages, which may incentivize the adoption of new labor-saving technologies. Additionally, increases in productivity in other sectors may increase the opportunity cost of labor in agriculture, inducing the adoption of new labor-saving technologies as well as affecting agricultural employment and fertility. Finally, these same concerns may also suggest that systematic measurement error, such as misreporting of earnings, may likely be correlated with actual adoption. To overcome these empirical challenges, I exploit plausibly exogenous variation in the adoption of these technologies.

For variation across space, I use the difference in potential yields from high to low technology regimes, effectively capturing the favorability of weather and soil characteristics for the adoption of these new technologies. For variation over time, I use the legalization of GE soy in 2003.

Other studies have used the FAO-GAEZ potential yields as a source of spatial variation, such as Gollin et al. (2021) and Bustos et al. (2016), the latter of which forms the basis for the identification used here. Due to the nature of my dependent variable, I require a different identifying assumption than that of Bustos et al. (2016), who require potential yields to be exogenous with respect to developments in the industrial sector. I require that the timing of adoption and potential yields are exogenous with respect to fertility decisions and labor market conditions by gender. While this is *a priori* reasonably satisfied, I nevertheless implement a stronger specification than that of the previous studies. Since municipalities may be more similar along unobservables within states, I include state-year fixed effects in addition to municipality fixed effects, which then identify the fertility and gendered market effects from the variability in soy technological change over time and across municipalities within states. This is important given the fact that changes in state level policies (e.g. Bolsa Familia), regional differences in norms and culture, cross state migration, or recessionary effects (e.g. Buckles et al. (2021)) may differentially influence trends in the outcomes.<sup>50</sup> To invalidate my

---

50. The presence of smuggling of GE seeds could raise concerns for identification. There were reports of GE soy seed smuggling in Brazil as early as 2001, with much of the smuggling taking place in the southern state of Rio Grande do Sul (Bentham, 2013). In the main baseline specifications, the pre-year data is created from the 2000 Brazilian Demographic Census Survey data, which predates the legalization as well as the reports of the smuggling of GE soy. It could be possible that smugglers moved to areas with higher geographic suitability or that higher suitability areas were areas more likely to smuggle for other unobserved reasons. More plausibly, between state migration may occur as farmers move south where it is easier to smuggle seeds from bordering Argentina. These could constitute threats to identification. The latter possibility can be adjusted for with state-year fixed effects. The former possibilities can be verified in the presence of differential pre-trends, which I can examine for demographic outcomes.

design, these factors would have to vary systematically among higher and lower soil suitability municipalities within a state-year. I additionally run specifications where I control for covariates allowing for differential trends in municipalities with heterogeneous baseline characteristics.

Figure 2 shows the spatial variation in the potential soy yields measure. In my specifications, I will also include potential maize, defined analogously to that of potential soy, to control for the simultaneous adoption of other technologies as mentioned in Section 3.2. With this in mind, Figure 2 plots the variation in the potential soy measure linearly uncorrelated with potential maize.<sup>51</sup> The dark bold outlines identify states—the inclusion of state-year fixed effects ensures that I am comparing municipalities within states.

A first order question is whether potential yields actually predict the actual adoptions of GE soy. I can confirm this relationship with data from the Agricultural Census. The 2006 and 2017 Agricultural Census contain data on the share of land harvested with GE soy. There are two considerations to take into account before using this data. First, the timing of these Census waves unfortunately do not correspond with the timing of the Demographic Census. Second, the Agricultural Census uses different reference dates for each wave, so each survey captures different crop cycles. Variables created from each wave then are not always comparable over time. With this in mind, I estimate two separate cross-sectional regressions of the share of land harvested with GE soy on potential yields for each year.

Appendix Table 2 reports coefficients of these regressions. The dependent variable is the share of land harvested with GE soy. All specifications include all the baseline controls listed in Table 1 and state fixed effects, representing partial correlations using the identifying variation. These confirm that municipalities *within states* with higher increases in potential soy yields have larger adoptions of GE soy. An increase in one ton per hectare of potential soy increases the share of the total harvested area of all crops reaped with GE soy by 1.5 percentage points in 2006 and by 3.8 percentage points in 2017.<sup>52 53</sup>

## 5.2 Difference-in-Differences for Decadal Labor Market and Fertility Outcomes

My baseline model uses outcomes constructed from the 2000 and 2010 sample supplement survey from the Brazilian Population Census. Let  $y_{m,t}$  be the outcome of interest (for example, the birth rate) in municipality

---

51. This is done by estimating  $Potential\ Soy_m = \beta_0 + \beta_1 Potential\ Maize_m + \epsilon_m$ , where  $m$  indexes municipalities, and plotting the residual. There is a high degree of correlation between these two measures ( $\rho = .79$ ). I add back the mean of potential soy for scaling.

52. The associated F statistic on the soy coefficient is 11.88 for the 2006 regression, and 20.29 for 2017. Moreover, with the above caveats in mind, simply pooling the 2006 and 2017 data and estimating a first difference model between these waves finds that a one ton per hectare increase in potential soy leads to a 1.9 percentage point increase (p-value: 0.007) in the share reaped with GE soy.

53. One may wonder about spatial concentration of soy in Brazil. The results from the 2006 Agricultural Census (table 824) show that while GE soy was planted more heavily in the traditional soy producing regions of the south and centerwest, all regions of Brazil adopted transgenic soy, allowing for sufficient within-state variation in adoption.

m at year t (in state s). I begin with equation:

$$y_{mt} = \alpha_m + \mu_{s,t} + \kappa \text{Potential Soy}_m * t + X'_{m,t} \delta + \psi_{m,t} \quad (8)$$

Where  $\alpha_m$  and  $\mu_{s,t}$  are municipality fixed effects and state-time fixed effects, respectively. The presence of state time fixed effects allows for different states to have arbitrarily different time evolutions in y. Potential Soy is the measure of predicted potential yields from moving from low to high technology regimes. I interact this measure with time. The parameter  $\kappa$  here has an intention-to-treat (ITT) interpretation.  $X_{m,t}$  is a vector of controls, which includes municipality baseline characteristics from the 1991 census (seen in Panel D of Table 1) interacted with time trends, allowing for municipalities with differential baseline characteristics to have different trends. It also includes measures of technical change in maize constructed analogously to that of potential soy, to control for the simultaneous adoption of other agricultural technologies over this time period. Since this equation is estimated with two time periods, I take a within municipality first difference to adjust for any time invariant heterogeneity across municipalities, purging the municipality fixed effect and yielding equation (9):

$$\Delta y_m = \Delta \mu_s + \kappa \text{Potential Soy}_m + X'_m \delta + \Delta \psi_m \quad (9)$$

Which would identify  $\kappa$  from the within-state cross-municipality variation over time. Baseline specification standard errors are all clustered at the microregion level to allow for spatial correlation.<sup>54</sup> The identifying assumption is that conditional on municipality and state-year fixed effects and baseline municipality characteristics, potential yields based on soil and weather characteristics are uncorrelated with other unobserved determinants of labor market outcomes by gender or fertility.<sup>55</sup>

### 5.3 Event Study for Fertility Outcomes

Labor market outcomes constructed from Census data can only be estimated with specification (9). However, using annual fertility data from 1997-2019, and population counts and inter-census population estimates from DATASUS, I also can estimate:

$$y_{m,t} = \lambda_m + \tau_{s,t} + \text{Potential Soy}_m^{\text{Med}} * \mathbb{1}\{t \geq 2003\} \sigma + X'_{m,t} \beta + \eta_{m,t} \quad (10)$$

where  $\lambda_m$  and  $\tau_{s,t}$  are municipality and state-time fixed effects, respectively, and potential yields are inter-

---

54. Microregions are groups of geographically contiguous municipalities created by the IBGE for statistical purposes, analogous to commuting zones in the United States. There are 554 microregions used in the regression sample.

55. The measures of earnings would likely contain measurement error. Consistent estimation of causal effects is still possible as long this error is not correlated with potential yields.

acted with a post-legalization indicator equal to unity for 2003 onwards, and baseline controls are interacted with a time trend.

Furthermore, since the existence of differential pre-trends in fertility by higher and lower suitability municipalities would invalidate the use of low suitability areas as a control, I replace the post indicator with a vector of year indicators, creating an event-study specification. This allows me to examine pre-trends and trace out dynamic effects. Hence, I estimate the following equation:

$$y_{m,t} = \lambda_m + \tau_{s,t} + \sum_{k \neq 2002} Potential\ Soy_m^{Med} * \mathbb{1}\{t = k\} \sigma_k + X'_{m,t} \beta + \eta_{m,t} \quad (11)$$

For interpretation, I present results from estimating (11) replacing the continuous measure of potential yields with an indicator which equals one if the increase in potential yields from a high to low technology regime for a given municipality is above the median increase across municipalities. This is indicated by the superscript “Med”. I omit the interaction in 2002, prior to the legalization of GE crops. All reported standard errors allow for arbitrary correlation between municipalities over time within micro-regions. This specification allows for testing threats to identification for fertility outcomes, namely whether fertility trends in higher and lower potential yield municipalities evolved differently prior to the legalization of GE crop technologies. I present figures plotting the estimates of  $\sigma_k^{Med}$ , which trace out effect of soy technical change on fertility in above median municipalities compared to below in a given year  $k$  relative to the base year 2002 (the year prior the legalization of GE soy).

## 6 Results

### 6.1 Earnings Effects by Gender

Table 2 presents the results of estimating equation (9) with gender-specific measures of agricultural earnings. Columns 1, 2, and 3 use the log of the average of women’s earnings in agriculture, the average of men’s earnings in agriculture, and the log of the average of the relative earnings of women to men in agriculture,<sup>56</sup> respectively, as the dependent variables.<sup>57</sup>

The estimates in Table 2 all confirm and establish that soy technological change in Brazil is gender biased, and more specifically, that it led to large reductions in female earnings. An increase in one ton per hectare of potential soy leads to a statistically and economically significant 11.5% reduction in women’s earnings in

<sup>56</sup> Relative earnings may be of theoretical interest. For example, Siegel (2017) argues the gender wage gap drives trends in U.S. fertility and female employment, as it changes the division of labor within the household.

<sup>57</sup> There are municipalities where agricultural earnings for women is zero in a given year. Given that relative earnings are mostly less than one, the inverse hyperbolic sine is no longer a close approximation to  $\log(2x)$ , so all dependent variables are defined as shifted log transforms for consistency across regressions. All results shown here are robust to either transformation.

agriculture, a 3.3% increase (p-value=.066) in men’s earnings in agriculture, and ultimately a 10% decline in women’s relative earnings in agriculture.

Column 4 shows the effect on overall household earnings, across all sectors, calculated from the household level survey from the Population Census. The coefficient on the soy shock confirms that municipalities with higher GE soy adoption within states experienced an overall increase in household earnings, consistent with the massive productivity gains from these agricultural technologies at the aggregate level. Families and municipalities are getting richer overall. However, the complementarities of these technologies with occupations that different genders perform in agriculture yield heterogeneous effects at the group level.

These tables present the ITT effects. Another way to interpret these results is to think of the effects scaled by the ‘first stage’, or the effect of increases in potential soy yields on actual GE soy adoption. As mentioned earlier, the timing of the Demographic and Agricultural Censuses do not correspond with one another. However, a linear interpolation between the two regression coefficients for 2006 and 2017 from Appendix Table 2 imply a coefficient in 2010 of .023. In other words, an increase of one ton per hectare of potential soy leads to a 2.3 percentage point increase in the share of land reaped with GE soy by 2010. Scaling the coefficients in Table 2 by this ‘first stage’ estimate suggests that a 1 percentage point increase in the share of land harvested with GE soy leads to a 5% decrease in female agricultural earnings, a 1.4% increase in male agricultural earnings, and about a .9% increase in overall earnings.

One may be concerned with this broader definition of earnings that captures all sources. I construct several different measures of average municipality agricultural earnings. I restrict the age of women (and men) to ‘fertile’ age as defined above from the Brazilian Ministry of Health (ages 10-49) as well as to ages 16-49 to avoid any child labor, and also create measures restricted to only direct labor income. This mitigates any concerns about capturing pension incomes, for instance. Appendix Table 4 shows the results using these measures. Results for women’s earnings are robust, and sometimes larger, when restricting income to the more restricted age range and to only earned labor income. Men’s earnings remain largely the same magnitude, but are less precise.

The gender-specific effects on other sectors of the economy from these agricultural shocks are not theoretically clear. The overall impact on gender-specific manufacturing sector earnings, for instance, depends on the production technologies in the manufacturing sector, particularly on the elasticity of labor demand, and on the mobility of labor across sectors. In response to the negative demand shock from soy in the agricultural sector, we may expect to see an *overall* reduction in the earnings of other sectors if labor reallocates into those sectors.<sup>58</sup> However, there is no clear mechanism ex-ante to expect a strong *gender-specific* effect in any direction. With these considerations in mind, I also estimate equation (9) with gender-specific earnings in

---

58. Indeed, this is empirically confirmed in Bustos et al. (2016).

manufacturing and services. Coefficients and 95% confidence intervals are presented in Appendix Figure 1. I find no statistically significant effects on these sectors. Focusing on just the overall (all sectors) municipality ratio of women’s earnings to men’s earnings, I find a negatively signed, but imprecisely estimated effect.

It may also be the case that there are individuals who live in one municipality and work in another. If some workers living in lower potential yield municipalities commute to higher potential yield municipalities, the higher treatment intensity effect would spill over into recorded earnings in the lower treatment places. If some workers live in higher yield municipalities and work in lower yield municipalities, then this would reflect in lack of movements in earnings in higher treatment intensity regions. In both cases, one would expect these to attenuate the effects.

One may also be concerned with measurement error in earnings affecting estimates. Idiosyncratic measurement errors are partly mitigated from averaging individual earnings to the municipality level, and since earnings are the dependent variable in regressions, then idiosyncratic error would reflect in higher error variance. Any problematic measurement error would have to be systematically correlated with potential soy yields within states to constitute a threat to identification. One potential source may be systematic misreporting. While the census attempts to ask each individual about their earnings, the requirement is at least one person answers for the interview. One possible concern is that men may be more likely to answer the census survey and systematically misreport their spouse’s earnings or share of profits of the enterprise. In light of this concern, I re-estimate (9) using earnings only from female-headed households, and find similarly sized reductions in female earnings.<sup>59</sup>

## 6.2 Gender-Specific Sector Shares

Here, I estimate (9) with the share of paid workers who are female in each sector as the dependent variable. While the shares of women aged 10-49 would be consistent with the age range relevant for those making fertility choices in Brazil, I restrict the shares here to those 16 and older, as the legal minimum age of work in Brazil in this time period is 16 and I will examine child labor as a separate outcome of interest in Section 6.5. Thus, I define the dependent variable as the share of women out of total paid employment (i.e., the denominator includes both men and women) aged 16-49 in each sector. Results presented here are robust to using the broader age ranges.

Table 3 Panel A presents these estimated effects. An increase in one ton per hectare of potential soy leads to a decrease in women’s paid share in agriculture by .72 percentage points (p-value=.078), a 6.5% decrease off the baseline 2000 mean. In absolute value, this effect is a sizable 9% of the overall (positive)

---

59. Other work in Brazil, such as Dix-Carneiro and Kovak (2017), reassuringly finds similar results when estimating earnings regression with Census data and from administrative data sources.

change in female remunerated agricultural shares from 2000 to 2010.

Again, scaling these ITT estimates by the crude first stage estimate as in the previous section suggests that a one percentage point increase in the share of land reaped with GE soy leads to a .3 percentage point reduction in the share of women working in paid agricultural work. This represents a 2.8% decrease off the baseline mean. This negative effect size on female employment is roughly equal to a 4-percentage point increase in agricultural mechanization in India as estimated by Afridi et al. (2022).<sup>60</sup>

Panel B examines female agricultural employment in more detail. The dependent variable in column 1 is the share of overall female employment, including both paid and unpaid work, in agriculture (i.e., the denominator is the total number of female workers). Column 2 then limits this measure to only paid workers. Column 3 uses the share of women in agriculture working for unpaid work (i.e., the denominator is all female workers in paid and unpaid work in agriculture), which is defined as either helping another household member with their work or working in cultivation for a household's sustenance for no pay. And finally, Column 4 uses hours in the main occupation within agriculture.<sup>61</sup>

Column 1 shows that there is no significant movement of women out of the agricultural sector overall, suggesting women are staying within the agricultural sector.<sup>62</sup> Limiting this measure to only paid workers in column 3, there is a clear reduction in the share of paid female workers in agricultural work. Column 4 then examines movements from paid to unpaid work. The point estimate in Column 3 indicates that an increase in one ton per hectare of potential soy leads to a 4.7 percentage point increase in women working in unpaid work. Finally, using hours in the main occupation within the agricultural sector for women as the dependent variable, I find statistically and economically insignificant point estimates. Overall, the employment results suggest a negative demand shock for women's work in agriculture, where displaced women reallocate into unpaid agricultural work. The null effects on hours also provide evidence against a case where women are consuming more leisure.

The coefficients in Panel A from other sectors of the economy are statistically insignificant, but negative in sign. These combined with the negative (but also statistically insignificant) coefficient on male manufacturing earnings (Appendix Figure 1) could be consistent with a story where men who lost their jobs from the

---

60. The larger magnitude can be explained by the mechanization studied in Afridi et al. (2022) occurring in different stages of the production process, which they find to have downstream effects lowering the need for weeding work. In the case studied here, the soy technological change directly affects weeding tasks.

61. The main occupation refers to the occupation in which the individual devoted the most hours to in the reference period. The survey in 2000 asks about hours in the main occupation as well as other occupations, however in the 2010, it only asked about hours in the main work.

62. Appendix Table 3 show female re-allocations across the other sectors of the economy, including both paid and unpaid work. The dependent variables are the share of women working in a given sector (i.e., the denominator is the total number of female workers). Columns 2 and 3 do suggest a movement of women out of services into manufacturing, albeit, economically small in magnitude. A Wald test of whether the potential soy coefficients across models offset allowing for non-zero covariances between estimators fails to reject that they offset (p-value=.5). This suggests that if there are any expansions in overall manufacturing occupations for women arising from structural transformation induced by this shock, it is only women who were previously in services who are able to take these jobs.

technological change were more able to move into other sectors of the economy such as manufacturing, i.e. a movement down the demand curve.<sup>63</sup>

One may worry about the changing composition of the workforce as people migrate driving these labor market results. I present evidence on the lack of differential migration by gender in Section 6.5. Moreover, the Census survey asks how long an individual has lived continuously in their current municipality. As a robustness exercise, I restrict the sample in the two Census years to only those who have lived in the same municipality for 10 or more years and rerun the employment and earnings results. I find that the effects for women are maintained (and are larger in magnitude) with this restriction. For men, the coefficient, while still positive, is statistically insignificant. These results are reported in Appendix Table 5.

### 6.3 Fertility: Difference-in-Differences and Event Study

The labor market effects presented above confirm the hypothesized restrictions from Section 2.0.2 that  $\frac{dy^w(\alpha)}{d\alpha} < 0$  and  $\frac{dy^m(\alpha)}{d\alpha} > 0$ , suggesting an unambiguous increase in fertility from soy technical change. Table 4 presents the results of estimating equation (9) with the birth rate defined as the number of live births per 1000 women aged 10-49 as the dependent variable. The coefficient shows that an increase in one ton per hectare of potential soy leads to an *increase* of 2.3 births per 1000 women. This result coupled with the labor market effects from the previous section fully sign and confirm the prediction of equation (6).

For comparisons to studies in the U.S. context, I redefine the birth rate as the number of births per 1000 women aged 20-39. I find that an increase in one ton per hectare of potential soy leads to an increase of about 5.3 births per 1000 women, or a 4.6% increase off a baseline mean of 115 births per 1000 women. This is roughly two times Autor et al. (2019)’s effect size from a one unit negative trade shock from China to female-intensive employment in the US of 2.3%.<sup>64</sup> It is slightly smaller than the magnitude of the positive income effect associated with a \$1000 increase in simulated fracking production from the fracking boom, a positive predominantly male-biased shock, which led to a 5.9% increase in fertility (Kearney and Wilson, 2018). Moreover, it is also smaller in magnitude than a 10% increase in earnings associated with the Appalachian coal boom, another primarily male-driven earnings shock, which led to a 7% increase in the birth rate (Black et al., 2013). The effects I find here are large and economically significant. The fact that these effect sizes are smaller in magnitude than primarily male dominated sectoral shocks is in accordance with economic theory. Negative income effects from the lower female earnings, in addition to attenuating gains in total family earnings, would dilute the incentives to increase fertility compared to a large primarily male dominant shock, which would work through pure positive income effects. In line with these differential magnitudes,

---

63. Consistent with this, using the shares of men across sectors as the dependent variable analogously to that in Panel B, there are statistically significant reallocations of men out of agriculture into manufacturing.

64. Their coefficient is 2 off a baseline mean of 87 births per 1000 women.

Schaller (2016) finds that improved labor market opportunities for men increase fertility, whereas increases in female labor demand negatively impact fertility, but by smaller magnitudes than that of the estimated male effects.

Appendix Table 10 shows the results of estimating equation (9) using measures of fertility directly from the Census. Increases in potential soy lead to an increase in the number of completed births to women, as well as the number of children under 5 in a municipality, consistent with the estimates using birth rates from the administrative records.<sup>65</sup>

I now exploit the longer time series data available for fertility. This gives the advantage of allowing for visual inspection of pre-trends and tracing out any dynamic effects from the adoption of these new technologies. For instance, it is plausible that this could be a short-term increase in fertility that eventually recedes as opposed to a sustained increase in fertility. Tracing out the dynamic effects of the shock can distinguish between these possibilities.

Figure 3 shows the result of the estimation of equation (11) for potential soy. Recall that these estimates trace the evolution of effects in municipalities with above median compared to below median increases in potential yields. I omit 2002, the year prior to the legalization of GE crop technologies. The figure shows the effect in a given year relative to the base year in the time period of 1997-2019. It contains a balanced panel of 4,254 municipalities with 97,842 observations. As before, standard errors are clustered at the microregion level to allow for arbitrary correlation of errors across municipalities and over time within microregions. The shaded bands represent 95% confidence intervals.

We can see there is an absence of problematic pre-trends, supporting the identifying assumptions.<sup>66</sup> Examining the dynamic effects following the legalization, there is a statistically significant and discernible effect on fertility starting around 3 years after the legalization of GE crops, and it continues to grow over this time period. The initial lag following the legalization is sensible in that fertility choices are not realized immediately. Moreover, recall in Section 3.2 that while Brazil legalized the first plantings of GE soy in 2003, it was not until 2005 where a full regulatory body formed to facilitate the commercialization of the genetically modified technology. Thus, the ascent in fertility beginning roughly around 2005 is expected.

Moreover, the increase in fertility is sustained up to 17 years following the legalization of the technology. Areas with above median increases in potential soy had 2 more births per 1000 women compared to below median municipalities in 2019 relative to the year before legalization. This also rules out the possibility of a re-timing of births, as the initial increase in births are not compensated by a decrease in births later. Given

---

65. Further, it may be possible that as municipalities get richer, they are able to register more births, rather than increasing actual fertility. This is likely not the case as the vital statistic data is comprehensive in the time period studied here, and the increases in fertility in the sample survey to the Census also mitigate such concerns about administrative records.

66. The F statistic for joint significance of pre- years is .79.

the backdrop of fertility decline across this period in Brazil, these results suggest that the soy technological change is generating these effects by stagnating the decline in fertility in high potential yield municipalities.

In order to think about the magnitudes of these effects, I turn to the 2017 Agricultural Census data to directly instrument for actual GE soy adoption in the fertility regressions. Estimating a specification akin to (10), where I instrument for the share of farmland harvested with GE soy with the potential yields measure,<sup>67</sup> I find that a one standard deviation, or 11 percentage point, increase in the share of farmland harvested with GE soy increased fertility by 9.5 births per 1000 women.<sup>68</sup> For comparability to more commonly used measures, I redefine the birth rate as the number of births per 1000 people, and find that a one standard deviation increase in the share of land harvested with GE soy leads to an increase of about 3.5 births per 1000 people. This is approximately equal to the total decline in fertility in sub-Saharan Africa between 2005 and 2017.<sup>69</sup> This is particularly relevant as this is a region where these technologies are likely to be adopted next and where policy makers are concerned about stalled fertility transitions.<sup>70</sup>

Appendix Figure 2 shows that the results remain robust to weighting by the population of women aged 10-49 as well as using the log number births as the dependent variable (while also controlling for the log population of women of fertile age). Finally, a recent literature documents the shortcomings of two-way fixed effects specifications. Estimating the annual specification (10) directly yields an increase in fertility of 1.4 births per 1000 women. While the timing of adoption in this setting is not staggered, De Chaisemartin and d’Haultfoeulle (2020) shows that the underlying weighted sum of average treatment effects may still contain negative weights even in this setting. Reassuringly, I find that only 7.5% of weights are negative, and the sum of the negative weights is equal to -.018.<sup>71</sup>

## 6.4 Age-Specific Fertility Effects

As discussed in Section 3.0.1, the Brazilian fertility experience is characterized by relatively high rates of adolescent fertility and a young fertility schedule. Moreover, Brazil’s fertility decline has been driven primarily by the stopping of births, rather than delaying first births. Consistent with these regularities documented by a large demographic literature on Brazil, La Ferrara et al. (2012), for example, find no movements in fertility in response to telenovelas in Brazil for younger age groups, but quantitatively large declines for the age groups 25-34 and 35-44. Thus, one would expect ex ante the older age groups to be more

67. I impute 0 for the pre year for GE soy, as it was not yet available in Brazil at that time.

68. The associated Kleibergen-Paap F-Statistic is 19.5.

69. The birth rate fell from 40 to 36 births per 1000 people over this time period. *Source:* World Bank World Development Indicators, Id:SP.DYN.CBRT.IN

70. Other factors such as child labor usage in weeding tasks could also mediate these effects in other developing countries.

71. Appendix Section 3 also explores the fact that certain regions of Brazil share similar agricultural structures to that of other developing economies, such as large amounts of smaller family farms, while other regions have primarily large farms with more hired labor. The effects are similar across regions, suggesting these results are not driven by peculiarities of farm structures in Brazil.

responsive in their fertility choices.

I test for this possibility by creating age-specific birth rates based on the age of the mother available from live birth certificates from SINASC. First, I construct the adolescent birth rate as the number of live births per 1000 women aged 10-19. I then construct the remaining birth rates in 10-year increments. Figure 4 plots the estimated coefficients on the potential soy measure in equation (11). Panel A uses adolescent fertility as the dependent variable, and Panel B and C use age-specific birth rates for the age ranges 20-29 and 30-39, respectively.<sup>72</sup> Panel A shows, consistent with the demographic and economic literature in Brazil, that there is no discernible impact of the soy technology shock on adolescent fertility.

Panels B and C show that all of the movement in fertility is driven by older age groups, consistent with La Ferrara et al. (2012). Municipalities with an above median increase in potential soy experience sizable and persistent increases in fertility at all age ranges compared to below median municipalities. Panel B shows that in 2019, above median soy potential municipalities had 5 more births per 1000 women aged 20-29 compared to below median municipalities relative to the base year, which amounts to 5% increase over the 2002 mean. Equivalently, Panel C shows in 2019 a 7.4% increase in above median compared to below median municipalities.<sup>73</sup>

The fertility dynamics presented in Figure 3 and Figure 4 speak more to the longer run effects of this technological change. The sustained increase in older age-specific rates shown in Figure 4 is consistent with a situation in which women who lose work opportunities from the adoption of soy technologies have more births throughout their lifetimes. For instance, consider if women in high potential yield municipalities increased fertility while in the 20-29 age range, but then no had more births afterwards. This would generate no changes (or possibly negative effects if it reflects a re-timing of births) in the age-specific birth rates for the older ages in later years after the legalization as these women enter into those older age ranges. However, for the 30-39 age group, we see sustained and still increasing birth rates in the latest years in the sample. These also are suggestive of longer run labor market effects. Unfortunately, the 2020 Demographic Census has been delayed due to Covid-19, so I cannot examine gendered labor market results on a longer time horizon. But the fertility dynamics do shed some light on these effects. If it were the case that the shock only temporarily displaced female workers in 2010, and these workers found employment in the years post 2010, we would not expect to see sustained increases in fertility 17 years following the legalization.

---

72. The coefficients using the birth rate per 1000 women aged 40-49 are positive and statistically significant, but very small in magnitude.

73. Note that estimating earnings regressions by these age groups shows similar magnitudes of declines for those 20-29 and 30-39 for women, but larger increases for men in the 30-39 age group. This could drive the slightly higher magnitudes of the 30-39 age range in 2010.

## 6.5 Ruling Out Other Channels

In addition to the opportunity cost and income effect channels, a large literature has explored other incentives for fertility choices. This section explores the following channels: child labor, infant mortality, and differential migration.

### 6.5.1 Child Labor

The returns from child labor, derived directly from their earnings or indirectly from participating in income generating activities, offset the pecuniary price of having children.<sup>74</sup> If children are also performing similar tasks in agriculture as women, it is possible that the technological adoption could be child labor-saving as well, providing competing incentives to lower fertility.

I obtain two measures of child labor. First, I obtain measures of the child labor rate in a municipality in 2000 and 2010, computed by DATASUS from the Brazilian Demographic Census. This represents the proportion of children aged 10-15 working or looking for work. Additionally, from the 2000 and 2010 Demographic Censuses, I directly compute the share of all agricultural workers between the ages 10-55 who are aged 10-15 working in agriculture. I choose 16 as the cut off age, as this is the legally defined minimum age of work in Brazil during this period. The mean child labor rate and share of agricultural workers who are children are 17% and 6.6% in the year 2000, respectively.

Columns 1 and 2 of Table 5 report the results of estimating equation (9) with these two measures of child labor as the dependent variable. I find no statistically significant effects on these two variables. The point estimates are also economically small in magnitude. Taken together, these results provide evidence that changes in child labor across high and low potential soy municipalities are not playing a significant role in driving these fertility outcomes.

### 6.5.2 Infant Mortality

These agricultural technologies contributed to large earnings and productivity gains in Brazilian agriculture. Increased development may decrease child mortality (Bharadwaj et al., 2020), which may in turn affect fertility decisions (Guinnane, 2011; Ager et al., 2018). Moreover, if municipalities with larger expansions in GE soy experienced differential infant mortality trends prior to the legalization of GE crops, that could call into question the exogeneity of the potential yields, suggesting that perhaps some other unobservables that affect infant mortality can also be influencing fertility, as they share similar determinants (Schultz, 1997).

---

74. Doepke (2004) finds that child labor laws played a key role in generating the demographic transition. Previous work has examined the impact of transitory economic shocks on child labor in modern developing countries. For instance, Beegle et al. (2006) finds that child labor works as a consumption smoothing mechanism in response to negative income shocks. And in Brazil specifically, Kruger (2007) shows that child labor increases in response to transitory earnings opportunities for children.

Or they could be indicative of an infant mortality transition itself acting as a potential driver of the fertility change.

To test for this possibility, I take data from the Mortality Information System of Brazil, a system managed by the Department of Health Situation Analysis which compiles information from death certificates.<sup>75</sup> I use data from 1998-2019.<sup>76</sup> I define the infant mortality rate as the number of deaths of children below age one per 1000 live births. The mean infant mortality rate in Brazil in 2000 stood at 23.5.<sup>77</sup> Panel A of Figure 5 shows the results of estimating equation (11) from 1998 to 2019, omitting the year 2002, with the infant mortality rate as the dependent variable. There is a clear absence of differential trends prior to the legalization of GE soy, and no effects in the post period.

### 6.5.3 Migration

Here, I consider the possibility that the adoption of these technologies induced by favorable soil and weather characteristics differentially affected migration choices of men and women across municipalities. The presence of such responses would change the composition of the municipalities and could alter interpretations of the estimates above.<sup>78</sup> To test for this possibility, I turn back to the 2000 and 2010 Demographic Censuses. In the sample questionnaires for both years, respondents answered how many complete years of residence they have lived without interruption in the city of current residence. From these questions, I create a measure of recent migrants as the log number of people in a municipality who have lived there for less than 10 years. I do this for all ages and also restrict the measure to the more relevant age group of 10-49. Columns 3-6 of Table 5 show the results of estimating the equation (9) with these outcomes. I find no significant effects on any of these measures.<sup>79</sup>

## 6.6 Child Quality

As municipalities are overall getting richer from this shock, we may expect improvements in educational and health outcomes for children. These outcomes are common measurements of the ‘quality’ of children. However, the quantity-quality framework links quality investment in children as jointly determined with the number, or quantity, of children (Becker 1960). When parents have more children the cost of quality increases, as they must invest more in order for each child to achieve a given level of quality. Thus the ‘price’ of child quality is an increasing function of the number of children. On the other hand, increases in the level

75. This data can be obtained from DATASUS.

76. While there are earlier years, a large number of municipalities have missing values/incomplete records before this year.

77. This is about the infant mortality rate of the United States in 1965.

78. For instance, Wanamaker (2012) finds the introduction of textile mills in South Carolina in the late 19th century led to lower fertility, but these effects were driven by the rising opportunity costs of children for migrating households who were separated from their families.

79. Recall the inclusion of state-year fixed effects adjusts for any interstate migration.

of quality increase the price of children, as each additional child requires more investments in order to achieve that higher level of quality. This creates a non-linearity in the budget constraint of the household. *Ceteris paribus* decreases in the price of children would induce substitution away from quality and towards quantity - suggesting that this soy shock may induce offsetting incentives for lower investments in child quality.<sup>80</sup>

I first turn to measures of child quality from the Demographic Census. I look at three measures of educational investment: the fraction of children aged 5 to 15 who are literate, the fraction attending school (including pre-school), and the fraction of those aged 18-24 who have obtained at least a high school degree. Moreover, following Ponczek and Souza (2012), I look at other measures of child quality that may be relevant for developing economies, such as the fraction of children who are helping other members of the household with work for no pay and the fraction who work in cultivation and vegetation for household sustenance. These forms of child labor take time and effort that children cannot use towards school and would serve as indicators of lower child quality investment.<sup>81</sup>

Table 6 reports estimates from equation (9) using these measures of quality as dependent variables. I find statistically insignificant and economically small point estimates. Taken together, these estimates suggest that there was no discernible change in child quality, at least by these measurements. The lack of effects for some of these outcomes could be specific to the Brazilian setting. For instance, the fraction of those who are attending school is high at .89 in 2000. For the household labor results, these could also suggest that any incentives to increase investments in child quality from growing earnings are offset by the increase in the number of children through the quantity-quality trade off.

Another type of child quality investment is infant health. We may expect that increased agricultural productivity and overall earnings provide more readily accessible food and nutrition. This would potentially improve children's health outcomes. To the extent that parents can exert influence over child health outcomes (and that these are time intensive and/or monetarily costly), it is also possible to see a reduction on these measures as parents substitute away from quality towards quantity.<sup>82</sup>

To explore infant health, I also obtain data from SINASC to construct the fraction of all births with low birth weights, defined as birth weights less than 2500 grams, from 1997 to 2019. Low birth weight is not a health outcome expected more than others to be impacted by these changes. Rather, I choose this measure because of both data availability and the epidemiological interest in the drivers of Brazil's negligible improvement in reducing low birth weights over the past few decades (Silveira et al., 2019).

---

80. Thus the nature of this shock, which decreases the price of children directly, yields different predictions on child quality than that of Foster and Rosenzweig (1996) or Gehrke and Kubitz (2021), where the shocks studied induced increases in the returns to education, either directly or indirectly.

81. Note that these measures intersect with the child labor measures in Section 6.5.1.

82. Further, through household bargaining channels, we may expect to see family resources shifted away from health investments in children. For instance, Atkin (2009) finds that women induced to work in new factories have stronger bargaining power within their household and have taller children.

Panel B of Figure 5 plots the estimated coefficients of equation (11) with the fraction of births that are low birth weight as the dependent variable. Bands represent 95% confidence intervals. Again, there are no significant pre-trends, as well as no effects on this outcome following the adoption of these technologies. While this does not rule out any changes in infant health, it does suggest that on an important margin where we may expect to see changes, there are none. This is consistent with offsetting effects from the channels mentioned above.<sup>83</sup>

## 7 Other Robustness

In Appendix Section 3, I report additional robustness checks. First in regard to inference, I report standard errors clustered at a larger level of aggregation than the Microregion - the Mesoregion, which were defined by the IBGE to group together regions with similar characteristics. There are 113 Mesoregions in the sample, meeting standard thresholds for asymptotics. Moreover, I report Conley (1999) standard errors for the Census sample using distance cutoffs of 50 and 200 km.<sup>84</sup> Appendix Table 6 reports these standard errors and show only men’s agricultural earnings and the paid agricultural shares are sensitive to clustering at the Mesoregion and using the 200km cutoff Spatial-HAC standard errors. Further, Appendix Table 7 reports the main estimates weighted by population in 2000. All results are similar with these weights. In Appendix Table 8, I add additional controls for market access. Since soy is often transported via trucks and trains, I estimate the distance from a municipality’s center to existing highways as of 2000 and railways using shapefiles from the Brazilian Ministry of Infrastructure. The coefficients are not sensitive to the addition of these controls. I also report the implied coefficients accounting for omitted variables under the assumption that additional unobservables are equally related to the treatment as the observables (Oster, 2019). And finally, I relax the assumption that the municipality approximates the local markets by aggregating the main variables to a larger level of observation, the microregion. Re-estimating the main results and clustering at the Mesoregion level, all point estimates remain the same sign, however only the female agricultural earnings and fertility results remain significant, with a decrease in female earnings of 7%, and an increase in fertility of about 2.7 births per 1000 women.

---

83. Given that male fetuses are more vulnerable to adverse conditions in utero, I also split infant mortality and low birth weight by gender. I find no significant changes in either gender.

84. To compute Conley standard errors, I use code from Hsiang (2010).

## 8 Discussion

I established that this agricultural technological change is gender-biased, and specifically that it served as a negative demand shock to women’s work. And consistent with economic models of fertility, this female labor-saving technological change led to increases in fertility. In this Section, I discuss how we can interpret these results. Specifically, I address potential reasons why women were not able to move into other sectors of the economy and what these findings may mean for their welfare.

### 8.1 Why are Women Not Allocating into Other Sectors?

My labor market results show that in response to the new soy technology, women were not reallocating into a different sector of employment but instead were reallocating into unpaid agricultural work. A natural question arises as to why women were not able to move into other sectors of the economy – such as services. One possible explanation is the difference in educational and skill requirements across sectors in Brazil. Alvarez (2020) finds that, in Brazil, the substantially lower earnings in the agricultural sector are more likely due to sorting, whereby lower- skilled workers select into agriculture. Female agricultural workers, who tend to have low levels of education (Helfand et al. (2015)), may not have the ability to allocate into service sector occupations. Further, recall that overall household earnings were *increasing*. Thus, there may have been no incentive to invest in more skills and move to other geographic locations (such as cities) for alternative employment, as opposed to reallocating into home production or unpaid work for another household member’s more profitable enterprise.<sup>85</sup>

This explanation is also consistent with Bustos et al. (2019), who find that the expansion of soy technologies in Brazil primarily expanded *lower-skilled* manufacturing sectors. The results from Table 3 and Appendix Figure 1 (as discussed in Section 6) provide evidence, albeit imprecisely estimated, of a reallocation of men into manufacturing and of lower male earnings in manufacturing. This is suggestive of men in agriculture affected by these technologies being able to move into low-skilled manufacturing sectors that may be more physically intensive and suited for male workers.<sup>86</sup>

---

85. Additionally, this time period in Brazil was characterized by lower returns to education (Ferreira et al., 2017), which would further disincentivize investments in human capital.

86. Further, Bustos et al. (2016) find no evidence of local spillovers expanding service sector employment from the soy shock, suggesting that this case is different than that of the expansion of palm oil in Indonesia, in which spillovers from the increased profits expanded local service sector employment (Gehrke and Kubitza, 2021). If there were local spillovers in terms of increased demand for manufactured goods, such as automobiles, this could increase male labor demand in manufacturing. Coupled with an increase in the supply of any male workers who may have been “pushed” out of agriculture, the new equilibrium manufacturing employment would be higher, but the equilibrium earnings are ambiguous. This outcome is consistent with the results in Table 3 and Appendix Figure 1.

## 8.2 Welfare

As the soy shock generates winners and losers within the same household, the welfare effects are theoretically ambiguous. If taking a pure income perspective, households as a unit are better off. However, this shock works directly against two major development goals: improving economic conditions for women and lowering fertility (Bank, 2015).<sup>87</sup> As noted earlier, Brazil ranks low in measures of gender equality that include reproductive rights. Data from the DHS reveals that women prefer fewer children than their husbands and have relatively high numbers of unwanted births. Given this context, this technological shock may lower welfare within the household.<sup>88</sup>

There is also evidence from literature outside of economics that corroborates a negative welfare interpretation of such effects. Kramer and McMillan (2006) find that in small-scale subsistence societies without much access to labor market opportunities, labor-saving technologies led to women having higher completed fertility, as families substitute from production to reproduction. Paris and Chi (2005) show that plastic row seeder technologies eliminated weeding jobs in Vietnam, which were performed primarily by women, leading to increased time in childcare. In interviews, lower-skilled women expressed subjectively worse conditions, as they lacked alternative employment opportunities, faced increases in debt, and lost independent streams of income.

The significantly reduced labor costs for weeding operations extend not only to the soy studied here but to other herbicide tolerant crops such as maize and cotton. All of these new technologies offer massive potential productivity and development gains. In this light, policy makers and economists are actively promoting new GE crops for economies in Africa that did not experience the Green Revolution (Pehu and Ragasa, 2008; Carter et al., 2021). By 2010, South Africa, Burkina Faso, and Egypt became leaders in their respective regions of Africa by legalizing and field-testing new GE crops. Many policy makers note the “strategic importance” that their example plays in encouraging other countries across the continent to adopt these technologies (James, 2011). Many of these countries have stalled fertility transitions and larger agricultural shares than that of Brazil. My results suggest that policy makers should be cautious about incentivizing and promoting agricultural technological change without also adopting measures to counter potential reductions in female reproductive rights and autonomy.

---

87. Moreover, these goals are intrinsically related due to the link between high fertility rates and low female reproductive rights and autonomy. For example, see: <https://www.worldbank.org/en/news/feature/2013/06/14/invest-in-fertility-decline-to-boost-development-in-pakistan>.

88. This argument is more salient if bargaining channels are present, whereby the technological change lowers the bargaining power of women. However, even assuming the unitary household model from Section 2.0.1, the welfare implications are ambiguous, and could still be negative. Taking the full Lagrangian value function, and, for simplicity, assuming that only women bear the time cost of child care ( $\tau^m = 0$ ), the envelope theorem reveals that the net effect of this specific technological change is only unambiguously positive when  $n^* \tau^w > 1$ .

## 9 Conclusion

Technological change is a key driver of growth, but it has not provided opportunity for all workers. When particular demographic groups systematically fill certain occupations, technological changes' interactions with those occupations can disproportionately alter the earnings opportunities for those groups. While our understanding of these unequal effects is growing for advanced economies, we know little about the impacts of technological change in modern developing economies. Here, the technological frontier is in agriculture, a sector where occupations are divided along gendered lines.

My paper focuses on the bias of agricultural technological change in Brazil. I show that surprising results can arise in this less considered but increasingly relevant context. I first provide new evidence of how agricultural technological change that promotes overall growth can disproportionately lower labor market opportunities for women, exacerbating existing gender inequalities. I then show that when technological change eliminates female economic opportunities, it can lead to economically meaningful increases in fertility. These results run contrary to historical experiences of development as documented widely in the literature, whereby fertility declines when economies grow due to the acceleration of technological progress. However, by embedding technological change into a Beckerian framework relating gender-specific earnings to fertility, I show that these increases in fertility are fully anticipated through the same economic channels which, in the past, led to fertility decline.

## References

- Acemoglu, Daron, and David Autor.** 2011. "Skills, tasks and technologies: Implications for employment and earnings." In *Handbook of Labor Economics*, 4:1043–1171. Elsevier.
- Afridi, Farzana, Monisankar Bishnu, and Kanika Mahajan.** 2022. "Gender and mechanization: Evidence from Indian agriculture." *American Journal of Agricultural Economics*.
- Ager, Philipp, Benedikt Herz, and Markus Brueckner.** 2020. "Structural change and the fertility transition." *The Review of Economics and Statistics* 102 (4): 806–822.
- Ager, Philipp, Casper Worm Hansen, and Peter Sandholt Jensen.** 2018. "Fertility and early-life mortality: Evidence from smallpox vaccination in Sweden." *Journal of the European Economic Association* 16 (2): 487–521.
- Alesina, Alberto, Paola Giuliano, and Nathan Nunn.** 2013. "On the origins of gender roles: Women and the plough." *The Quarterly Journal of Economics* 128 (2): 469–530.
- Alvarez-Cuadrado, Francisco, and Markus Poschke.** 2011. "Structural change out of agriculture: Labor push versus labor pull." *American Economic Journal: Macroeconomics* 3 (3): 127–58.
- Alvarez, Jorge A.** 2020. "The agricultural wage gap: Evidence from Brazilian micro-data." *American Economic Journal: Macroeconomics* 12 (1): 153–73.
- Anelli, Massimo, Osea Giuntella, and Luca Stella.** 2019. "Robots, labor markets, and family behavior." *Forthcoming, Journal of Human Resources*.
- Arias, Diego, Pedro Abel Vieira, Elisio Contini, Barbara Farinelli, and Michael Morris.** 2017. "Agriculture productivity growth in Brazil: Recent trends and future prospects." Brazil Productivity Growth Flagship Report, <https://openknowledge.worldbank.org/handle/10986/29437>.
- Atkin, David.** 2009. "Working for the future: Female factory work and child health in Mexico." <http://economics.mit.edu/files/6727>, *Unpublished Manuscript, Yale University*.
- Autor, David, David Dorn, and Gordon Hanson.** 2019. "When work disappears: Manufacturing decline and the falling marriage market value of young men." *American Economic Review: Insights* 1 (2): 161–78.
- Bank, World.** 2015. *Global monitoring report 2015/2016: Development goals in an era of demographic change*.
- Beaudry, Paul, and Ethan Lewis.** 2014. "Do male-female wage differentials reflect differences in the return to skill? Cross-city evidence from 1980-2000." *American Economic Journal: Applied Economics* 6 (2): 178–94.
- Becker, Gary S.** 1960. "An economic analysis of fertility." In *Demographic and Economic Change in Developed Countries*, 209–240. Columbia University Press.
- Becker, Gary S.** 1992. "Fertility and the Economy." *Journal of Population Economics* 5 (3): 185–201.
- Becker, Gary S, and H Gregg Lewis.** 1973. "On the Interaction between the Quantity and Quality of Children." *Journal of Political Economy* 81 (2, Part 2): S279–S288.
- Beegle, Kathleen, Rajeev H Dehejia, and Roberta Gatti.** 2006. "Child labor and agricultural shocks." *Journal of Development Economics* 81 (1): 80–96.
- Benthem, Leander van.** 2013. "Soybean and maize production in Brazil." PhD diss., M. Sc. thesis), Wageningen University & Research Center Department of Business Economics.
- Bharadwaj, Prashant, James Fenske, Namrata Kala, and Rinchan Ali Mirza.** 2020. "The Green revolution and infant mortality in India." *Journal of Health Economics* 71:102314.
- Black, Dan A, Natalia Kolesnikova, Seth G Sanders, and Lowell J Taylor.** 2013. "Are children "normal"?" *The Review of Economics and Statistics* 95 (1): 21–33.

- Black, Dan A, Terra G McKinnish, and Seth G Sanders.** 2003. "Does the availability of high-wage jobs for low-skilled men affect welfare expenditures? Evidence from shocks to the steel and coal industries." *Journal of Public Economics* 87 (9-10): 1921–1942.
- Black, Sandra E, and Alexandra Spitz-Oener.** 2010. "Explaining women's success: technological change and the skill content of women's work." *The Review of Economics and Statistics* 92 (1): 187–194.
- Boserup, Ester.** 1970. *Woman's Role in Economic Development*. Earthscan from Routledge; republished 2007.
- Brueckner, Markus, and Hannes Schwandt.** 2015. "Income and population growth." *The Economic Journal* 125 (589): 1653–1676.
- Brumer, Anita.** 2004. "Gênero e agricultura: a situação da mulher na agricultura do Rio Grande do Sul." *Revista Estudos Feministas* 12:205–227.
- Brumer, Anita.** 2008. "Gender relations in family-farm agriculture and rural-urban migration in Brazil." *Latin American Perspectives* 35 (6): 11–28.
- Buckles, Kasey, Daniel Hungerman, and Steven Lugauer.** 2021. "Is fertility a leading economic indicator?" *The Economic Journal* 131 (634): 541–565.
- Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli.** 2016. "Agricultural productivity and structural transformation: Evidence from Brazil." *American Economic Review* 106 (6): 1320–65.
- Bustos, Paula, Juan Manuel Castro-Vincenzi, Joan Monras, and Jacopo Ponticelli.** 2019. "Industrialization without Innovation." NBER WP No. 25871, <https://www.nber.org/papers/w25871>.
- Carney, Conor, and Monica Harber Carney.** 2018. "Impact of soil conservation adoption on intra-household allocations in Zambia." *Review of Development Economics* 22 (4): 1390–1408.
- Carranza, Eliana.** 2014. "Soil endowments, female labor force participation, and the demographic deficit of women in India." *American Economic Journal: Applied Economics* 6 (4): 197–225.
- Carter, Michael, Rachid Laajaj, and Dean Yang.** 2021. "Subsidies and the african green revolution: Direct effects and social network spillovers of randomized input subsidies in mozambique." *American Economic Journal: Applied Economics* 13 (2): 206–29.
- Cattelan, AJ, and A Dall'Agnol.** 2018. "The rapid soybean growth in Brazil." *OCL-Oilseeds and Fats, Crops and Lipids* 25 (1).
- Conley, Timothy G.** 1999. "GMM estimation with cross sectional dependence." *Journal of Econometrics* 92 (1): 1–45.
- Connelly, Rachel, Deborah S DeGraff, and Deborah Levison.** 1996. "Women's employment and child care in Brazil." *Economic Development and Cultural Change* 44 (3): 619–656.
- De Chaisemartin, Clément, and Xavier d'Haultfoeuille.** 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American Economic Review* 110 (9): 2964–96.
- Dettling, Lisa J, and Melissa S Kearney.** 2014. "House prices and birth rates: The impact of the real estate market on the decision to have a baby." *Journal of Public Economics* 110:82–100.
- Diebolt, Claude, and Faustine Perrin.** 2013. "From stagnation to sustained growth: the role of female empowerment." *American Economic Review* 103 (3): 545–49.
- Dix-Carneiro, Rafael, and Brian K Kovak.** 2017. "Trade liberalization and regional dynamics." *American Economic Review* 107 (10): 2908–46.
- Doepke, Matthias.** 2004. "Accounting for fertility decline during the transition to growth." *Journal of Economic Growth* 9 (3): 347–383.
- Doepke, Matthias.** 2015. "Gary Becker on the quantity and quality of children." *Journal of Demographic Economics* 81 (1): 59–66.

- Doepke, Matthias, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt.** 2022. *The economics of fertility: A new era*. Technical report. National Bureau of Economic Research.
- Ferreira, Francisco HG, Sergio Firpo, and Julián Messina.** 2017. “Ageing poorly?: accounting for the decline in earnings inequality in Brazil, 1995–2012.” *Accounting for the decline in earnings inequality in Brazil* 2012.
- Flaskerud, George K.** 2003. “Brazil’s soybean production and impact.” North Dakota State University, Department of Agribusiness and Applied Economics, *Extension Bulletins*.
- Foster, Andrew D, and Mark R Rosenzweig.** 1996. “Technical change and human-capital returns and investments: evidence from the green revolution.” *The American economic review*: 931–953.
- Funke, Todd, Huijong Han, Martha L Healy-Fried, Markus Fischer, and Ernst Schönbrunn.** 2006. “Molecular basis for the herbicide resistance of Roundup Ready crops.” *Proceedings of the National Academy of Sciences* 103 (35): 13010–13015.
- Galor, Oded.** 2012. “The demographic transition: causes and consequences.” *Cliometrica* 6 (1): 1–28.
- Galor, Oded, and David N Weil.** 1996. “The gender gap, fertility, and growth.” *The American Economic Review* 86 (3): 374.
- Gehrke, Esther, and Christoph Kubitza.** 2021. “Agricultural Productivity and Fertility Rates: Evidence from the Oil Palm Boom in Indonesia.” *Journal of Human Resources*: 0520–10905R1.
- Goldani, Ana Maria.** 2009. “What will happen to Brazilian fertility.” *Completing the Fertility Transition*. New York: United Nations: 331–45.
- Goldin, Claudia.** 1994. *The U-shaped female labor force function in economic development and economic history*.
- Goldin, Claudia.** 2006. “The quiet revolution that transformed women’s employment, education, and family.” *American Economic Review* 96 (2): 1–21.
- Goldin, Claudia, and Kenneth Sokoloff.** 1982. “Women, children, and industrialization in the early republic: Evidence from the manufacturing censuses.” *The Journal of Economic History* 42 (4): 741–774.
- Gollin, Douglas, Casper Worm Hansen, and Asger Mose Wingender.** 2021. “Two blades of grass: The impact of the green revolution.” *Journal of Political Economy* 129 (8): 000–000.
- Grassi, Flavia, Josefine Landberg, and Sophia Huyer.** 2015. “Running out of time: The reduction of women’s work burden in agricultural production.” *Rome: Food and Agricultural Organization of the United Nations*.
- Guanziroli, Carlos, Antonio Buainain, and Alberto Sabbato.** 2013. “Family farming in Brazil: evolution between the 1996 and 2006 agricultural censuses.” *Journal of Peasant Studies* 40 (5): 817–843.
- Guinnane, Timothy W.** 2011. “The historical fertility transition: A guide for economists.” *Journal of Economic Literature* 49 (3): 589–614.
- Heath, Rachel, and A Mushfiq Mobarak.** 2015. “Manufacturing growth and the lives of Bangladeshi women.” *Journal of Development Economics* 115:1–15.
- Helfand, Steven M, Ajax RB Moreira, and Edward W Bresnayan Jr.** 2015. “Agricultural productivity and family farms in Brazil: Creating opportunities and closing gaps.” Sustainable Development Department, *World Bank Report*.
- Hsiang, Solomon M.** 2010. “Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America.” *Proceedings of the National Academy of Sciences* 107 (35): 15367–15372.

- James, Clive.** 2011. "Brief 42: Global Status of Commercialized Biotech/GM Crops: 2010." *ISAAA Brief*, no. 44.
- Jensen, Robert.** 2012. "Do labor market opportunities affect young women's work and family decisions? Experimental evidence from India." *The Quarterly Journal of Economics* 127 (2): 753–792.
- Jones, Larry E, Alice Schoonbroodt, and Michele Tertilt.** 2008. *Fertility theories: can they explain the negative fertility-income relationship?* Technical report. National Bureau of Economic Research.
- Katz, Lawrence F, and Kevin M Murphy.** 1992. "Changes in relative wages, 1963–1987: supply and demand factors." *The Quarterly Journal of Economics* 107 (1): 35–78.
- Kearney, Melissa S, and Riley Wilson.** 2018. "Male earnings, marriageable men, and nonmarital fertility: Evidence from the fracking boom." *The Review of Economics and Statistics* 100 (4): 678–690.
- Kebede, Endale, Anne Goujon, and Wolfgang Lutz.** 2019. "Stalls in Africa's fertility decline partly result from disruptions in female education." *Proceedings of the National Academy of Sciences* 116 (8): 2891–2896.
- Kitchens, Carl T, and Luke P Rodgers.** 2020. "The Impact of the WWI Agricultural Boom and Bust on Female Opportunity Cost and Fertility." NBER WP No.27530, <https://www.nber.org/papers/w27530>.
- Klawon, Emily, and Jill Tiefenthaler.** 2001. "Bargaining over family size: The determinants of fertility in Brazil." *Population Research and Policy Review* 20 (5): 423–440.
- Kramer, KarenL, and GarnettP McMillan.** 2006. "The effect of labor-saving technology on longitudinal fertility changes." *Current Anthropology* 47 (1): 165–172.
- Kruger, Diana I.** 2007. "Coffee production effects on child labor and schooling in rural Brazil." *Journal of Development Economics* 82 (2): 448–463.
- La Ferrara, Eliana, Alberto Chong, and Suzanne Duryea.** 2012. "Soap operas and fertility: Evidence from Brazil." *American Economic Journal: Applied Economics* 4 (4): 1–31.
- Lastarria-Cornhiel, Susana.** 2017. "Feminization of agriculture: Trends and driving forces." <https://openknowledge.worldbank.org/handle/10986/9104> .
- Lima, Claudia Risso de Araujo, José A. Escamilla, Otaliba Libânio de Moraes Neto, and Vinicius Pawlowski Queiroz.** 2006. "Successful Brazilian Experiences in the Field of Health Information Final Report."
- Marteletto, Leticia J, Gilvan Guedes, Raquel Z Coutinho, and Abigail Weitzman.** 2020. "Live births and fertility amid the Zika epidemic in Brazil." *Demography* 57 (3): 843–872.
- Martine, George.** 1996. "Brazil's fertility decline, 1965–95: a fresh look at key factors." *Population and Development Review*: 47–75.
- Moscona, Jacob.** 2019. "Agricultural development and structural change within and across countries." <http://economics.mit.edu/files/13482>, *Unpublished Manuscript*.
- Ngai, L Rachel, and Claudia Olivetti.** 2015. "Structural transformation and the u-shaped female labor supply." In *2015 Meeting Papers*, vol. 1501.
- Nunn, Nathan, and Nancy Qian.** 2011. "The potato's contribution to population and urbanization: evidence from a historical experiment." *The Quarterly Journal of Economics* 126 (2): 593–650.
- Oster, Emily.** 2019. "Unobservable selection and coefficient stability: Theory and evidence." *Journal of Business & Economic Statistics* 37 (2): 187–204.
- Paris, Thelma, and Truong Thi Ngoc Chi.** 2005. "The impact of row seeder technology on women labor: A case study in the Mekong Delta, Vietnam." *Gender, Technology and Development* 9 (2): 157–184.

- Pehu, Eija, and Catherine Ragasa.** 2008. "Agricultural biotechnology: Transgenics in agriculture and their implications for developing countries." World Bank Report, <https://openknowledge.worldbank.org/handle/10986/9042>.
- Ponczek, Vladimir, and Andre Portela Souza.** 2012. "New evidence of the causal effect of family size on child quality in a developing country." *Journal of Human Resources* 47 (1): 64–106.
- Qian, Nancy.** 2008. "Missing women and the price of tea in China: The effect of sex-specific earnings on sex imbalance." *The Quarterly Journal of Economics* 123 (3): 1251–1285.
- Rasul, Imran.** 2008. "Household bargaining over fertility: Theory and evidence from Malaysia." *Journal of Development Economics* 86 (2): 215–241.
- Romer, Paul M.** 1990. "Endogenous technological change." *Journal of Political Economy* 98 (5, Part 2): S71–S102.
- Rosenzweig, Mark R.** 1977. "The demand for children in farm households." *Journal of Political Economy* 85 (1): 123–146.
- Rotella, Elyce J.** 1981. "The transformation of the American office: changes in employment and technology." *The Journal of Economic History* 41 (1): 51–57.
- Schaller, Jessamyn.** 2016. "Booms, Busts, and Fertility: Testing the Becker Model using Gender-Specific Labor Demand." *Journal of Human Resources* 51 (1): 1–29.
- Schultz, T Paul.** 1985. "Changing world prices, women's wages, and the fertility transition: Sweden, 1860–1910." *Journal of Political Economy* 93 (6): 1126–1154.
- Schultz, T Paul.** 1997. "Demand for children in low income countries." *Handbook of Population and Family Economics* 1:349–430.
- Schultz, T Paul.** 2001. "Women's roles in the agricultural household: bargaining and human capital investments." *Handbook of Agricultural Economics* 1:383–456.
- Siegel, Christian.** 2017. "Female relative wages, household specialization and fertility." *Review of Economic Dynamics* 24:152–174.
- Silveira, Mariangela F, Cesar G Victora, Bernardo L Horta, Bruna GC da Silva, Alicia Matijasevich, and Fernando C Barros.** 2019. "Low birthweight and preterm birth: trends and inequalities in four population-based birth cohorts in Pelotas, Brazil, 1982–2015." *International Journal of Epidemiology* 48 (Supplement\_1): i46–i53.
- Silvestrin Zanon, Raquel, and Maria Sylvia Macchione Saes.** 2010. "Soybean production in Brazil: main determinants of property sizes." <https://ageconsearch.umn.edu/record/100476/files/22-Saes-Zanon.pdf>.
- Sofa, T, and Cheryl Doss.** 2011. "The Role of Women in Agriculture: An ESA Working Paper no. 11-02, March 2011." *Agricultural Development Economics Division, the Food and Agriculture Organization of the United Nations, Viale delle Terme di Caracalla, Rome, Italy.*
- Solow, Robert M.** 1956. "A contribution to the theory of economic growth." *The Quarterly Journal of Economics* 70 (1): 65–94.
- Staff, LADB.** 2003. "Brazil Approves Genetically Modified Soy." <https://digitalrepository.unm.edu/notisur/13192>.
- Wanamaker, Marianne H.** 2012. "Industrialization and fertility in the nineteenth century: Evidence from South Carolina." *The Journal of Economic History* 72 (1): 168–196.
- Westoff, Charles F, et al.** 2010. "Desired number of children: 2000–2008." *DHS comparative reports*, no. 25.

Table 1: Summary Statistics

Panel A: Fertility					
Variable (2000)	Mean	SD	2010-2000 Mean	SD	N
Birth Rate (Births per 1000 women aged 10-49)	55.48	16.56	-10.91	15.22	4254
Panel B: Monthly Municipality Agricultural Earnings (in 2010 Reals)					
Variable (2000)	Mean	SD	Mean Log Difference: 2010-2000	SD	N
Female Earnings	281.25	307.7	0.72	0.99	4254
Male Earnings	608.3	485.3	0.24	0.34	4254
Relative Earnings (Female/Male)	0.51	0.43	0.34	0.56	4254
Panel C: Monthly Municipality Earnings and Other Sectors (in 2010 Reals)					
Variable (2000)	Mean	SD	Mean Log Difference: 2010-2000	SD	N
Female Earnings (All Sectors)	517	249.7	0.38	0.29	4254
Female Earnings (Manufacturing)	499.27	392.6	0.5	1.48	4254
Female Earnings (Services)	594.94	226.9	0.37	0.24	4254
Male Earnings (All Sectors)	850.8	503.5	0.2	0.25	4254
Male Earnings (Manufacturing)	865.19	722.7	0.17	0.64	4254
Male Earnings (Services)	1209	574.2	0.17	0.3	4254
Panel D: Baseline Covariates 1991					
Variable (1991)	Mean	SD			N
Share of Rural Population	0.46	0.23			4254
Female to Male Ratio of Literacy Rates	1.1	0.29			4254
Log Population Density	3.2	1.35			4254
Log Income Per Capita	4.52	0.59			4254
% of Children in Low Income Households	87.11	11.61			4254
Potential Maize	3.06	1.81			4254
Panel E: Instrument: Potential Yields					
Variable	Mean	SD			N
Potential Soy	1.8	0.85			4254

Table shows summary statistics where the unit of observation are municipalities in Brazil. Earnings data and covariates are from the sample supplement to the Brazilian Census. Earnings data is in Brazilian Reals, deflated to 2010. Birth rates are calculated as the number of live births per 1000 women aged 10-49, defined as 'fertile' age by the Brazilian Ministry of Health. Data on live births come from live birth certificates compiled by SINASC, and the population counts come from projections provided by DATASUS.

Table 2: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Earnings

VARIABLES	(1) ΔAgr. Female Earnings	(2) Δ Agr. Male Earnings	(3) Δ Agr. Relative Earnings	(4) Δ Family Earnings
ΔPot. Soy	-0.115*** (0.0432)	0.0326* (0.0177)	-0.0999*** (0.0277)	0.0202** (0.00848)
Observations	4,254	4,254	4,254	4,254
R-squared	0.062	0.067	0.126	0.223
Controls	YES	YES	YES	YES
$\bar{Y}$ (2000)	281.25	608.30	.51	1236.92

Robust standard errors clustered at the Microregion level are in parentheses

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table reports first difference estimates using Decennial Census data from 2000 and 2010. In the first column, the dependent variable is the log of the average female municipality agricultural earnings. The second column uses the log of the average male municipality agricultural earnings. The third column uses the log of the ratio of women to men's agricultural earnings. Data are taken at the individual level, and aggregated to the municipality using weights provided by the IBGE. The last column reports the average overall family earnings, taken from the household survey of the demographic census, and includes all sectors. All regressions include the following baseline controls: the share of the population that is rural, the female and male literacy rates, the population density, the log income per capita, and the percent of children living in low income households, all from their levels in 1991. Regressions additionally control for measure of potential yields in maize analogous to that of the potential soy measure, to adjust for the simultaneous expansion of technologies in other crops. Finally, all regressions include municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. The last row reports the mean (in levels) of the dependent variable in the year 2000.

Table 3: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Female Employment Shares

Panel A: Share of Workers in Paid Work who are Female				
VARIABLES	(1) $\Delta$ Agriculture	(2) $\Delta$ Manufacturing	(3) $\Delta$ Services	(4) $\Delta$ Light Industry
$\Delta$ Pot. Soy	-0.00715* (0.00406)	-0.00375 (0.00656)	-0.00184 (0.00229)	-0.00254 (0.0151)
Observations	4,254	4,254	4,254	4,254
R-squared	0.215	0.043	0.123	0.016
Controls	YES	YES	YES	YES
$\bar{Y}$ (2000)	0.11	0.28	0.55	0.70

Panel B: Female Agricultural Employment				
VARIABLES	(1) $\Delta$ Share in Agriculture	(2) $\Delta$ Share in Paid Agriculture	(3) $\Delta$ Unpaid Agriculture	(4) $\Delta$ Agricultural Hours
$\Delta$ Pot. Soy	-0.00569 (0.00477)	-0.0144*** (0.00421)	0.0470*** (0.0142)	-0.420 (0.412)
Observations	4,254	4,254	4,254	4,254
R-squared	0.107	0.170	0.211	0.025
Controls	YES	YES	YES	YES
$\bar{Y}$	0.21	0.12	0.48	37.24

Robust standard errors clustered at the Microregion level are in parentheses

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table reports first difference estimates using Decennial Census data from 2000 and 2010. In Panel A, the dependent variables are the female sector shares i.e. the numerator is the number of women working for pay in a sector, and the denominator is the number of men and women working for pay in that sector. The dependent variable in first column is the share within agriculture, the second column is the share in manufacturing, the third column is the share in services, and the fourth is the share in light industry, a subset of manufacturing that includes textile manufacturing sectors that heavily employ women. Panel B focuses on agriculture. The dependent variable in column 1 is the share of women in agriculture, i.e. the denominator includes only women. Column 2 shows the share of paid female employment that is employed in agriculture, column 3 uses the number of hours worked in the main work within agriculture, and column 4 shows the share of female employment in unpaid agriculture. Column 4 of Panel B is the share of women within agriculture working in unpaid work. Unpaid work includes helping another household member with their work or working in cultivation for a household's sustenance for no pay. The means of the dependent variables in 2000 are reported in the last row of each panel. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state.

Table 4: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Birth Rates per 1000 Women

VARIABLES	(1) $\Delta$ Fertility
$\Delta$ Pot. Soy	2.346*** (0.907)
Observations	4,254
R-squared	0.128
Controls	YES
$\bar{Y}$ (2000)	55.48
Effect size	4.23
Standard errors clustered at the Microregion level	
*** p<0.01, ** p<0.05, * p<0.1	

Table reports first difference estimates where the dependent variable is the number of live births per 1000 women aged 10-49. The mean of the dependent variable in 2000 is reported in the second to last row, and the final row reports the effect size, defined as the coefficient on potential soy divided by the mean in 2000, multiplied by 100. Data on live births are taken from administrative records from SINASC, and population projections from DATASUS. The regression includes all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state.

Table 5: Testing Other Channels that Influence Fertility

	(1) $\Delta$ Child Labor	(2) $\Delta$ Agr. Child Share	(3) $\Delta \log(\text{Female Migrants})$	(4) $\Delta \log(\text{Male Migrants})$	(5) $\Delta \log(\text{Female Migrants}^{10-49})$	(6) $\Delta \log(\text{Male Migrants}^{10-49})$
$\Delta \text{Pot. Soy}$	0.000266 (0.00312)	0.00236 (0.00170)	0.0120 (0.0162)	0.0191 (0.0215)	0.0166 (0.0165)	0.0176 (0.0203)
Observations	4,245	4,254	4,254	4,254	4,254	4,254
R-squared	0.057	0.111	0.052	0.052	0.050	0.047
Controls	YES	YES	YES	YES	YES	YES
$\bar{Y}$ (2000)	0.17	0.07	3489.51	3365.05	2391.35	2247.38

Robust standard errors clustered at the Microregion level are in parentheses

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table shows first differences estimates. Column 1 uses the share of children (10-15) working or seeking work in the reference period as the dependent variable. Column 2 uses the share of all agricultural workers who are 10-15 as the dependent variable. Column 3-6 all use the log of the number of migrants by gender. Migrants are defined as people who have lived in the municipality for less than 10 years (the range between census waves). Columns 3 and 4 include all ages, and columns 5 and 6 include only migrants ages 10-49. All regressions include the full set of controls. Means in the bottom row represent dependent variables in levels in 2000. All regressions include municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. Standard errors are clustered at the Microregion level.

Table 6: First Difference Estimates of Potential Soy on Child Quality Investments

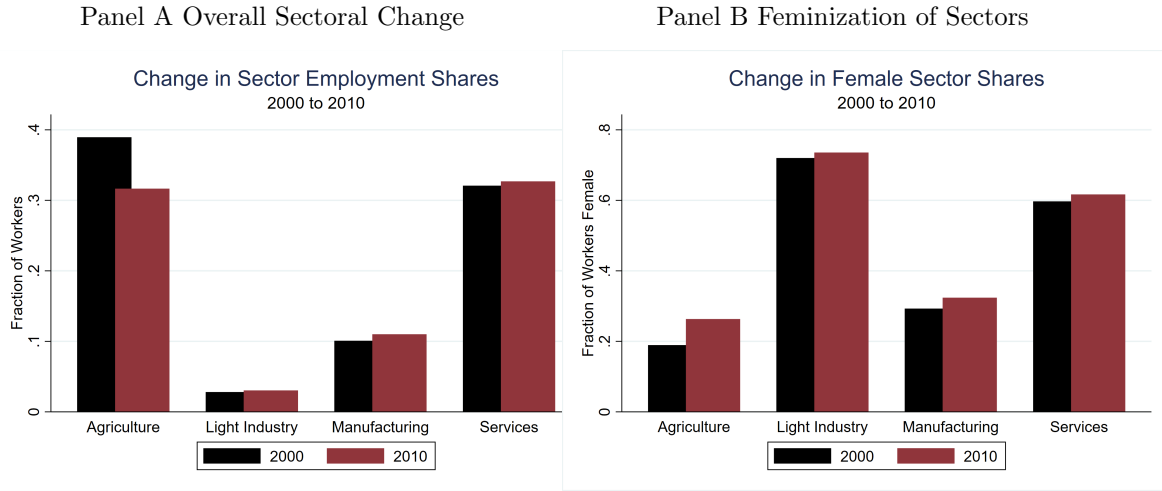
	(1)	(2)	(3)	(4)	(5)
	$\Delta$ Literate	$\Delta$ Attend School	$\Delta$ High School +	$\Delta$ Unpaid Labor	$\Delta$ Cultivation
Pot. Soy	0.00158 (0.00253)	-0.000270 (0.00265)	-0.00156 (0.00391)	0.00186 (0.00132)	-0.000720 (0.000659)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.327	0.266	0.157	0.137	0.138
Controls	ALL	ALL	ALL	ALL	ALL
$\Delta \bar{Y}$	0.10	0.07	0.20	-0.01	0.01
$\bar{Y}_{2000}$	0.75	0.89	0.20	0.04	0.01

Standard Errors Clustered at the Microregion Level

\*\*\* p&lt;0.01, \*\* p&lt;0.05, \* p&lt;0.1

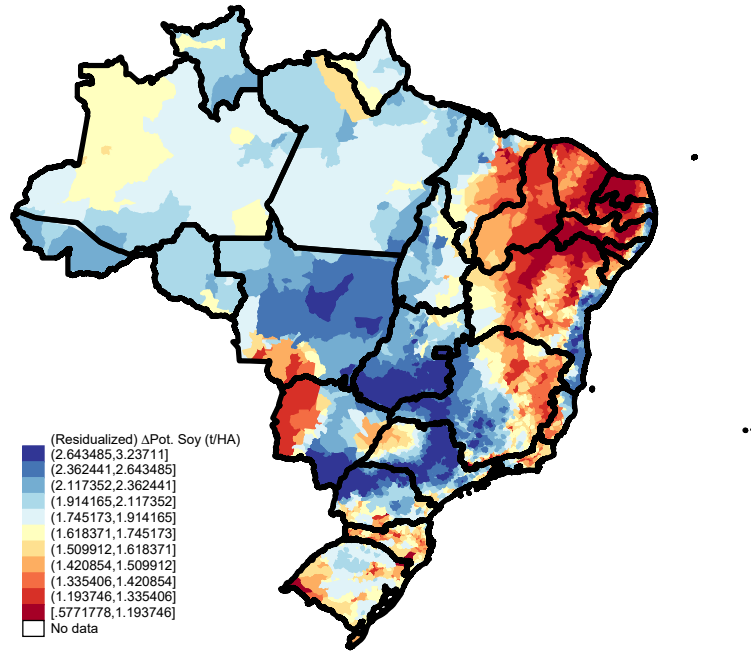
Table shows first difference estimates. Column 1 uses the share of all children 5-15 who are literate as the dependent variable. Column 2 uses the share of children who are attending school, including primary school. Column 3 uses the share of people 18-24 who have completed at least a high school degree. Column 4 uses the share of children 5-15 who participated in unpaid labor for a household member, and column 5 looks at the share who helped a household member without pay in cultivation or vegetation for household sustenance. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. Standard errors are clustered at the Microregion level.

Figure 1: Changes in Employment Across Sectors, 2000 to 2010



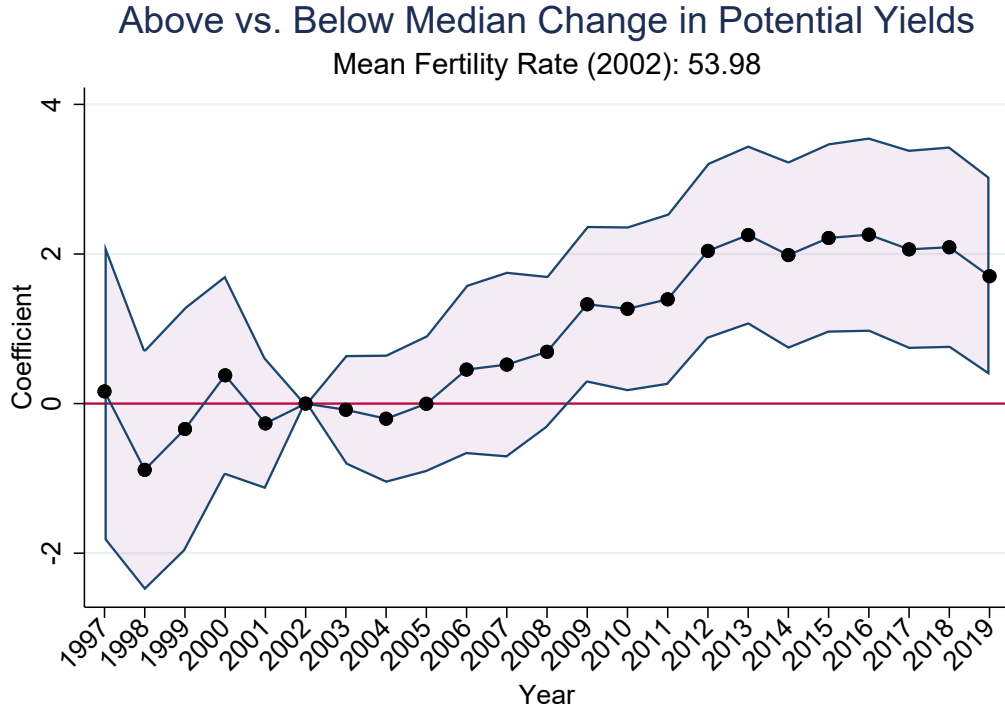
Panel A reports the averages in 2000 and 2010 of the employment shares of all workers across municipalities in a given sector. Panel B shows the share of workers aged 15-55 who are female working in a given sector in a given year. The numerator is the total number of women working in that sector, and the denominator is the total number of workers (men and women). The figures include both paid and unpaid work. Data are taken from the sample supplement of the Demographic Census from 2000 and 2010, and aggregated to the municipality level using weights provided by the IBGE. Each bar represents an average across municipalities (N=4254)

Figure 2: Spatial Suitability for Soy Technological Change



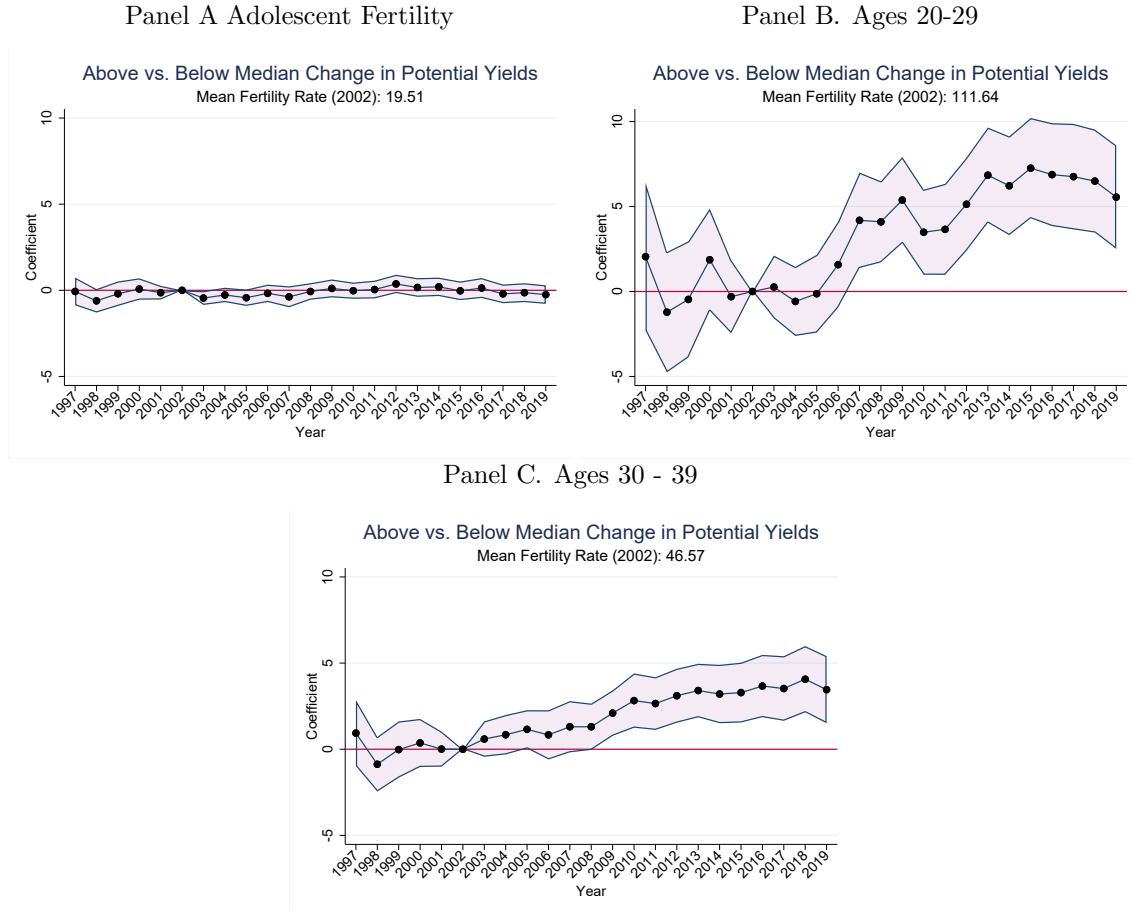
Map shows the measure of soy technological change: the difference in potential yields between high and low technology regimes, from the FAO-GAEZ database, by municipality. Dark bold outlines represent the Brazilian states. Municipality borders are omitted for visualization. Darker shades of blue represent higher potential yields from adopting new technologies, and darker shades of red represent lower potential soy yields. The soy yields residual out potential maize yields to account for the simultaneous expansion of other technologies in maize, and the mean is added back to the measure for interpretation. Potential yields are measured in tons per hectare.

Figure 3: Event Study Estimates of Potential Yields on Birth Rates per 1000 Women



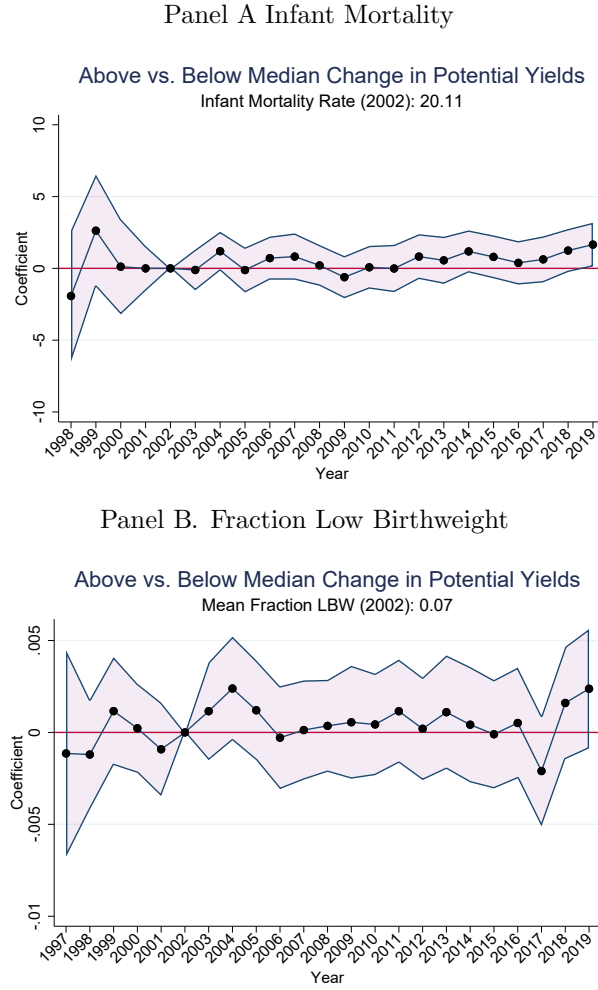
This figure plots estimated coefficients from Equation (11) where the dependent variable is the Birth Rate, defined as the number of live births per 1000 women aged 10-49. The data run from 1997-2019 and have a balanced sample of 4254 municipalities, creating 97,842 observations. The year prior to the legalization of GE soy, 2002, is omitted. The mean birth rate in the omitted year 2002 is presented above the figure. The shaded bands represent 95 percent confidence intervals. The regression includes all baseline controls interacted with a time trend, as well as the measure of potential maize interacted with a vector of year indicators. The regression includes municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. Standard errors are clustered at the Microregion level.

Figure 4: Event Study Estimates of the Effect of Soy Potential Yields on Age-Specific Fertility



These sub-figures plots estimated coefficients from (11). Panel A uses the adolescent birth rate as the dependent variable, defined as the number of live births per 1000 women aged 10-19, and the remaining panels B and C use the birth rate per 1000 women aged 20-29 and 30-39, respectively. The data run from 1997-2019 and have a balanced sample of 4254 municipalities, creating 97,842 observations. The year prior to the legalization of GE soy, 2002, is omitted, and the mean of the dependent variable in this omitted year is given above the figure. The shaded bands represent 95 percent confidence intervals. The regression includes all baseline controls interacted with a time trend, as well as the measure of potential maize interacted with a vector of year indicators. All regressions include municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. Standard errors are clustered at the Microregion level.

Figure 5: Event Study Estimates of the Effect of Soy Potential Yields on Infant Health Outcomes



These sub-figures plots estimated coefficients from (11) where the dependent variable in the first figure is the infant mortality rate, defined as the number of deaths of children less than 1 years old per 1000 live births, and the second is fraction of children with low birth weight, where low birth weight is defined as less than 2500 grams. The data run from 1998-2019 for infant mortality and have a balanced sample of 4254 municipalities, creating 93,585 observations. The fraction low birth weight specification runs from 1997 to 2019 with 97,815 observations. The year prior to the legalization of GE soy, 2002, is omitted. The shaded bands represent 95 percent confidence intervals. The regression includes all baseline controls interacted with a time trend, as well as the measure of potential maize interacted with a vector of year indicators. All regressions include municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. Standard errors are clustered at the Microregion level.

# Appendix

## Section 1: Data Appendix

### Labor Market Variables

Sectors for occupations are taken from the activity codes of the 2000 Census for consistency and applied to the 2010 census wave, since occupation codes change over time. Agricultural work is defined using codes  $\in [1101, 05002]$ , which include Agricultural and aquaculture production. Manufacturing is defined using codes in the interval  $[15010, 37000]$ . Services are defined using codes defined in  $[53000, 99000]$ , which include occupations such as retail trade, housing and hospitality services, financial intermediation, public administration, and education services. Light manufacturing is defined as a subset of manufacturing, using codes in the interval  $[17000, 20000]$ . These include textile products, clothing, leather and shoe manufacturing.

The census asks about monthly earnings and labor market activity for individuals aged 10 or above. All variables in the sample supplement are at the individual level. I aggregate them to the municipality level using individual weights provided by the IBGE. All income measures are deflated to 2010 real using the deflator provided by the IBGE. Each Census wave was conducted at the same time in the year (August through October), which mitigates any concerns about differential crop cycles when examining labor outcomes in the agricultural sector. The definition of employment and the reference period for economic variables changed substantially starting in the 2000 Census, making the key dependent variables not comparable with previous Census surveys.

My main measure is earnings from all occupations, as well as measures of income from all sources. Other sources in this latter measure generally can include rental income, financial investments, social transfers, and pensions. Restricting ages to those under the retirement age avoids the influences of pensions on social transfers. Rental income is of interest either directly or as a proxy for land values. Increased agricultural productivity could result in higher values of land which can be rented out, altering earnings opportunities and through wealth effects can influence fertility. It is not possible, however, to isolate rental income in both census years. The questionnaire in the 2000 Census survey directly asks about components of ‘other sources of income’, including rent, however the 2010 Census includes rental income inside a larger group simply defined as ‘other sources’. For comparability, I take a consistent measure of ‘other sources’ which includes the broad category. Individuals working in unpaid work are recorded as having zero labor income. Note, however, individuals could be classified as being in unpaid work if their occupation they spent the most time in in the reference period was unpaid, but still report positive earnings from other jobs, if say, they also had seasonal employment.

There are municipalities where there are no recorded women working in agriculture. For these municipalities, I impute average agricultural income of women as zero.

I use activity codes from the 2000 Census to form consistent measures of sectors, as definitions of sectors change across Census years. I define three broad sectors- agriculture, manufacturing, and services. I additionally define the sector of light manufacturing as a subset of manufacturing which incorporates industries such as textiles and leather goods manufacturing which largely employ women. An individual is assigned to a particular sector if the work they devoted the most hours to in the reference period falls into that sector.

## Live births

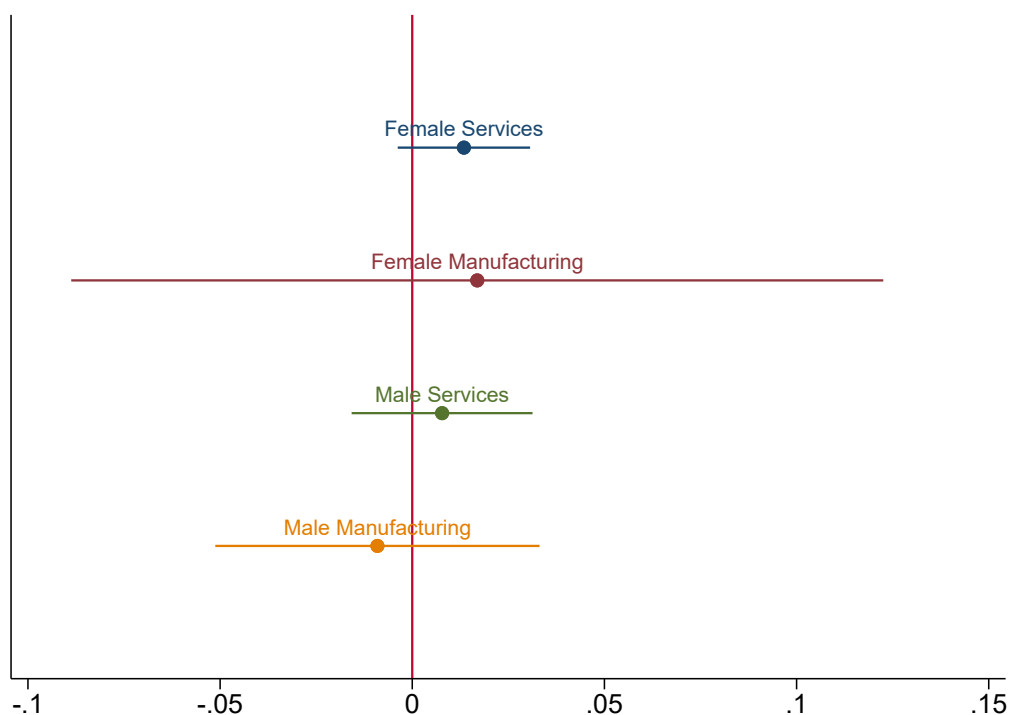
Although data collection for SINASC began in 1994, implementation was gradual and not all municipalities in years before 2000 have available data. When constructing data for annual level specifications, I omit years 1994-1996, which had a significant number of missing municipalities to avoid compositional bias in results. Moreover, delayed implementation in some states (such as Minas Gerais) occurred until 1997. After 1997, SINASC began to contain as much or greater coverage than data obtained from other sources such as the civil registry from different localities (Lima et al., 2006). Thus, the annual specifications will use birth rates from 1997-2019.

When obtaining data on live births and population by year and municipality, there are a few municipalities where population data is available, however there is no information on live births. Occurrences such as this only happen in years prior to 2000. In these situations, I assume there were no births for these observations and impute a zero for live births (i.e., the numerator of the birth rate) when creating birth rates.

## Population Counts

Population counts from 1996, 2000, and 2010 come from the 1996 population count, the 2000, and the 2010 Census, respectively. In the inter-census years, I use estimates of the population which are conducted by the IBGE and adjusted by the Ministry of Health for different ages and sexes. I additionally use projections that extend to 2019, made to be consistent with official population statistics in 2018. It is important to note that for different year ranges, the calculation of estimates contain different adjustments. To the extent that these differences in estimation are constant across municipalities in a given year (or to municipalities within a given state-year), this can be adjusted for by the inclusion of state-time fixed effects in the empirical analysis.

Figure A.1: Other Sectors Income



Figures show first difference estimates for average municipality earnings in the non-agricultural sectors of the economy. All data is taken from the individual level from the sample supplement of the population census, and aggregated to the municipality level using sample weights provided by IBGE. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state.

## Section 2: Context in Brazil

President Luiz Inácio Lula da Silva was initially opposed to GE crops when elected in 2002 (Bentham, 2013). In mid-2003, Lula's Chief of Staff Jose Dirceu reaffirmed this opposition, citing health and environmental concerns. However, a few months later Vice President Jose Alencar signed a decree legalizing the first use of GE soy for the upcoming harvesting season (Staff, 2003). It also was not

Brazilian agriculture is often characterized by its large, heavily mechanized farms and remarkable productivity growth compared to other sectors of the economy. The agricultural sector grew by over 105.6% from 2000 to 2013, in part driven by the adoption of new technologies. While these features distinguish Brazilian agriculture from other developing country contexts, smaller scale farms between 0-10 Ha, many of which are family farms, still dominate much of Brazilian agriculture in terms of pure numbers and employment. Moreover, the productivity gains were realized by both the smallest and largest firms (Arias et al., 2017). In 2006, these smaller farms employed about 75% of the working agricultural population, about 12.3 million workers, and constituted 84% of all farms, despite accounting for only 24% of the cultivated area (Arias et al., 2017). Larger farms have also been noted for their high levels of hired labor (Flaskerud, 2003).

### Family Farms- Employment and Ownership

Comprehensive examination of family farms began with the 2006 Agricultural Census. The Census reveals that family farms accounted for about 92% of farms less than 10 HA in size (Helfand et al., 2015).<sup>89</sup> Family farms in Brazil use labor extensively. Consistent with experiences of historical development and current developing economies, the 2006 Agricultural Census (tables 1113 and 1114) shows that family farms provide a large source of employment opportunities for women. They employed 74% of women who work on farms, and 66% of men in 2006.<sup>90</sup>

Family farms in Brazil are mostly run by men, and female family farmers are more likely to be illiterate- with illiteracy rates of .33 and .26 for women and men, respectively (Helfand et al., 2015). There is regional heterogeneity as well- most family farms are located in the North East, South, and South East regions of Brazil, while they make up the majority of establishments in all regions.

### Family Farms Participated in the Agribusiness Boom

Arias et al. (2017) show that although much of the boom in agriculture arose from large farms typically associated with Brazil, small farms (less than 5 Hectares) thrived over this period, combining 'state-of the art technology with abundant family labor.' Family farms in Brazil grew in precisely the areas where the agribusiness boom was the strongest, expanding at a similar pace as non-family or corporate farming, and are largely integrated into production chains with non family farms (Guanziroli et al., 2013).

### Soy Production in Brazil

Soybean is the largest crop in Brazil by scale and value, and family farms participate heavily in their production. For instance, Guanziroli et al. (2013), applied a methodology to expand analysis of family farming to the 1996 census and estimate that family farms had a share of 1/3rd of soybean production in 1996, and continued to have a 1/4th share in 2006. Table A1 shows that despite this small (albeit still

---

89. Family Farms are legally defined in Brazil as farms whose land holdings less than four fiscal units, derive most of their household income from agriculture, primarily use household labor, and manage farm activities themselves.

90. Note that family farms also used hired labor Guanziroli et al., 2013.

significant) share of total production of soybeans, family farms constituted about 76% of all farms that produced soy.

Soy production in Brazil is often associated with the remarkable development and large farms of the Cerrado region,<sup>91</sup> whereby massive investment by Embrapa and other research institutes along with entrepreneurial farmers migrating from the traditional soy producing regions in the South led the way in the soy agribusiness boom over the past few decades. Large and small soy farms alike utilize family and non family labor (Silvestrin Zanon and Saes, 2010). Small scale and family farms are heavily involved in soy production in the Centerwest, and dominate in terms of number of establishments in the traditional soy producing South. In fact, Table A1 shows that a majority of all soy farms are family farms in the South region of Brazil.

Table A.1: Family and Non Family Farms in Soy Production

Region	No. Family Farms	No. Non Family Farms	Share
Centerwest	4223	9563	0.31
North	229	501	0.31
Northeast	194	1123	0.15
South	156944	38022	0.80
Southeast	2425	3791	0.39
Brazil	164015	53000	0.76

Table reports the number of family and non family farms involved in soy production. Data are taken from the 2006 Agricultural Census (table 949), which is the first agricultural census that investigates formally family farming. The final column is the proportion of farms that are family farms.

---

91. The Cerrado lands largely incorporate the Centerwest region, which includes the state of Mato Grosso.

## GE Soy

Table A.2: Potential Yields Predict the Actual Adoption of GE Soy

	(1)	(2)
	GE soy <sup>2006</sup>	GE soy <sup>2017</sup>
$\Delta$ Pot. Soy	0.0148*** (0.00430)	0.0377*** (0.00836)
Observations	4,045	4,241
R-squared	0.475	0.377
Controls	ALL	ALL
Standard Errors Clustered at the Microregion Level		
*** p<0.01, ** p<0.05, * p<0.1		

Table reports cross sectional regression results showing that potential yields predict the adoptions of actual GE soy yields. The dependent variable is the share of all harvested area with GE soy, taken from the 2006 Agricultural Census, table 824, and the 2017 Agricultural Census, table 6958. Each column is a separate regression, with column (1) using the GE soy share in 2006, and column (2) using the share in 2017. These are run separately as the Agricultural Census changes its reference date over these two waves, making the variables not comparable across years. The lower number of observations reflect data availability from the Agricultural Census. All regressions include all controls and state fixed effects, which make this a comparison between municipalities in the same state. Pooling both the 2006 and 2017 years into one regression, I estimate that a one ton per hectare increase in potential soy leads to a 1.9 percentage point increase (p-value: 0.007) in the share of land harvested with GE soy.

## Section 3: Further Results and Robustness

### Alternative Specifications for Labor Market Outcomes

Table A.3: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Female Employment

Share of Women in Each Sector: Structural Transformation				
VARIABLES	(1) $\Delta$ Agriculture	(2) $\Delta$ Manufacturing	(3) $\Delta$ Services	(4) $\Delta$ Unpaid Agriculture
$\Delta$ Pot. Soy	-0.00569 (0.00477)	0.0117*** (0.00324)	-0.00953** (0.00452)	0.0470*** (0.0142)
Observations	4,254	4,254	4,254	4,254
R-squared	0.107	0.113	0.117	0.211
Controls	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.21	0.09	0.68	0.48

Robust standard errors clustered at the Microregion level are in parentheses

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table reports first difference estimates using Decennial Census data from 2000 and 2010. The dependent variables are the share of women in each sector, i.e. the denominator includes only women. Column 4 of Panel B is the share of women within agriculture working in unpaid work. Unpaid work includes helping another household member with their work or working in cultivation for a household's sustenance for no pay. The means of the dependent variables in 2000 are reported in the last row of each panel. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state.

Table A.4: First Difference Estimates of the Effect of Potential Soy on Agricultural Earnings: Ages 10-49

	(1)	(2)	(3)	(4)
	$\Delta \log(\text{Overall Earnings}^f)$	$\Delta \log(\text{Labor Earnings}^f)$	$\Delta \log(\text{Overall Earnings}^m)$	$\Delta \log(\text{Labor Earnings}^m)$
Pot. Soy	-0.156*** (0.0463)	-0.117** (0.0518)	0.0371* (0.0191)	0.0251 (0.0196)
Observations	4,254	4,254	4,254	4,254
R-squared	0.108	0.102	0.077	0.082
$\Delta \log \bar{Y}$	.78	.57	.24	.23
$\bar{Y}_{2000}$	233.77	210.80	505.89	483.35

Standard Errors Clustered at the Microregion Level  
\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table reports first difference estimates using Decennial Census data from 2000 and 2010. Variables here are restricted to those aged 10-49, the ‘fertile’ age population as defined by the Ministry of Health. In the first column, the dependent variable is all agricultural earnings for women, and the second column are agricultural earnings directly from labor only. The third and fourth column do the same but for men. Data are taken at the individual level, and aggregated to the municipality using weights provided by the IBGE. The last row reports the mean (in levels) of the dependent variable in the year 2000. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state.

Table A.5: First Difference Estimates of the Effect of Potential Soy on Agricultural Earnings: Long Term Residents

Primary Labor Market Outcomes				
	(1)	(2)	(3)	(4)
VARIABLES	$\Delta$ Paid Female Agriculture Share	$\Delta \log(\text{Female earnings}^{ag})$	$\Delta \log(\text{Male earnings}^{ag})$	$\Delta \log(\text{Family earnings})$
$\Delta$ Pot. Soy	-0.0195*** (0.00592)	-0.171*** (0.0622)	0.0176 (0.0224)	0.0220** (0.0108)
Observations	4,254	4,254	4,254	4,254
R-squared	0.102	0.049	0.024	0.123
Controls	YES	YES	YES	YES
$\bar{Y}$ (2000)	0.14	328.46	809.85	1380.09

Robust standard errors clustered at the Microregion level are in parentheses  
\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table reports first difference estimates using Decennial Census data from 2000 and 2010. Variables here restrict the sample to those who have lived in the same municipality for at least 10 years. In the first column, the dependent variable is the share of women in paid agricultural work. The second column uses the log of the average female municipality agricultural earnings. The third column does the same using the average male municipality earnings. The fourth column uses overall family earnings. Data are taken at the individual level, and aggregated to the municipality using weights provided by the IBGE. The last row reports the mean (in levels) of the dependent variable in the year 2000. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state.

## Inference

Table A.6: Inference: Alternative Standard Errors

	(1) $\Delta$ Paid Agriculture <sup>f</sup>	(2) $\Delta$ Agr. Female Earnings	(3) $\Delta$ Agr. Male Earnings	(4) $\Delta$ Agr. Relative Earnings	(5) $\Delta$ Family Earnings	(6) $\Delta$ Fertility
<b>Pot. Soy</b>						
Microregion	-0.00715 (0.004)	-0.115 (0.043)	0.0326 (0.018)	-0.0999 (0.028)	0.0202 (0.0085)	2.346 (0.907)
Mesoregion	(0.005)	(0.049)	(0.020)	(0.030)	(0.009)	(1.016)
Conley SE, 50km	(0.003)	(0.040)	(0.015)	(0.023)	(0.008)	(0.693)
Conley SE, 200km	(0.005)	(0.050)	(0.021)	(0.030)	(0.010)	(0.998)
Observations	4,254	4,254	4,254	4,254	4,254	4,254
R-squared	0.215	0.062	0.067	0.126	0.223	0.128

Standard Errors in Parenthesis

Table reports first difference estimates for main results using the Decennial Census data from 2000 and 2010, showing sensitivity to different assumptions on standard errors. The dependent variables from column (1)-(6) are the paid female agricultural share (Ages 16-49), the log of female agricultural earnings, the log of male agricultural earnings, the log of the relative agricultural earnings of women to men, the log of overall family earnings (taken from the household survey of the Demographic Census), and the birth rate per 1000 women aged 10-49. Data are taken at the individual level, and aggregated to the municipality using weights provided by the IBGE. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. Under each coefficient are 4 different standard errors. The first clusters by the Microregion level, the second row clusters by the Mesoregion level, and the third and fourth row use Conley Standard Errors with a cut off of 50km and 200km, respectively.

## Robustness to Different Weighting and Additional Controls

Table A.7: Main Results: Weighted by Population

Panel A: Weighted						
	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta$ Paid Agriculture <sup>f</sup>	$\Delta$ Log Agr. Female Earnings	$\Delta$ Log Agr. Male Earnings	$\Delta$ Log Agr. Relative Earnings	$\Delta$ Log Family Earnings	$\Delta$ Fertility
Pot. Soy	-0.00687* (0.00402)	-0.106*** (0.0391)	0.0324* (0.0177)	-0.0953*** (0.0277)	0.0191** (0.00813)	2.505*** (0.902)
Observations	4,254	4,254	4,254	4,254	4,254	4,254
R-squared	0.208	0.070	0.067	0.129	0.233	0.132
Panel B: Unweighted						
	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta$ Paid Agriculture <sup>f</sup>	$\Delta$ Log Agr. Female Earnings	$\Delta$ Log Agr. Male Earnings	$\Delta$ Log Agr. Relative Earnings	$\Delta$ Log Family Earnings	$\Delta$ Fertility
Pot. Soy	-0.00715* (0.00406)	-0.115*** (0.0432)	0.0326* (0.0177)	-0.0999*** (0.0277)	0.0202** (0.00848)	2.346*** (0.907)
Observations	4,254	4,254	4,254	4,254	4,254	4,254
R-squared	0.215	0.062	0.067	0.126	0.223	0.128

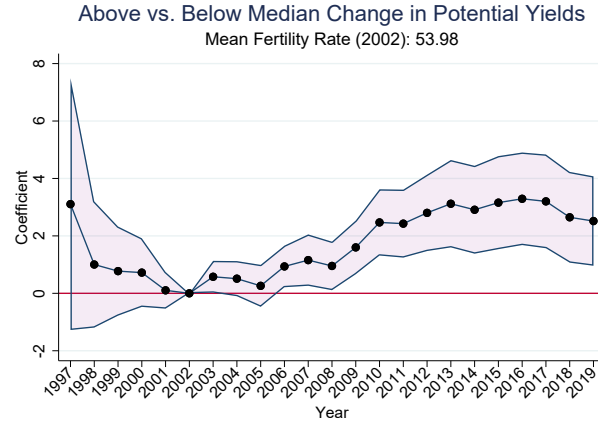
Standard Errors Clustered at the Microregion Level

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

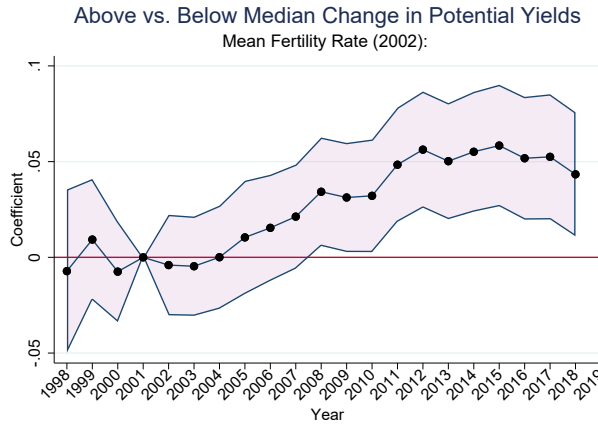
Table reports first difference estimates using Decennial Census data from 2000 and 2010. The dependent variables form column (1)-(6) are the paid female agricultural share (Ages 16-49), the log of female agricultural earnings, the log of male agricultural earnings, the log of the relative agricultural earnings of women to men, the log of overall family earnings (taken from the household survey of the Demographic Census), and the birth rate per 1000 women aged 10-49. Panel A shows regressions weighted by the log of the population in 2000. Panel B shows the unweighted coefficients from the main specifications. Data are taken at the individual level, and aggregated to the municipality using weights provided by the IBGE. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state.

Figure A.2: Robustness

Panel A Weighted by Female Population (10-49)



Panel B. Ln(Births)



These sub-figures plots estimated coefficients from (11) where the dependent variable in the first figure is the birth rate, defined as the number of live births per 1000 women aged 10-49, and the second is the natural log of the number of births. The data run from 1997-2019 and have a balanced sample of 4254 municipalities, creating 97,842 observations. The year prior to the legalization of GE soy, 2002, is omitted. The shaded bands represent 95 percent confidence intervals. The regression includes all baseline controls interacted with a time trend, as well as the measure of potential maize interacted with a vector of year indicators. The specification using the log of births also includes the log of the female population aged 10-49 as a control. All regressions include all controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. Standard errors are clustered at the Microregion level.

Table A.8: Additional Controls

	(1) Δ Paid Agriculture <sup>f</sup>	(2) Δ Log Agr. Female Earnings	(3) Δ Log Agr. Male Earnings	(4) Δ Log Family Earnings	(5) Δ Fertility
Pot. Soy	-0.00683* (0.00408)	-0.114*** (0.0433)	0.0317* (0.0174)	0.0203** (0.00836)	2.429*** (0.901)
Observations	4,253	4,253	4,253	4,253	4,253
R-squared	0.217	0.063	0.070	0.225	0.133
Controls	ALL + Market Access	ALL + Market Access	ALL + Market Access	ALL + Market Access	ALL + Market Access
Oster (2019) Coefficient	-0.002	-0.085	0.028	0.025	2.808

Standard Errors Clustered at the Microregion Level

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table reports first difference estimates for main results using the Decennial Census data from 2000 and 2010, adding in additional controls for pre-existing market access. I include the controls for distance of a municipality to railways and distance to highways in existence in 2000s. The dependent variables form column (1)-(5) are the paid female agricultural share (Ages 16-49), the log of female agricultural earnings, the log of male agricultural earnings, the log of overall family earnings (taken from the household survey of the Demographic Census), and the birth rate per 1000 women aged 10-49. Data are taken at the individual level, and aggregated to the municipality using weights provided by the IBGE. All regressions include all other controls and municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. The last row reports estimations from Oster (2019), the coefficient assuming that unobservables would increase the r-squared by 30% and that the selection on unobservables is equal to the selection on observables.

## Fertility: Descriptive Statistics from the DHS and Alternative Measurement

When thinking about fertility outcomes, one must consider the use and availability of contraceptives. Brazil presents an interesting case, as it is characterized by a large number of female sterilizations (Martine, 1996), contributing to the stopping of child births. This would serve to mute any effects on increasing fertility. The Demographic Health Surveys from Brazil taken in 1996 (this is the latest survey conducted in Brazil) asks about contraceptive use. Appendix Table 9 shows descriptive statistics of mothers who are sterilized vs others. While sterilization is prevalent (an estimated 27.3% of mothers), these women tend to be older and have higher completed fertility than others. Sterilized mothers tend to be older (37 years old vs 27), and tend to have larger families (3.5 children vs 1.3). These suggest that sterilization is likely used to stop births on the intensive margin, rather than the extensive margin. Moreover, it typically occurs after already having 3 children. These descriptive statistics as suggesting that the contraceptive setting in Brazil leave the results found in this paper to be plausible.

Table A.9: Mean Characteristics of Mothers by Sterilization Status

	Sterilized	Non Sterilized
Mean Age	37.09	27.35
Mean No. of Children	3.55	1.34

Table reports descriptive statistics for the age and number of children of mothers by sterilization status. Not sterilized includes all other contraceptive methods including not using any. Data are taken from the 1996 DHS wave III survey, taken from the mothers questionnaire. Statistics are weighted by probability weights provided by the DHS.

Table A.10: First Difference Estimates of Soy Potential Yields on Fertility Measures From the Census

	(1) $\Delta \log(Completed)$	(2) $\Delta \log(Children)$
$\Delta$ Pot. Soy	0.0159* (0.00860)	0.0245*** (0.00947)
Observations	4,254	4,254
R-squared	0.172	0.218
$Y_{2000}$	17527	3463
Standard Errors Clustered at Microregion Level		
*** p<0.01, ** p<0.05, * p<0.1		

Table reports first difference estimates using Decennial Census data from 2000 and 2010. In the first column, the dependent variable is the log of the total number of children (of all ages) as reported by households in a municipality, and the second column is the log number of children (less than 5) in a municipality. Data are taken at the individual level, and aggregated to the municipality using weights provided by the IBGE. The last row reports the mean (in levels) of the dependent variable in the year 2000. All regressions include municipality and state by year fixed effects, which make this a comparison between municipalities in the same state.

## Fertility: Regional Heterogeneity

Appendix Table 1 shows 76% of soy farmers are family farmers, and there is regional heterogeneity in the composition of farmers. For instance, the two largest producers by number of establishments of soy include the South, the original soy producing region of Brazil, and the Centerwest, the oft-cited driver of the large soy-driven agribusiness boom. While other regions do have many farmers growing soy, these two regions represent 90% and 6%, respectively, of all establishments producing soy. Moreover, the southern region comprises 80% family farmers in soy, whereas the Centerwest is more dominated by larger commercialized nonfamily farms, who make up 69% of soy farmers in the region. Thus, the Southern regions of Brazil are more similar to farm structures in other developing economies, whereas the Centerwest has a higher number of large farms.

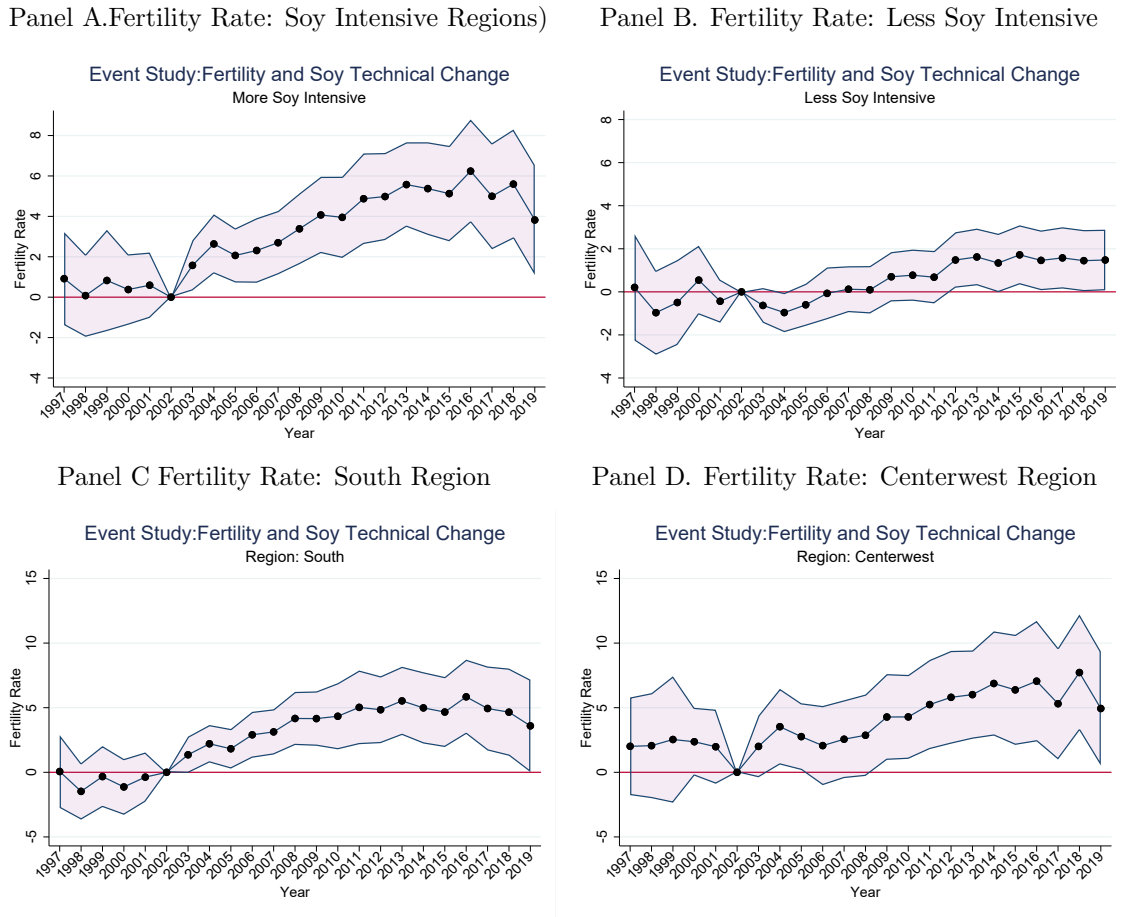
I exploit this regional heterogeneity to test whether the effects are driven by any peculiarities of the agricultural structure of Brazil. I first estimate the equation (11) on two broadly defined regions, more soy intensive and less intensive, to first verify that the results are primarily driven by these regions. The soy intensive regions include the Centerwest and South, and the less soy intensive regions are all remaining regions.<sup>92</sup> We would expect to see most of the increases in fertility driven by the more soy intensive regions. Figure 3 panel A and B plot the coefficients with 95% confidence intervals for both of these regions and each panel represents a separate regression. It is clear that most of the effects on fertility are being driven by these more soy intensive regions.

Within the soy intensive regions, the Southern regions, more dominated by smaller family farms, are

92. Note that the Appendix Table is based off of the 2006 Agricultural census data- collected after the legalization of GE technology- making this selection criteria of soy intensive and non-intensive seemingly made off of a possibly endogenous outcome. However, I choose this division because the Southern and Centerwest regions of Brazil have been the dominant players in Soy production even before the legalization of GE crops(Cattelan and Dall'Agnol, 2018).

likely to be more similar in composition to other developing country regions' agricultural structure. If the impacts of the adoption of these technologies are peculiar to Brazil's agricultural structure, we may expect to see the fertility response dominated by the Centerwest regions. Panels C and D estimate regressions interacting all event study interactions and controls with region indicator variables and plot the coefficients for the Southern and Centerwest regions, respectively. Both panels indicate that despite the difference in establishment composition, both regions experience increases in fertility following the adoption of GE soy technologies. This exercise demonstrates that the effects of these agricultural technologies are occurring regardless of the underlying agricultural structure.

Figure A.3: Event Studies: Regional Heterogeneity



These sub-figures plots estimated coefficients from (11) where the dependent variable birth rate, defined as the number of live births per 1000 women aged 10-49, split by regions. Panels A and B run the regressions on the Soy Intensive (Southern and Centerwest) and Less Soy Intensive (remaining regions) regions, respectively. The soy intensive specification contains 26,014 observations, and the less soy intensive region contains 71,829. Panels C and Panel B divide the Soy Intensive regions into the Southern region, more dominated by small scale family farm agriculture, and the Centerwest region, more dominated by larger non-family farms, respectively. The year prior to the legalization of GE soy, 2002, is omitted. The light purple bands represent 95 percent confidence intervals. The regression includes all baseline controls interacted with a time trend, as well as the measure of potential maize interacted with a vector of year indicators. All regressions include municipality and state by year fixed effects, which make this a comparison between municipalities in the same state. Standard errors are clustered at the Microregion level.