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

Vivek S. Moorthy

College of the Holy Cross

May 2024

Abstract

I study how 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, J43, 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 As-sociation 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 vmoorthy@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 the 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, á 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 disproportionately performed by women (Sofa and Doss, 2011; Grassi et al., 2015).

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.[1,2] 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 9% reduction in women's earnings in agriculture with no movements into other sectors of the economy. Within agriculture, the effects are driven by agricultural employees, suggesting that larger farms with wage labor are the primary, or early, adopters these technologies and are shedding their labor force.[3] Women in soy-producing regions reallocated from paid work into work for another member of the household's job or directly for household sustenance, all for no pay. This establishes a novel result: GE soy technological change was female labor-saving without creating alternative employment opportunities for women in other sectors of the economy.

These results are striking given the context of overall rising female labor force participation in Brazil over this time period.[4] On the other hand, I find evidence of increases in male earnings, and that men affected by the technologies reallocated into non-agricultural employment. Further, this is not limited to higher educated men – lower educated men are also able to move sectors, whereas neither higher or lower educated women leave the agricultural sector. 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

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. Paris and Chi (2005) document a similar change in a case study in Vietnam, where row seeder technologies in rice largely displaced female wage laborers in hand-weeding occupations.

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

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. 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. Given the context of declining overall fertility in Brazil, this change is driven by slowing the decline in soy-producing municipalities. In other words, municipalities that adopted these technologies would have experienced faster declines in fertility in the absence of this technology. I further find that technological change led to a small increase in marriages and that the increases in fertility are driven by marital births. Finally, I find no changes in differential migration, infant mortality, child labor usage, or the easing of credit constraints in driving these results.

This finding at first seems counterintuitive because fertility tends to decline as countries develop and technologies alter the structure of the labor market. I present a simple model of fertility that incorporates gender-specific effects of 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.[5] 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.[6]

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

5. 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).

6. 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. 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 from higher overall earnings.

that typically dominate the income effects (Schultz, 1997; Kitchens and Rodgers, 2023).[7] 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.

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 have gendered effects 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). Closely related to my work, 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. 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.[8]

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 (2023) 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.[9] Most evidence of the effects of negative labor market shocks for women on fertility comes from specific

7. 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.

8. 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.

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

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 (2023)'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.[10]

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.[11] They create a theoretical model whereby declines in agricultural employment come with reductions in female labor force participation. I provide causal micro-evidence of this relationship in a modern developing country. Further, Ager et al. (2020) and Gehrke and Kubitza (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, Section 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.

10. 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.

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

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 this technological change may alter the incentives surrounding fertility choice. 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-specific effects of 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 9, 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.[12] 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.[13] Assume that the cost of children enters entirely through the opportunity cost of time used in raising them,[14] 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, τ^w, and the fraction of the male's unit time endowment, τ^m, required in child rearing.[15] The household's budget constraint is then:

$$(\tau^w y^w + \tau^m y^m)n + c \leq y^w + y^m \tag{1}$$

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

13. 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).

14. 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.

15. 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.[16] For simplicity of exposition, assume household preferences are represented by the following log-linear utility function:

$$u(n, c) = \gamma\, ln\, n + (1 - \gamma)\, ln\, c; \qquad \gamma \in (0, 1) \tag{2}$$

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

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

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:

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

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.:

$$\frac{\partial n^*}{\partial y^m} > 0 \tag{5}$$

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 α represents the technological level. The change in fertility in response to technological change is given by the total derivative of

16. 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.

$n^*(\tau^w, \tau^m, y^w, y^m)$ with respect to α:

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

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 α and manipulating yields the following elasticity representation:

$$\epsilon_{n,\alpha} = \epsilon_{n,y^w} \, \epsilon_{y^w,\alpha} + \epsilon_{n,y^m} \, \epsilon_{y^m,\alpha} \tag{7}$$

where $\epsilon_{y,x}$ represents the elasticity of variable y with respect to x. Thus, the overall responsiveness of fertility to 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., $|\epsilon_{n,y^w}| < \epsilon_{n,y^m}$.[17] 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.[18]

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 $\frac{dy^m(\alpha)}{d\alpha} > 0$. In other words, the soy technology interacts in

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

18. 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.

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.[19]

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 is driven by the stopping of births, rather than spacing or delaying births.[20] Thus, the Brazilian demographic transition is characterized by by negligible changes in the 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).[21]

It is plausible that if the GE soy technology also destroys occupations that children perform, there could be an offsetting effect, lowering the demand for children.[22] Moreover,

19. In the framework of Mookherjee et al. (2012), this would drive a positive relationship between parental earnings and fertility *within* agricultural occupations regardless of the relative strength of the substitution and income effect, as they move in the same direction.

20. 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 2.

21. 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.

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

increased economic growth may also reduce fertility by lowering infant mortality. Thus, I explore these as outcomes of interest in the analysis.

Due to data constraints, I am unable to empirically identify any effects from changing bargaining power within the household. However, bargaining channels would likely further reinforce the opportunity cost and income effect channels in this context. Data from the 1996 Demographic Health Survey[23] indicates that, on average, men prefer more children than their wives, suggesting that household bargaining channels are likely present.[24,25] 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. [26]

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).[27] Qualitative research in Brazil and reports from the Brazilian Agribusiness Ministry confirm that women disproportionately perform these tasks in Brazil as well, in both small and large farms alike. Men are more commonly involved in management, contact with agronomists, and the investments and use of new technologies (Brumer, 2008; Brumer, 2004; Lastarria-Cornhiel, 2017).[28,29]

Women typically earn income from selling products that they process in the family establishment directly on the market or by selling their labor to a third party for wages (Brumer, 2008). Strong gender norms typically dominate the landscape within agricultural areas in Brazil. Deere (2005) notes that while increased wage earning opportunities for women in agribusiness affords them increased economic independence, in

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

24. 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.

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

26. Finally, land values can affect fertility through the bargaining channels mentioned above or through increased income if the land is rented. Table 320 of the 2006 Agricultural Census shows that over 70% of farms in Brazil are owned by men, thus any increase in land values would likely increase the bargaining power of men or be reflected in men's earnings directly.

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 Section 2, I discuss Brazilian farm structures and agricultural 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.

some cases, married women working in wage employment in agriculture faced increased conflict within the household. In both family farms and wage labor, women were typically relegated to low skilled tedious tasks such as weeding. Lastarria-Cornhiel (2017) notes that this gender division leaves women with little opportunity to acquire the necessary skills or tacit knowledge that can be employed in alternative opportunities. On the other hand, male workers often gain useful skills from managerial roles in the family establishment and are the main beneficiaries of formal training within agro-industry that teaches management and the operation of machinery. Due to their low skill acquisition, women face limited opportunities in urban employment (Waltz, 2016) and their most prominent outside option is domestic work (Brumer, 2004).[30] Further, gender norms dictate that men inherit farms, and daughters typically only inherit farms if their husband is also a farmer (Arends-Kuenning et al., 2021). As rural women often only receive training and education for working on a farm, Paulilo et al. (2003) notes that many can only ensure continued work on a farm throughout their working life by marrying.

Table 1 Panel B presents demographic and economic characteristics about agricultural employment from the 2000 Brazilian Census, which asks workers about the type of work in their main occupations.[31] In 2000, about 35% of women in agriculture were employees for other agricultural establishments with or without formal contracts in their main economic activity, and a further 15% reported being self-employed workers. The remaining 48% were not remunerated in their main occupation.[32] The corresponding numbers for male agricultural workers were 51%, 29%, and 19%, respectively. Panel B of Table 1 shows that female employees tend to be younger and are less likely to be married. While they are marginally more educated than other occupational groups, educational attainment is very low across the entire sector. A majority of those self-employed as well as unpaid in their main occupation are married.

3.2 Legalization of GE Soy

Soy has grown into one of the major crops in Brazil over the past few decades. In 2000, soy constituted 29% of all harvested area in Brazil, and grew to 44% of the total harvested land by 2019.

Genetically engineered soy from Monsanto was commercially released in 1996 but was not legalized in Brazil until 2003.[33] While the technology was legalized in 2003, the

30. Further, in my own conversations with domestic workers in Brazil, they cite large barriers for rural women to gain employment in domestic work, such as long commutes and a lack of networks to connect them with potential employers.

31. Main occupations refer to the economic activity respondents current/most recent job, or the activity in which they spent most of their time in during the reference period.

32. 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 48% then is likely an *overestimate* of the total number of female workers solely working in unpaid work in agriculture.

33. GE soy was among a new class of agricultural technologies unveiled in the mid-1990s. Innovations

passage of a Biosafety Bill[34] in 2005 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 timeline 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).[35] 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 (Benthem, 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.[36] Thus, this technology directly eliminates female dominated occupations.[37]

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.

34. Law no. 11.105.

35. 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.

36. 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.

37. 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. An increase in machine usage may then mitigate any potential reductions in demand for men working in vulnerable operations as men disproportionately operate machinery.

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 Geografia e Estatística (IBGE).[38] 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).[39] Finally, I obtain data on potential GE soy yields from the FAO Global Agro-Ecological Zones database (FAO-GAEZ).[40]

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.[41] 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.[42] There are also measures of family earnings, which are the sum of the earnings of all individuals who constitute a 'family': Groups of 2 or more people related to a defined household head (e.g., a household head and their spouse).

I construct the monthly earnings and the share of workers in a given sector by gender.

38. 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/.

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

40. 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.

41. The way in which economic activity was measured changed significantly starting in the 2000 Census, including who counts as economically active, making it difficult to create comparable measures with previous waves. For example, earlier Census waves had requirements that individuals who worked for no pay in their main occupation needed a minimum of 15 hours per week over the course of the year to be counted/recorded as employed.

42. 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.

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 aggregate the latter two into the 'non-agricultural' sector for the main analysis. 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.

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.[43] 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. DATASUS also provides inter-census population counts, which I use for the denominator of birth rates. 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. Further, the parameters for the calculations of these potential yields are based on field experiments at agricultural research stations, rather than actual agricultural inputs and outputs.[44] 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

43. 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).

44. Costinot et al. (2016) provides a detailed description for more information.

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.[45] 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.[46]

I define my treatment as the difference in potential yields between the high and low regimes. In other terms: ΔPot. Soy = Potential Yields^{high} - 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. Appendix Table A1 Panel D shows the mean value of the potential yield differential from high to low technology in soy is 1.8 tons per hectare.[47]

4.4 Summary Statistics

Appendix Table 1 shows that the demographic transition is still underway in Brazil over this time period. In 2000, there were on average 66.24 births per 1000 women aged 16-49. This declined by 14 births per 1000, constituting a 21% decline, by 2010.[48]

Figure 1 shows the average of municipality (N=4254) sector employment shares in Brazil during this same decade. Sector shares are defined as the fraction of all workers (i.e. both men and women) aged 15-55 working in that particular sector.[49] 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.[50] Light industries made approximately 3% of the employment share in 2000 and grew by about 8% over the decade. Overall, Figure 1 suggests that industries that employ women, such as services and light manufacturing, may not be expanding rapidly enough to absorb

45. Further,the FAO-GAEZ researchers did in fact use GE soy for calculating the high yield potential soy measure, see Bustos et al. (2019) for more discussion.

46. 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, more information panel (under the GAEZ V3.0 GLOBAL AGRO-ECOLOGICAL ZONES box), and User's Guide.

47. Brazil's soy production in 2000/2001 was estimated at about 2.7 tons per hectare. Prices of a ton of soy around this time period were about 236.8 US dollars per ton, which would translate to a sizeable 7100 dollar increase in revenue for a medium sized farm of 30 hectares. 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/cropexplorer/pecad_stories.aspx?regionid=br&ftype=prodbriefs.

48. 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.

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

50. 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.

rapid reductions of employment in agriculture.

Figure 1 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. Appendix Table 1 also shows that while there remain large gender inequalities in earnings across the entire Brazilian economy, all sectors experienced proportionally larger increases in women's earnings than that of men's over the decade.[51]

Table 1 Panel A shows demographic characteristics of men and women across the three major sectors of the economy in 2000. Agricultural workers tend to be less educated than the other sectors of the economy, with women having about 4 years of education on average. Further women in agriculture are more likely to be married, and have the highest levels of fertility across the sectors with 4 children on average.

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.

51. In the agricultural sector, women made on average 43% of the earnings of men in 2000.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 $124 per month in 2010 dollars, and for men, $290 per month.

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.[52] To invalidate my design, these factors would have to vary systematically among higher and lower soil suitability municipalities within a state-year.

Table 2 presents statistics on baseline factors from the 1991 Census for municipalities that fall above the median increase in potential yields to those below the median. Column 3 presents a within state average difference in these characteristics between these municipalities. Higher potential yield municipalities tend to be more developed on average than lower potential yield municipalities at baseline, including being less rural and have higher income per capita. Any pre-legalization differences in Table 2 are not inherently problematic unless they translate into differential trends. Figure 2 Panel A displays a descriptive graph on the evolution of fertility between above and below median municipalities within states.[53] Consistent with higher development and urbanization, above median potential yield municipalities have lower levels of fertility prior to the legalization of GE soy, however the changes in fertility seem to evolve similarly. As the time series progresses post legalization, this gap narrows until the levels of fertility begin to equalize in the later years. This is suggestive of GE soy decelerating the decline in fertility.

52. 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 (Benthem, 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 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.

53. In practice, this is constructed by first regressing potential yields on the set of state fixed effects, and splitting municipalities into above and below values of the residuals. Appendix Figure 1 presents the same graph without partialling out state effects, showing similar trends.

In my empirical analysis, I will control for these covariates allowing for differential trends in municipalities with heterogeneous baseline characteristics, and evaluate the sensitivity of any results to the inclusion of these factors. I also must account for the simultaneous adoptions of other agricultural technologies that could be occurring over this time period. As mentioned in Section 3.2, the introduction of new seed technologies and machinery in maize production during this same period allowed for the expansion of the crop into multiple seasons in the year. To adjust for the simultaneous adoption of these technologies I will also include a measure of the increases in potential maize from new technologies, defined analogously to that of potential soy. With this in mind, Figure 2 Panel B plots the variation in the potential soy measure linearly uncorrelated with potential maize.[54] 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 predict the actual adoption 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.

Table 3 reports coefficients of these regressions.[55] Panel A uses data from the 2006 Agricultural Census and Panel B uses the 2017 Census. The dependent variable in Columns (1) and (2) is the percent of farm land harvested with GE soy. All regressions include state fixed effects ensuring that these are comparisons within the same state, and Column (2) additionally adds all the baseline controls listed in Table 2 and the measure of potential maize, representing partial correlations using the identifying variation. Both panels 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 1.2 percentage points in 2006 and by 3.7 percentage points in 2017.[56,57] Columns (3) and (4) in both panels then examine the basic correlations in the data between the actual

54. 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 (ρ=.79). I add back the mean of potential soy for scaling.

55. The lower sample size in these estimates reflects the coverage of the Agricultural Census.

56. The associated F statistic on the soy coefficient is 12.87 for the 2006 regression, and 19.67 for 2017.

57. 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.

adoption of GE soy and fertility. In 2006, before the inclusions of controls, there is a significant negative correlation between the actual adoption of GE soy and fertility, however including the controls lowers the coefficient, and there is no longer a statistically significant relationship. This is suggestive that the (early) adopters of GE soy are regions that are richer or more urbanized, and consequentially have lower fertility at baseline. However by 2017, there is significant positive association between actual GE adoption and fertility.

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

$$y_{mt} = \alpha_m + \mu_{s,t} + \kappa Potential\ Soy_m * t + X'_{m,t}\delta + \psi_{m,t} \tag{8}$$

Where α_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.[58] The parameter κ 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 Table 2) 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):

$$\Delta y_m = \Delta\mu_s + \kappa Potential\ Soy_m + X'_m\delta + \Delta\psi_m \tag{9}$$

Which would identify κ from the within-state cross-municipality variation over time. Baseline specification standard errors are all clustered at the microregion level to allow for

58. Rather than interacting the change in potential soy with time, another way to conceptualize this specification is to allow the potential soy measure to take on the values for low input regimes in the pre-period (2000) and the high input regimes in the post period (2010) when these technologies are now available.

spatial correlation.[59] 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.[60]

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 between 1997-2019 from DATASUS, I also can estimate:

$$y_{m,t} = \lambda_m + \tau_{s,t} + Potential\ Soy_m^{Med} * \mathbb{1}\{t >= 2003\}\sigma + X'_{m,t}\beta + \eta_{m,t} \tag{10}$$

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

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

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 σ_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).

59. 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.

60. 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.

6 Labor Market Results

Before examining the gendered impacts in agriculture, Appendix Table 4 summarizes the overall economic impact of soy technological change. Overall, the technology induced structural transformation out of agriculture: Municipalities with higher GE soy adoption within states experienced an overall increase in household earnings and a reduction in agricultural employment.[61] There is also no change in overall female or male labor force participation or unemployment rates.

Table 4 presents the results of estimating equation (9) with gender-specific measures of agricultural earnings and hours. Columns 1 and 2 use the log of the average of women's earnings in agriculture and the number of hours women in agriculture work in their main occupation as the dependent variables.[62] Columns 3 and 4 use these same measures but for men, and Column 5 looks at overall family earnings for households working in agriculture. Families here are defined as groups of two or more people related by birth or marital status and residing together, such as an individual living with their spouse or partner. The estimates in Table 4 all confirm and establish that soy technological change differentially affected each gender. Specifically, it led to a large reduction in female earnings. An increase in one ton per hectare of potential soy leads to a statistically and economically significant 11% reduction in women's earnings in agriculture, with no change in the number of hours worked.[63] There is a statistically insignificant but positively signed 3% increase in men's agricultural earnings. Column 5 shows the effect on family earnings at the household level for those working in agriculture. Overall, agricultural households see a 2.7% increase in their earnings (p-value: .08).[64]

These tables present the ITT effects. Another way to interpret these estimates is to scale them 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 Table 3 imply a coefficient in 2010 of 2.1. In other words, an increase of one ton per hectare of potential soy leads to a 2.1

61. Note this result is essentially a replication of a major finding from Bustos et al. (2016) with a more restrictive specification using state by year fixed effects. Thus, structural transformation is also occurring within states in Brazil.

62. There are 21 municipalities in the sample where agricultural earnings for women is zero in a given year. I discuss this further in Appendix Section 1. All dependent variables are defined as shifted log transforms for consistency across regressions, however the results shown here are robust to the inverse hyperbolic sine transformation.

63. The null effects on hours provide evidence against a case where women are consuming more leisure.

64. Some individuals may live alone and/or do not constitute a 'family' in the data, i.e., no relation to a defined household head. Thus household earnings in this context do not necessarily correspond to a simple aggregation of all men and women in the sector. I however perform this exercise with 'overall agricultural earnings' in Panel B of Appendix Table 4, finding no significant changes in earnings. Note that this does not contradict the household earnings result and the predictions from 6 remain given positive income and substitution effects from men and women's earnings, respectively.

percentage point increase in the share of land reaped with GE soy by 2010. Scaling the coefficients in Table 4 by this 'first stage' estimate suggests that a 1 percentage point increase in the share of land harvested with GE soy leads to a 5.2% decrease and 1.3% increase in female and overall agricultural household earnings, respectively.

All regressions in Table 4 include the controls discussed earlier. Given the baseline differences in levels between higher and lower potential soy municipalities shown in Table 2, Appendix Table 16 examines the sensitivity of these coefficients from sequentially adding controls to the regression. Each regression from Table 4 is estimated first with no controls besides the state-year and municipality fixed effects, then adding the controls for percent rural and technological change in maize, and finally adding the remaining controls. Reassuringly, the estimated effects on female earnings are quite stable across these specifications. Both male and overall agricultural household earnings increase in response to the inclusion of the controls.[65] 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.[66]

One may worry that migratory responses drive these labor market results. I present evidence on the lack of differential migration by gender in Appendix Section 3. 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

65. Further, Appendix Table 14 reports coefficients on the potential maize measure, and Appendix Section 3 discusses the potential effects of maize technical change on the main outcomes of interest.

66. 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.

employment and earnings results. I find that the effects for women are maintained (and are larger in magnitude) with this restriction. For family earnings, the coefficient, while still positive, is statistically insignificant. These results are reported in Appendix Table 17.

Overall, these estimates suggest that 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.

6.0.1 Gender-Specific Sectoral Changes

Table 5 examines whether men and women impacted by the shock reallocate outside of the agricultural sector. In the first two columns, the dependent variable is the share of all female workers, in both paid and unpaid work, in agricultural and non-agricultural employment.[67] The second two columns use analogously defined variables for male workers.

Columns 1 and 2 show that there were no changes in either agricultural or non-agricultural employment shares for women - women are staying within agriculture. On the other hand, Columns 3 and 4 show that the soy technological change induced a movement of men outside of agriculture into other sectors of the economy. A one ton per hectare increase in potential soy led to a 2.8 percentage point decline in the share of men in agriculture and a corresponding 2.2 percentage point increase of men into non-agricultural employment, representing a 6.3% and 5.3% change over the baseline means, respectively. This suggests that the overall structural transformation induced by the soy technological shock found previously by Bustos et al. (2016) are driven by the movement of men outside of the agricultural sector.

Table 6 now examines reallocations within the agricultural sector. Panel A looks at the share and earnings of women in agriculture who are employees and self-employed, and the share who were unpaid in their main occupation. Unpaid work includes helping another household member with their work or working in cultivation for a household's sustenance for no pay. Noticeably, the women most affected by the soy technology were employees. A one ton per hectare increase in potential soy leads to a reduction in the share of female agricultural employees of 3.8 percentage points, and a 4.8 percentage point increase in the share of women working in unpaid work. Looking at earnings for these employees, there is a negatively signed but statistically insignificant coefficient. Thus, the overall declines in female agricultural earnings in Table 4 are driven by shifts from paid employment to unpaid work in the household.

67. The specifications here leave out some occupations, such as 'unspecified' occupations that may not fit into services or manufacturing. Thus, the baseline averages in Table 5 do not add up to 1. However, all the following results are similar if defining non-agricultural employment to mean any occupation not that is not explicitly defined as agricultural.

Panel B examines men. Similar to the results for women, the changes within agriculture are driven by employees. A one ton per hectare increase in potential soy leads to a 2-percentage point reduction in the share of employees, and a 5.5% *increase* in the earnings of male employees. There is no change in the employment or earnings of self-employed workers.

Overall, the estimates from Table 6 indicate that the new technologies were likely adopted primarily (or first) by larger farms with wage employees. Farms in regions that benefited more from the adoption of GE soy reduced their wage employment, with large negative effects for women who were likely more directly involved in the tasks being replaced (such as weeding). These displaced women then reallocated into unpaid work. While men may also be displaced from these technologies, the overall impact is smaller than that of women: Table 6 suggests a 11% decline in the share of female employees off the baseline mean compared to a 4% decline for men. Further, the opposing movement of employment and earnings for men suggest that the movement is also partially driven by a supply shift, where male employees are leaving agriculture to pursue new opportunities in non-agricultural employment, as shown in Table 5.

Panel B does show a rise in the share of unpaid male workers. However, if male employees are leaving the agricultural sector (i.e., shrinking the denominator of the shares), this would mechanically increase the share of unpaid male workers in the sector.[68]

So far, I have not considered specific sectors outside of agriculture. The gender-specific effects within the non-agricultural sector 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. A negative demand shock in the agricultural sector may reduce male manufacturing earnings if men reallocate into the sector. However, if richer municipalities also demand more male-intensive manufacturing goods (e.g., automobiles), these spillovers could increase male labor demand in the sector, increasing the new equilibrium manufacturing employment but leaving the change in equilibrium earnings ambiguous. In total, there is no clear mechanism ex-ante to expect a strong *gender-specific* effect in any direction. With these considerations in mind, Appendix Table 5 reports estimates from equation (9) using gender specific earnings and employment in manufacturing and services as the dependent variable. Overall, manufacturing is expanding and absorbing male agricultural workers. There is however movement of women from the service sector into manufacturing, although small in magnitude. 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. I discuss this further in Section 9.

68. Further, when examining changes in the total number of unpaid workers, I find increases in the number of female unpaid workers, but no significant change in the number of male unpaid workers.

7 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 7 presents the results of estimating equation (9) with the birth rate defined as the number of live births per 1000 women aged 16-49 as the dependent variable. Column 1 includes only state-year and municipality fixed effects, Column 2 adds controls the percent of the population that is rural and the potential maize measure, and Column 3 includes all the controls. Column 3 shows that an increase in one ton per hectare of potential soy leads to an *increase* of 2.7 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).[69] In light of the declining fertility in both higher and lower potential soy municipalities seen in Panel A of Figure 2, the soy technical change is generating these effects by slowing the decline in fertility in higher potential yield municipalities. In other words, fertility would have declined faster in regions adopting GE soy in the absence of the technologies.

For comparisons to studies in the U.S. contexts, Appendix Table 6 redefines the birth rate as the number of births per 1000 women aged 20-40. I find that an increase in one ton per hectare of potential soy leads to an increase of about 4 births per 1000 women, or a 4.9% increase off a baseline mean of 82 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%.[70] 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 can be accounted for theoretically. For instance, if women bear the full time cost of child care, or $\tau^m = 0$ in equation (1), then the 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. Assuming that men do make contributions to childcare, but that they are small, then the size of

69. The increase in the coefficients from Column 1 as controls are included is consistent with Figure 2 Panel A and Table 1, where higher potential yield municipalities had lower levels of fertility at baseline due to higher urbanization and incomes per capita.

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

the substitution effects from changes in male earnings will be small relative to the income effects. Consistent with this latter case, Schaller (2016) also 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 6 also estimates (9) using a measure of fertility directly from the Census: The log number of children under 5 in a municipality. Column 3 shows than an increase in one ton per hectare of potential soy increases the number of children under 5 by about 2%.[71] Finally, the last columns look at the share of women who are childless, measuring the extensive margin of fertility choice. A one unit change in potential soy leads to a .4 percentage point increase in the share of women who are childless, constituting a 1.3% increase. Thus, the main effects are occurring along the intensive margin. I return to this in Section 7.2.

I now exploit the longer time series data available for fertility to inspect pre-trends and trace 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. Examining the dynamics of the shock can distinguish between these possibilities.

Figure 3 plots the coefficients from estimating equation (11). Recall that these estimates trace the evolution of effects in municipalities with above median compared to below median increases in potential yields in a given year relative to 2002, the omitted base year prior to the legalization. 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.

Panel A uses the birth rate for women ages 16-49 as the dependent variable. We can see there is an absence of problematic pre-trends, supporting the identifying assumptions.[72] Examining the dynamic effects following the legalization, there is an ascent in fertility starting 3 years after the legalization of GE crops, which continues to grow into economically and statistically significant increases as time progresses. 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.

71. 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.

72. The F statistic for joint significance of pre- years is 1.4.

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 2.4 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. 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 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. The remaining Panels of Figure 3 plot the estimated coefficients on the potential soy measure in equation (11). Panel B uses adolescent fertility as the dependent variable, and Panel C and D use age-specific birth rates for the age ranges 20-29 and 30-39, respectively.[73] 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 C and D show that all of the movement in fertility is driven by the 20-29 and 30-39 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 both age ranges compared to below median municipalities. Panel C shows that in 2019, above median soy potential municipalities had 5.5 more births per 1000 women aged 20-29 compared to below median municipalities relative to the base year, which amounts to a 5% increase over the 2002 mean. Equivalently, Panel D shows in 2019 a 7.4% increase in above median compared to below median municipalities. Appendix Table 7 estimates equation (9) looking at age specific agricultural shares by gender, analogously defined as in Table 5, in order to see if the labor market outcomes align with the changes in the age specific rates. Panel A shows that for women, there is first no reallocation of any age group outside of the agricultural sector, and there is statistically significant 14% reduction in agricultural earnings for 30–39-year-olds and a 15% reduction in earnings for 40–49-year-olds. There is an economically large but imprecisely estimated decrease in earnings for 20–29-year-olds. For men, Panel B shows that male workers across all age

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

groups are leaving the agricultural sector, with a 5% increase in agricultural earnings for both 30-39- and 40–49-year-olds. The larger labor market effects for the 30-39 and older age ranges align with the proportionately larger fertility effects for the 30-39 age group.

The fertility dynamics presented in Figure 3 speak more to the longer run effects of this technological change. The sustained increase in older age-specific rates shown in Figure 3 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. 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.

In order to think about the magnitudes of these fertility effects, I turn to the 2017 Agricultural Census data and estimate a specification akin to (9), where I instrument for the share of farmland harvested with GE soy with the potential yields measure.[74] I find that a one standard deviation, or 11 percentage point, increase in the share of farmland harvested with GE soy increased fertility by 7.5 births per 1000 women.[75] 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 3.2 births per 1000 people. This is approximately equal to the total decline in fertility in sub-Saharan Africa between 2005 and 2017.[76] 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.[77]

Appendix Figure 4 shows that the results remain robust to weighting by the population of women aged 16-49 as well as using the log number births as the dependent variable (while also controlling for the log population of women). Finally, a recent literature documents the shortcomings of two-way fixed effects specifications. While the timing of adoption is not staggered, De Chaisemartin and d'Haultfoeuille (2020) shows that the underlying weighted sum of average treatment effects may still contain negative weights

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

75. The associated Kleibergen-Paap F-Statistic is 20.4.

76. 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

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

even in this setting. Reassuringly, I find that only 7% of weights are negative and that the sum of the negative weights is equal to -.018, suggesting that the regression coefficient is not subject to bias from negative weighting and heterogeneous treatment effects.[78]

7.1 Marriages

Marriage plays an important role in fertility choice and, as discussed in Section 3.1, on continued access to farm work. Appendix Table 11 shows that soy technical change led to a small movement from cohabitations towards marriages: A one ton per hectare increase in potential soy increases the share of those married by .47 percentage points, and lowers the share cohabiting by .7 percentage points, with no change in the share who are single or divorced. Further I utilize the demographic information from the administrative records on live births from SINASC to explore whether the fertility results are driven by marital or non-marital births. Columns 5 and 6 of Appendix Table 11 finds a one ton per hectare increase in potential soy leads to an increase in the marital birth rate of 4.3 births per 1000 women, with no change for non-marital births. While the overall impact on the marriage market is small, the increase in fertility is driven by marital births.

7.2 Child Quality

Appendix Section 3 discusses the theoretical considerations for changes in child quality in this context. The soy technology does not directly impact the returns to education. However, Bustos et al. (2019) find that the expansion of GE soy increased lower skilled manufacturing jobs (and reallocated workers away from more innovative manufacturing industries), which could indirectly lower the skill premium. However, increases in overall economic growth could induce increases in educational and health investments. On the other hand, it may lower these investments through the Quality-Quantity Trade-off, whereby increases in the number of children in the household make quality investments more costly.[79] Appendix Section 3 shows results from estimating equations (9) and (11) using measures of human capital and infant health (such as birth weight) as outcomes. I find no discernible impact on child quality, suggesting offsetting effects from these two channels. One may expect there to be an increase in fertility along the extensive margin, particularly if there is no change in the price of quality (e.g., Aaronson et al., 2014). As discussed earlier, there is an economically small increase in the share of women who

78. Checking the robustness of results with the alternative estimator in De Chaisemartin and d'Haultfoeuille (2020), I find effects consistent with (and even larger than) those estimated here. For instance, these suggest an increase in fertility of over 3 births per 1000 women, as compared to the 2.6 increase found in Table 7.

79. For example, increasing human capital or health investments is more costly if the household has to invest in 3 children versus 1 child. Thus the price of quality is an increasing function of the number of children in the household.

are childless. This then is consistent with the expansion of lower skilled manufacturing industries lowering the skill premium. However the small magnitudes along the extensive margin coupled with no discernible changes in child quality suggest that changes in the price of quality are likely not playing a major role.[80]

7.3 Other Channels

In addition to the opportunity cost and income effect channels, a large literature has explored other incentives for fertility choices. In Section 3 of the Appendix, I show that the soy technological change did not induce any changes in child labor usage, infant mortality, or migration. Further, Appendix Table 17 looks at the fertility of women who are 'stayers' defined earlier as those who have lived in the municipality continuously for at least 10 years. Given the limitations of fertility questions in the Census that fit this purpose, I look at the share of women who had a birth in the reference period, finding a statistically significant 4.2 percentage point increase.[81] Finally, increases in overall earnings may have eased credit constrained households who are now able to adjust to their desired level of fertility. If so, the impact of soy technical change on fertility would be smaller for households with more access to credit. To test this, I follow Costa et al. (2023) and proxy for access to credit by looking at the number of financial organizations, such as local commercial banks and rural credit cooperatives, per capita in a municipality using data from the registry of legal entities of the Brazilian Revenue Office (*Receita Federal*).[82] Appendix Table 9 reports the estimates from a modified version of equation (9) that adds the financial organizations per capita by itself and its interaction with the measure of soy technical change. I find no differential impacts of soy technical change on fertility or the other main earnings results, suggesting that credit constraints are not driving the results.[83] In total, these suggest that other channels outside of income and

80. The overall findings on fertility and quality are consistent with Mookherjee et al. (2012) who emphasize the importance of within and between occupational shifts in determining the fertility-wage correlation. In this setting, the positive income and substitution effects from increased household and reduced female earnings within agriculture combined with the low scope of human capital variation across sectors would drive a positive correlation between overall wages and fertility in municipalities.

81. Recall, the inclusion of state-year fixed effects adjusts for inter-state migration. Relevant for interpretation of the results, there is no differentially higher or lower migration (overall or gender-specific) between higher and lower soy suitability municipalities within states.

82. I follow Costa et al. (2023) and define financial organizations as those under activity codes 64212, 64221, 64239, 64247, 64336, 64344, 64361, and 64379. I take all the financial organizations that could provide access to credit that existed in 2000, prior to the legalization of GE soy.

83. Another possibility is that the GE soy increased the availability of time available to households, allowing them to have higher earnings while also now having more time to achieve their desired number of children. This argument may be more salient if the movements within agriculture were driven by family farms, who would be the residual claimant to the increased profits from the new technology. Recall that the employment effects driven by employees, suggesting that the effects are coming primarily from larger farms that hire wage labor, and that there were no reductions in total hours worked within agriculture. These suggest that women are not fully moving away from work towards enjoying more leisure or towards only child care.

substitution effects are not driving the fertility response.

8 Robustness and Other Results

The annual data in the analysis on fertility found no evidence of systematic pre-trends. This suggests that mechanization, any re-optimization of labor practices, or other factors that can generate the main labor market results that also drive fertility are not being captured by potential yields measure. Such a direct test is not possible with the labor market variables, given the changes in reference period and the definition of who is counted as employed, among other changes, that occurred starting in the 2000 Census (as mentioned in Section 4).[84] However, it is possible to construct comparable measures of employment across the 1980 and 1991 Census surveys, as these waves share the same methodology towards employment. Turning to the 1991 and 1980 Census data from the IBGE, I construct several measures of agricultural employment.[85] I first create a measure of the share of women and the share of men who are economically active, meaning they either had employment or were actively seeking work, across all sectors of the economy. To look more specifically within agriculture, I create the share of men and the share of women who are working in the agricultural sector, as well as the share of men and the share of women in agriculture who were not paid in their main occupation. Appendix Table 18 reports the estimates of equation (9) with these outcomes from 1980 to 1991. Reassuringly, the measure of potential soy has no statistically or economically significant effect on any of these outcomes.[86]

In Appendix Section 4, I report additional robustness checks. Appendix Table 19 provides further validation of the measure of soy technological change, showing that it only predicts the expansion of GE-Soy but not the expansion of non-GE-soy. Appendix Table 20 adds additional controls for market access. Since soy is often transported by 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. Appendix Table 21 shows that results are robust to clustering standard errors at higher levels of aggregation and using Conley (1999) standard errors with distance cutoffs of 50 and 200 km, and from relaxing the assumption that the municipality approximates the local markets by

84. See the discussion from Datazoom, a project from researchers at PUC-Rio to harmonize Brazilian data, for more detailed discussion on the differences.

85. The 1991 and 1980 Census survey data can be purchased in the form of CDs from the IBGE at `https://loja.ibge.gov.br/catalogsearch/result/?q=censo`

86. Recall also that all specifications in the main analysis include controls for changes in potential maize defined the same way (high technology potential yields minus the low potential yields) to capture the simultaneous expansion of maize technical change which included increased machinery. This would then control, to some extent, households finding ways to re-optimize labor and adopt machines generally over this time period that could be captured by the potential yields measure.

aggregating the main variables to a larger level of observation, the Mesoregion. Appendix Table 22 reports the main estimates weighted by population in 2000. All results are similar with these weights.

9 Discussion and Conclusion

I established that this agricultural technological change had different effects by gender. Specifically, they constituted a negative demand shock to women's work, with displaced women reallocating from paid to unpaid employment within the agricultural sector. Any men affected by the technologies moved into other sectors of the economy. Consistent with economic models of fertility, this labor-saving technological change led to increases in fertility.

These results indicate potential frictions preventing agricultural women from moving into non-agricultural employment. One such possibility is the low levels of education among agricultural workers. While low education is common to both men and women in agriculture (seen in Table 1), displaced women with higher levels of education might have the skills necessary to take up non-agricultural employment. Appendix Table 8 reports estimates from equation (9) using dependent variables analogously defined to those in Table 5, but split by education.[87] Thus, these shares are the number of workers of a given gender and education level in a sector, divided by the total number of workers of that gender and education level. The results show that neither higher or lower educated women are leaving the agricultural sector. On the other hand, both lower and higher educated men are able to move into non-agricultural employment.

Consistent with this, Alvarez (2020) finds that while the lower earnings in agriculture in Brazil is likely due to sorting, the agricultural wage gap cannot be fully explained by observed differences in education, suggesting other barriers for these workers. My results suggest that there is an important gender-specific component to these frictions.[88] As discussed in Section 3.1, one possibility is that the nature of women's work in agriculture does not equip them with transferable tacit knowledge or learning by doing that can facilitate movements outside of agriculture.[89]

87. While the 2000 Census asks directly about years of education, the 2010 Census only asks about ranges of education. To create a consistent measure across waves, I define low educated workers as those who have less than 10 years of education, which is 8th grade or less, and educated workers as those with more than 10 years of education. This is also similar to Bustos et al. (2019)'s definition of 'skilled' workers.

88. Further, Bustos et al. (2019) found that GE soy primarily expanded *lower-skilled* manufacturing sectors. It may be the case that men affected by these technologies are moving into low-skilled manufacturing sectors that may be more physically intensive and suited for male workers. Recall that Figure 1 shows that light industry, which heavily employs women, is not expanding much during this time period.

89. Further, Dix-Carneiro and Kovak (2017) show that imperfect regional mobility plays a key role in explaining the long run dynamics of labor market decline in response to lower trade barriers in Brazil. I find in Section 6.0.1 evidence of shifts between the manufacturing and services, which could be due

What does this mean for welfare? 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).[90] As noted in Section 3.0.1, women in Brazil tend to 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.[91] Work outside of economics also corroborates a negative welfare interpretation. Paris and Chi (2005) show that plastic row seeder technologies eliminated female-intensive weeding jobs in Vietnam, 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 underlying situation for women is similar to those in Brazil as discussed in Section 3.1.

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.

to regional mobility if service workers live in urban areas where new manufacturing opportunities arise. However, I leave further investigation of this to future work.

90. 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.

91. This argument is more salient if bargaining channels are present, whereby the technological change lowers the bargaining power of women.

References

Aaronson, Daniel, Fabian Lange, and Bhashkar Mazumder. 2014. "Fertility transitions along the extensive and intensive margins." *American Economic Review* 104 (11): 3701–3724.

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 microdata." *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.*

Arends-Kuenning, Mary, Akito Kamei, Marcos Garcias, Gisele Esser Romani, and Pery Francisco Assis Shikida. 2021. "Gender, education, and farm succession in Western Paraná State, Brazil." *Land Use Policy* 107:105453.

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.

Costa, Francisco, Angelo Marcantonio, and Rudi Rocha. 2023. "Stop suffering! Economic downturns and pentecostal upsurge." *Journal of the European Economic Association* 21 (1): 215–250.

Costinot, Arnaud, Dave Donaldson, and Cory Smith. 2016. "Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world." *Journal of Political Economy* 124 (1): 205–248.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American Economic Review* 110 (9): 2964–96.

Deere, Carmen Diana. 2005. *The feminization of agriculture? Economic restructuring in rural Latin America.* 1. UNRISD Occasional Paper.

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.

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 Bresnyan 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.

Kitchens, Carl T, and Luke P Rodgers. 2023. "The Impact of the WWI Agricultural Boom and Bust on Female Opportunity Cost and Fertility." *The Economic Journal* 133 (656): 2978–3006.

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.

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 Morais Neto, and Vinicius Pawlowski Queiroz. 2006. "Successful Brazilian Experiences in the Field of Health Information Final Report."

Marteleto, Letıcia 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.

Mookherjee, Dilip, Silvia Prina, and Debraj Ray. 2012. “A theory of occupational choice with endogenous fertility.” *American Economic Journal: Microeconomics* 4 (4): 1–34.

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.

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.

Paulilo, Maria Ignez S, et al. 2003. “Movimento de mulheres agricultoras: terra e matrimônio.” *Agricultura e espaço rural em Santa Catarina. Florianópolis: Ed. da UFSC:* 15–38.

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.

Pryor, Frederic L. 1985. “The invention of the plow.” *Comparative Studies in Society and History* 27 (4): 727–743.

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.

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.

Waltz, Amber. 2016. "The women who feed us: Gender empowerment (or lack thereof) in rural Southern Brazil." *Journal of rural studies* 47:31–40.

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: Demographic Characteristics by Sector (2000)

Panel A: Demographic Characteristics by Sector

	Agriculture	Manufacturing	Services
	(1)	(2)	(3)
Age: Female	38.32	36.1	33.11
Age: Male	37.01	32.23	35.69
Years of Education: Female	3.88	6.29	7.78
Years of Education: Male	3.84	5.73	7.02
Fraction Married: Female	0.51	0.47	0.42
Fraction Married: Male	0.46	0.43	0.51
Number of Children	4.01	3.13	2.76

Panel B: Demographic Characteristics by Employment Within Agriculture

	Employees	Self Employed	Not Paid
	(1)	(2)	(3)
Shares: Female	0.35	0.15	0.48
Shares: Men	0.51	0.29	0.19
Age: Female	32	36	33.52
Age: Men	30.33	34.79	26.63
Years of Education: Female	4.14	4.05	3.83
Years of Education: Men	3.53	4.45	4.71
Fraction Married: Female	0.41	0.55	0.59
Fraction Married: Men	0.37	0.57	0.19
Number of Municipalities: 4254			

Table shows summary statistics where the unit of observation are municipalities in Brazil. Data are from the sample supplement to the 2000 Brazilian Census. Employees are those who report working as an employee in the main occupation with either a formal or informal contract in their main economic activity. Self employed are those who report being self employed or an autonomous worker in their main economic activity. Unpaid work includes helping another household member with their work or working in cultivation for a household's sustenance for no pay.

Table 2: High Versus Low Soy Potential Yields

	Below Median	Above Median	Within State Differences
	(1)	(2)	(3)
	Mean Values		
Share of Rural Population	0.51	0.4	-0.08***
Female to Male Ratio of Literacy Rates	1.05	1.15	0.02***
Log Population Density	3.17	3.23	0.03
Log Income Per Capita	4.39	4.65	0.09***
% of Children in Low Income Households	89.27	84.94	-0.94***
Number of Municipalities	2127	2127	

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

Table shows differences in average baseline characteristics recorded in 1991 between above and below median soy potential municipalities. Column (3) presents the coefficient estimate and standard errors from a regression of the the characteristic on a dummy variable for being above the median potential yield measure, controlling for state fixed effects. Thus this represents the average difference *within* states.

Table 3: Partial Correlations in the Data

Panel A: 2006

VARIABLES	(1) Share GE Soy (%)	(2) Share GE Soy (%)	(3) Fertility	(4) Fertility
ΔPot. Soy	1.228***	1.162***		
	(0.405)	(0.324)		
GE Soy (%)			-0.0634*	-0.0420
			(0.0331)	(0.0357)
Observations	4,181	4,181	4,181	4,181
R-squared	0.342	0.354	0.427	0.453
Controls	NO	YES	NO	YES
State FE	YES	YES	YES	YES

Panel B: 2017

VARIABLES	(1) Share GE Soy (%)	(2) Share GE Soy (%)	(3) Fertility	(4) Fertility
ΔPot. Soy	2.706***	3.713***		
	(0.489)	(0.837)		
GE Soy (%)			0.0579***	0.0631***
			(0.0178)	(0.0183)
Observations	4,192	4,192	4,192	4,192
R-squared	0.365	0.375	0.290	0.298
Controls	NO	YES	NO	YES
State FE	YES	YES	YES	YES

Robust standard errors clustered at the Microregion level are in parentheses
*** 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 and show partial correlations between actual GE soy and fertility. Panel A uses data from the 2006 Agricultural Census (table 824) and Panel B from the 2017 Agricultural Census (table 6958). 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. The dependent variable in Columns (1) and (2) in both Panels is the percent of all harvested area with GE soy. In Columns (3) and (4), the dependent variable is fertility, measured as the number of births per 1000 women aged 16-49. All regressions include State fixed effects, which make this a comparison between municipalities in the same state. Columns (2) and (4) include all controls from Table 2 as well as an analogous measure of potential maize, adjusting for the simultaneously adoption of other crop technologies.

Table 4: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Agricultural Earnings and Hours

VARIABLES	(1) Δ Female Earnings	(2) Δ Female Hours	(3) Δ Male Earnings	(4) Δ Male Hours	(5) Δ Family Earnings
ΔPot. Soy	-0.110**	-0.433	0.0295	0.0338	0.0274*
	(0.0509)	(0.411)	(0.0198)	(0.203)	(0.0158)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.098	0.026	0.079	0.059	0.072
Controls	ALL	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES	YES
$\bar{Y}_{2000}$	218.22	37.20	508.81	45.03	893.33

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 number of hours worked by women in the main work within agriculture. The third column uses the log of the average male municipality agricultural earnings; the fourth column uses the number of hours worked by men in the main work within agriculture, and the fifth column uses the total households earnings of families working in agriculture. Families are defined as groups of two or more people related by birth or marital status and residing together, such as an individual living with their spouse or partner. Data are taken at the individual level, and aggregated to the municipality using weights provided by the IBGE. 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 5: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Agricultural and Non-Agricultural Employment

VARIABLES	(1) Δ Female Agricultural Shares	(2) Δ Female Non-Agricultural Shares	(3) Δ Male Agricultural Shares	(4) Δ Male Non-Agricultural Shares
ΔPot. Soy	-0.00569	0.00214	-0.0275***	0.0216***
	(0.00477)	(0.00503)	(0.00463)	(0.00419)
Observations	4,254	4,254	4,254	4,254
R-squared	0.107	0.117	0.150	0.145
Controls	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.21	0.77	0.44	0.41

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 Columns (1) and (2), the dependent variables are the share of women in specific sector i.e. the numerator is the number of women working for in a sector, and the denominator is total female employment. Columns (3) and (4) report the same for men. 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 6: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Agricultural Outcomes

Panel A: Female Agricultural Employment

VARIABLES	(1) Δ Share Employees	(2) Δ Employee Earnings	(3) Δ Share Self Employed	(4) Δ Self Employed Earnings	(5) Δ Share Unpaid
ΔPot. Soy	-0.0376***	-0.0208	-0.0147	0.170	0.0483***
	(0.0118)	(0.0612)	(0.00905)	(0.114)	(0.0143)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.080	0.054	0.168	0.060	0.217
Controls	ALL	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.35	301.89	0.15	365.59	0.48

Panel B: Male Agricultural Employment

VARIABLES	(1) Δ Share Employees	(2) Δ Employee Earnings	(3) Δ Share Self Employed	(4) Δ Self Employed Earnings	(5) Δ Share Unpaid
ΔPot. Soy	-0.0195***	0.0552***	-0.00568	0.0282	0.0253***
	(0.00619)	(0.0140)	(0.00589)	(0.0319)	(0.00532)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.109	0.121	0.085	0.028	0.105
Controls	ALL	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.51	428.09	0.29	797.38	0.19

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 in Column 1 and 2 of Panel A are the shares and the earnings of female employees at agricultural establishments with a work contract. Column 3 and 4 use the share and earnings of self-employed women in agriculture. Column 5 uses the share of women whose main economic activity is unpaid agricultural work helping another household member with their work or for cultivation for household sustenance. Panel B is analogous for men. All earnings variables are the logs of the average earnings. 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 7: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Birth Rates per 1000 Women

VARIABLES	(1) Δ Fertility	(2) Δ Fertility	(3) Δ Fertility
ΔPot. Soy	1.677***	2.750***	2.661**
	(0.473)	(0.984)	(1.050)
Observations	4,254	4,254	4,254
R-squared	0.136	0.140	0.147
Controls	No	Maize and Rural	All
State FE	YES	YES	YES
$\bar{Y}_{2000}$	66.24	66.24	66.24

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 where the dependent variable is the number of live births per 1000 women aged 16-49. The mean of the dependent variable in 2000 is reported in the last row. Data on live births are taken from administrative records from SINASC, and population projections from DATASUS. The first column contains no controls and only state fixed effects, the second includes controls for potential maize and the percent of the population that is rural, and the last column includes all controls and municipality and state by year fixed effects.

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

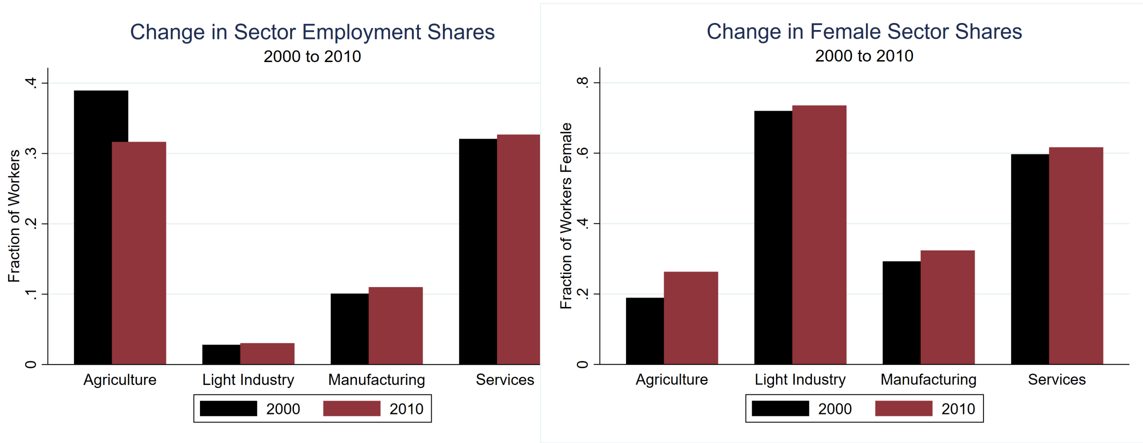

The first figure reports the average employment share in 2000 and 2010 across municipalities in a given sector. The second shows the share of female workers aged 15-55 by sector. The numerator is the total number of women working in that sector, and the denominator is the total number of workers (men and women). Data are taken from the sample supplement of the Demographic Census from 2000 and 2010. Each bar represents an average across municipalities (N=4254)

Figure 2: Soy Potential Yields

Panel A Evolution of Fertility: High and Low Potential Soy

Panel B Spatial Suitability for Soy Technological Change

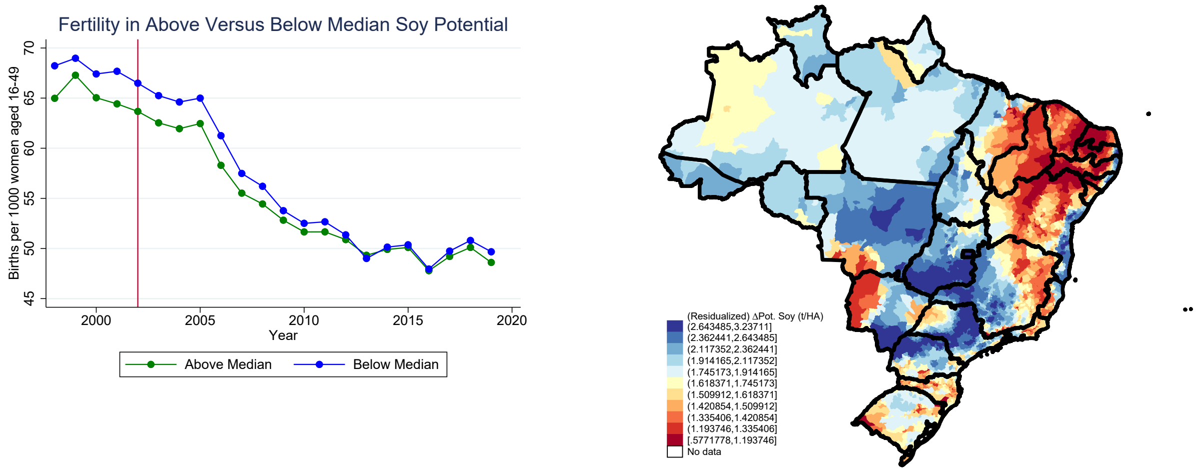

Panel A shows the evolution of fertility in above and below median potential yields over time. Each point represents the average fertility rate in a given year. The above and below median potential yield groups are calculated after partialling out state fixed effects from the measure of potential yields taken from the FAO-GAEZ database. Panel B presents a map showing potential yields by municipality with municipality borders omitted. Dark bold outlines represent the Brazilian states. Darker shades of blue represent higher potential yields and darker shades of red represent lower potential soy yields. The soy yields partial 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 the Effect of Soy Potential Yields on Age-Specific Fertility

Panel A: Ages 16-49

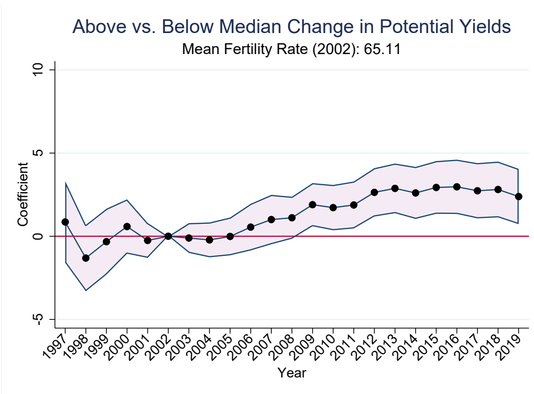

Panel B: Adolescent Fertility

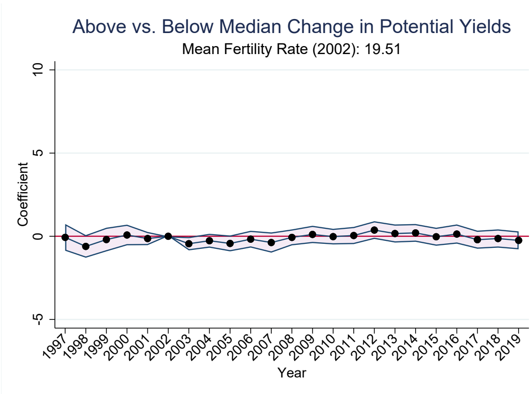

Panel C: Ages 20-29

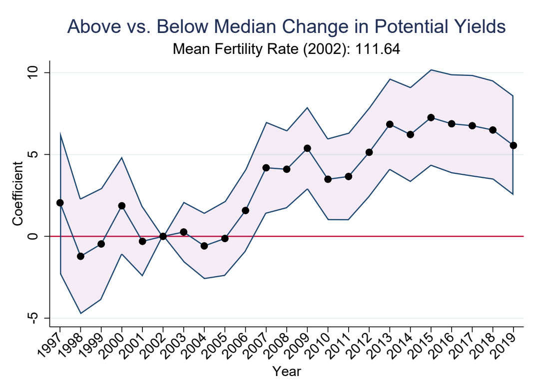

Panel D: Ages 30 - 39

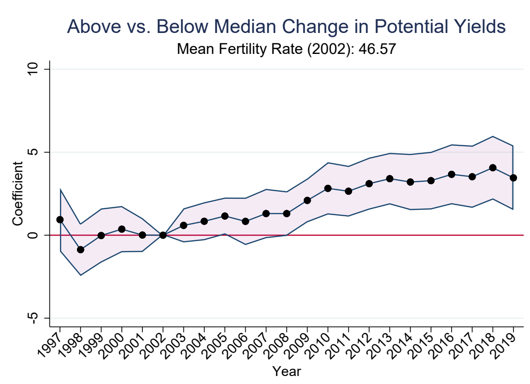

These sub-figures plots estimated coefficients from (11). Panel A uses the number of live births per 1000 women aged 16-49 as the dependent variable. Panel B uses the adolescent birth rate as the dependent variable, defined as the number of live births per 1000 women aged 10-19, and the remaining panels C and D 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.

Online 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, who counts as an employed worker, 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. 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. One may worry about whether these zeros are small data issues. The Census survey is the only source of data in Brazil that is representative at the municipality level. Examination of the 21 municipalities in the regression sample that are recorded as zero female earnings shows that these tend to be very small and highly service oriented municipalities. For example, these municipalities on average have 8000 women total, compared to the overall average being closer to 20,000. As a specific example, one of the municipalities is Nova Aurora in Goiás, which had a total population of about 1927 in 2000, with 50% of its workforce working in a service oriented industry.

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. The last two are also aggregated to form the 'non'-agricultural' sector. 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.

A.1: Summary Statistics

Panel A: Fertility					
Variable (2000)	Mean	SD	2010-2000 Mean	SD	N
Fertility Rate (Births per 1000 women aged 16-49)	66.24	19.71	-14.14	17.11	4254

Panel B: Monthly Municipality Earnings by Sector (in 2010 Reals)					
Variable (2000)	Mean	SD	Mean Log Difference: 2010-2000	SD	N
Female Earnings (Agriculture)	218.22	316.32	0.57	1.24	4254
Female Earnings (Manufacturing)	458.23	414.62	0.37	1.38	4254
Female Earnings (Services)	577.1	232.47	0.26	0.25	4254
Male Earnings (Agriculture)	508.81	388.46	0.21	0.39	4254
Male Earnings (Manufacturing)	802.2	732.42	0.16	0.7	4254
Male Income (Services)	1135.5	544.05	0.05	0.3	4254

Panel C: 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

Panel D: Instrument: Potential Yields			
Variable (1991)	Mean	SD	N
Potential Soy	1.8	0.85	4254
Potential Maize	3.06	1.81	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 16-49. Data on live births come from live birth certificates compiled by SINASC, and the population counts come from projections provided by DATASUS.

A.1: Evolution of Fertility in Above and Below Median Potential Yield: No State Fixed Effects

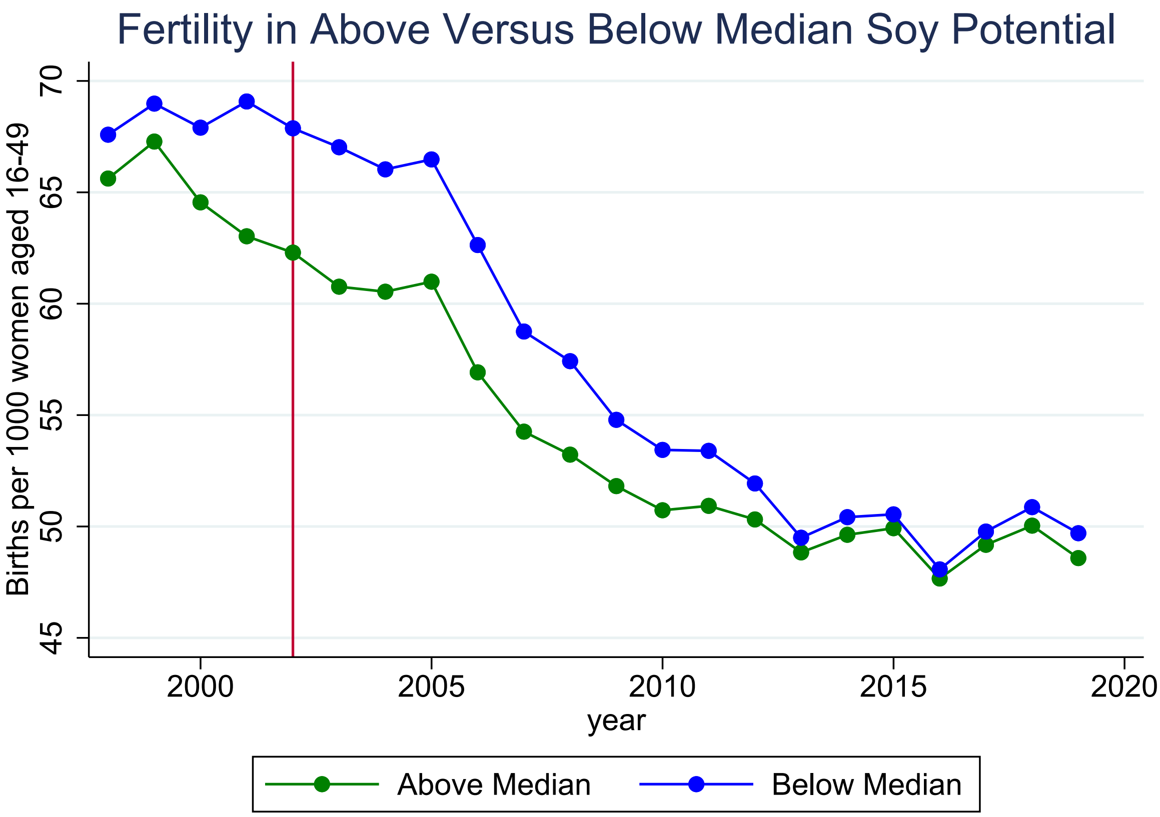

Section 2: Context in Brazil

President Luiz Inácio Lula da Silva was initially opposed to GE crops when elected in 2002 (Benthem, 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). [92] 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.[93]

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 A2 shows that despite this small (albeit still significant) share of total production of soybeans, family farms constituted about 76% of all farms that produced soy.

92. 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.

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

Soy production in Brazil is often associated with the remarkable development and large farms of the Cerrado region,[94] 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 A2 shows that a majority of all soy farms are family farms in the South region of Brazil.

A.2: 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.

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 a 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 3 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.

94. The Cerrado lands largely incorporate the Centerwest region, which includes the state of Mato Gosso.

A.3: 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.

Section 3: Further Results

Aggregate Effects

A.4: Overall Labor Market and Structural Transformation

Panel A: Overall Labor Market Outcomes

	(1)	(2)	(3)	(4)	(5)
	Δ Female Unemployment	Δ Male Unemployment	Δ Female Labor Force Participation	Δ Male Labor Force Participation	Δ Overall Family Earnings
ΔPot. Soy	-0.000911	0.00279	0.00312	0.00171	0.0276***
	(0.00358)	(0.00228)	(0.00236)	(0.00223)	(0.00896)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.105	0.229	0.146	0.112	0.144
Controls	ALL	ALL	ALL	ALL	ALL
$\bar{Y}_{2000}$	0.17	0.09	0.31	0.55	1185.48

Panel B: Sectoral Outcomes

	(1)	(2)	(3)	(4)	(5)
	Δ Agricultural Employment Shares	Δ Manufacturing Employment Shares	Δ Services Employment Share	Δ Unpaid Agricultural Work	Δ Overall Agricultural Earnings
ΔPot. Soy	-0.0212***	0.0162***	0.00225	0.0244***	0.0270
	(0.00415)	(0.00346)	(0.00207)	(0.00687)	(0.0203)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.142	0.155	0.185	0.195	0.096
Controls	ALL	ALL	ALL	ALL	ALL
$\bar{Y}_{2000}$	0.39	0.10	0.32	0.25	450.22

Standard Errors Clustered at the Microregion level are reported in parenthesis

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

A.5: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Labor Market Outcomes: Other Sectors

Panel A: Female Sectoral Reallocations

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ Share in Agriculture	Δ Share in Manufacturing	Δ Share in Services	Δ Manufacturing Earnings	Δ Services Earnings
ΔPot. Soy	-0.00569	0.0117***	-0.00953**	0.0380	0.0288***
	(0.00477)	(0.00324)	(0.00452)	(0.0477)	(0.00938)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.107	0.113	0.117	0.014	0.211
Controls	ALL	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.21	0.09	0.68	458.23	577.10

Panel B: Male Sectoral Reallocations

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ Share in Agriculture	Δ Share in Manufacturing	Δ Share in Services	Δ Manufacturing Earnings	Δ Services Earnings
ΔPot. Soy	-0.0275***	0.0200***	0.00160	0.00584	0.0209*
	(0.00463)	(0.00423)	(0.00258)	(0.0253)	(0.0123)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.150	0.175	0.193	0.022	0.099
Controls	ALL	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.44	0.11	0.30	802.19	1135.50

Robust standard errors clustered at the Microregion level are in parentheses

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

A.6: Additional Fertility Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
VARIABLES	Δ *Log*(Less than 5)	Δ *Log*(Less than 5)	Δ *Log*(Less than 5)	Δ Fertility:20-40	Δ Fertility:20-40	Δ Fertility:20-40	Δ Childless	Δ Childless	Δ Childless
ΔPot. Soy	0.0294***	0.0245***	0.0159*	2.913***	4.102***	3.973***	0.00268***	0.00574***	0.00401***
	(0.00574)	(0.00947)	(0.00870)	(0.615)	(1.241)	(1.318)	(0.000874)	(0.00150)	(0.00144)
Observations	4,254	4,254	4,254	4,254	4,254	4,254	4,254	4,254	4,254
R-squared	0.164	0.218	0.249	0.129	0.129	0.136	0.178	0.205	0.218
Controls	NO	Maize and Rural	All	NO	Maize and Rural	All	NO	Maize and Rural	All
State FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
$\bar{Y}_{2000}$	4112.39	4112.39	4112.39	82.38	82.38	82.38	0.31	0.31	0.31

Robust standard errors clustered at the Microregion level are in parentheses

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

A.7: First Difference Estimates of the Effects of Soy Potential Yields on Municipality Level Earnings and Employment by Ages

Panel A: Female Agricultural Labor Market: By Age

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES	Δ Share in Agr, 10- 19	Δ Agricultural Earnings, 10-19	Δ Share in Agr, 20-29	Δ Agricultural Earnings, 20-29	Δ Share in Agr, 30-39	Δ Agricultural Earnings, 30-39	Δ Share in Agr, 40-49	Δ Agricultural Earnings, 40-49
ΔPot. Soy	-0.00159	-0.194	0.00209	-0.131	-0.00414	-0.147**	0.00112	-0.168**
	(0.00183)	(0.119)	(0.00305)	(0.0919)	(0.00341)	(0.0748)	(0.00341)	(0.0758)
Observations	4,238	4,238	4,238	4,238	4,238	4,238	4,238	4,238
R-squared	0.066	0.105	0.072	0.067	0.073	0.066	0.083	0.058
Controls	YES	YES	YES	YES	YES	YES	YES	YES

Panel B: Male Agricultural Labor Market: By Age

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES	Δ Share in Agr, 10- 19	Δ Agricultural Earnings, 10-19	Δ Share in Agr, 20-29	Δ Agricultural Earnings, 20-29	Δ Share in Agr, 30-39	Δ Agricultural Earnings, 30-39	Δ Share in Agr, 40-49	Δ Agricultural Earnings, 40-49
ΔPot. Soy	-0.0103***	-0.0281	-0.0275***	-6.90e-05	-0.0237***	0.0532**	-0.0115***	0.0507**
	(0.00351)	(0.0456)	(0.00533)	(0.0252)	(0.00501)	(0.0239)	(0.00421)	(0.0250)
Observations	4,238	4,238	4,238	4,238	4,238	4,238	4,238	4,238
R-squared	0.161	0.149	0.163	0.077	0.146	0.058	0.133	0.066
Controls	YES	YES	YES	YES	YES	YES	YES	YES

Robust standard errors clustered at the Microregion level are in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

A.8: Sectoral Reallocations by Skill

Panel A: Female Sectoral Reallocations by Education

	(1)	(2)	(3)	(4)
VARIABLES	Δ Share LE in Agricultural Employment	Δ Share LE in Non-Agricultural Employment	Δ Share HE in Agricultural Employment	Δ Share HE in Non-Agricultural Employment
ΔPot. Soy	-0.000953	0.00177	-4.68e-05	0.00583
	(0.00330)	(0.00299)	(0.00269)	(0.00411)
Observations	4,254	4,254	4,254	4,254
R-squared	0.077	0.275	0.172	0.408
Controls	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.11	0.27	0.02	0.65

Panel B: Male Sectoral Reallocations by Education

	(1)	(2)	(3)	(4)
VARIABLES	Δ Share LE in Agricultural Employment	Δ Share LE in Non-Agricultural Employment	Δ Share HE in Agricultural Employment	Δ Share HE in Non-Agricultural Employment
ΔPot. Soy	-0.0227***	0.0191***	-0.0111***	0.00932*
	(0.00475)	(0.00377)	(0.00427)	(0.00565)
Observations	4,254	4,254	4,254	4,254
R-squared	0.134	0.141	0.092	0.193
Controls	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.36	0.29	0.12	0.63

Robust standard errors clustered at the Microregion level are in parentheses

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

A.9: Credit Constraints

	(1)	(2)	(3)	(4)
	Δ Fertility	Δ Family Earnings	Δ Female Earnings	Δ Male Earnings
ΔPot. Soy	2.541**	0.0191	-0.170***	0.0366*
	(1.169)	(0.0190)	(0.0590)	(0.0219)
ΔPot. Soy x Financial Institutions	-0.114	0.0208	-0.0296	0.00631
	(1.524)	(0.0321)	(0.162)	(0.0338)
Observations	4,254	4,254	4,254	4,254
R-squared	0.150	0.074	0.126	0.079
Controls	ALL	ALL	ALL	ALL

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

Fertility: Regional Heterogeneity

Appendix Table 2 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.[95] We would expect to see most of the increases in fertility driven by the more soy intensive regions. Figure 2 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

95. 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).

family farms, are 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.

A.2: Event Studies: Regional Heterogeneity

Panel A.Fertility Rate: Soy Intensive Regions)

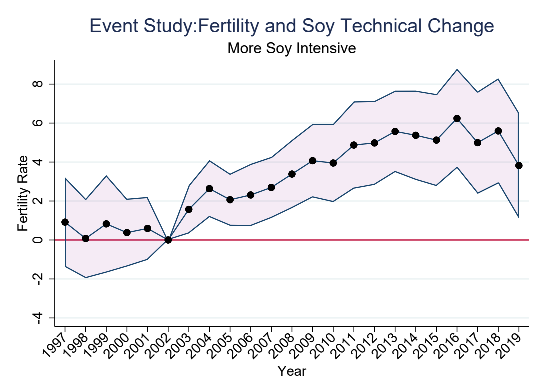

Panel B. Fertility Rate: Less Soy Intensive

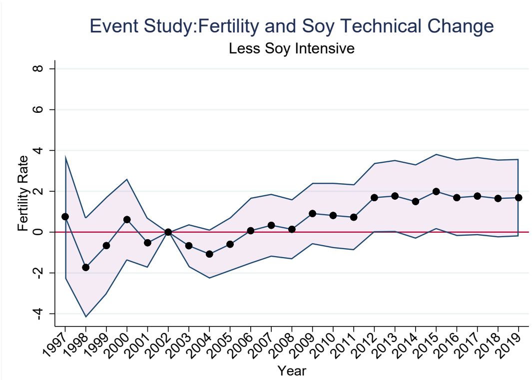

Panel C Fertility Rate: South Region

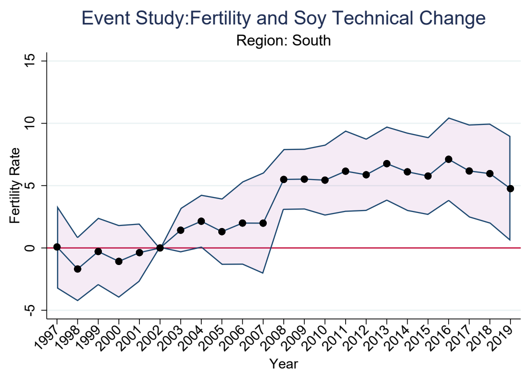

Panel D. Fertility Rate: Centerwest Region

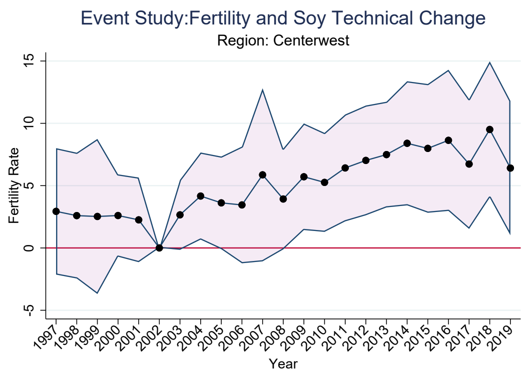

These sub-figures plots estimated coefficients from (11) where the dependent variable birth rate, defined as the number of live births per 1000 women aged 16-49, split by regions. Panels A and B run the regressions on the Soy Intensive (Southern and Centerwest) and Less Soy Intensive (remaining regions) 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.

Additional Channels

Marital Outcomes

A.11: First Difference Estimates of the Effects of Soy Potential Yields on Marital Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Married	Δ Cohabited	Δ Single	Δ Divorced	Δ Marital Fertility	Δ Non-Marital Fertility
ΔPot. Soy	0.00469**	-0.00744***	0.00172	-0.000137	4.330***	-0.310
	(0.00228)	(0.00217)	(0.00130)	(0.000221)	(1.548)	(1.575)
Observations	4,254	4,254	4,254	4,254	4,254	4,254
R-squared	0.267	0.268	0.184	0.160	0.152	0.163
Controls	All	All	All	All	All	All
$\bar{Y}_{2000}$	0.41	0.17	0.34	0.01	75.29	70.27

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 where the dependent variable is the share of individuals aged 16-49 of a given marital status. The dependent variable in the first column is the share who are married; the second uses the share who are cohabited; the third column uses the share who are single, and the fourth column uses the share who are divorced. These variables are constructed from the sample supplement to the Demographic Census. The final two columns construct the number of births per 1000 married women aged 16-49 and the number of births per 1000 unmarried women aged 16-49, respectively. The number of births by marital status are taken from administrative records from SINASC, and the population counts of each marital status are taken from the Census. The mean of the dependent variable in 2000 is reported in the last row.

	(1)	(2)	(3)	(4)	(5)
	Δ Husbands: Agriculture	Δ Husbands: Manufacturing	Δ Husbands: Services	Δ Age: Female	Δ Age: Male
ΔPot. Soy	-0.0263***	0.0142***	0.00462	-0.115**	-0.138**
	(0.00517)	(0.00406)	(0.00343)	(0.0499)	(0.0593)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.125	0.158	0.167	0.156	0.138
Controls	ALL	ALL	ALL	ALL	ALL
$\bar{Y}_{2000}$	0.42	0.11	0.33	34.55	36.56

Robust standard errors clustered at the Microregion level are in parentheses

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

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.[96] 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 12 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.

Migration

Here, I consider the possibility that the adoption of these technologies induced by favorable soil and weather characteristics differentially effected 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.[97] 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 12 show the results of estimating the equation (9) with these outcomes. I find no significant effects on any of these measures.[98]

96. 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.

97. 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.

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

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). 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.[99] I use data from 1998-2019.[100] 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.[101] Panel A of Figure 3 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.

99. This data can be obtained from DATASUS.

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

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

A.12: Testing Other Channels that Influence Fertility

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Child Labor	Δ Agr. Child Share	Δ *log*(Female Migrants)	Δ *log*(Male Migrants)	Δ *log*(Female Migrants$^{10-49}$)	Δ *log*(Male Migrants$^{10-49}$)
ΔPot. Soy	0.000266	0.00236	0.0120	0.0191	0.0166	0.0176
	(0.00312)	(0.00170)	(0.0162)	(0.0215)	(0.0165)	(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.

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 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.[102]

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.[103]

Table 13 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

102. 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 Kubitza (2021), where the shocks studied induced increases in the returns to education, either directly or indirectly.

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

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.[104]

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).

Panel B of Figure 3 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.[105]

104. 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.

105. 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.

A.13: First Difference Estimates of Potential Soy on Child Quality Investments

	(1)	(2)	(3)	(4)	(5)
	Δ Literate	Δ Attendance	Δ High School	Δ Unpaid Work	Δ Cultivation Help
ΔPot. Soy	0.00106	0.00149	-0.00120	0.00337	-0.00109
	(0.00135)	(0.00206)	(0.00432)	(0.00235)	(0.00113)
Observations	4,254	4,254	4,254	4,254	4,254
R-squared	0.549	0.200	0.149	0.151	0.133
Controls	ALL	ALL	ALL	ALL	ALL
StateFE	YES	YES	YES	YES	YES
$\Delta\bar{Y}$	0.03	0.04	0.18	-0.01	0.01
$\bar{Y}_{2000}$	0.93	0.92	0.18	0.06	0.01

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

A.3: Event Study Estimates of the Effect of Soy Potential Yields on Infant Health Outcomes

Panel A Infant Mortality

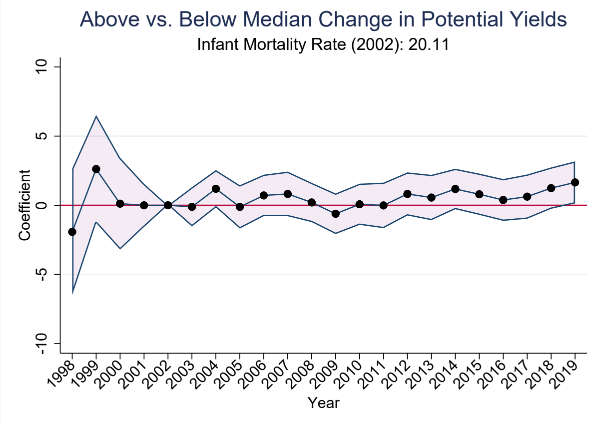

Panel B. Fraction Low Birthweight

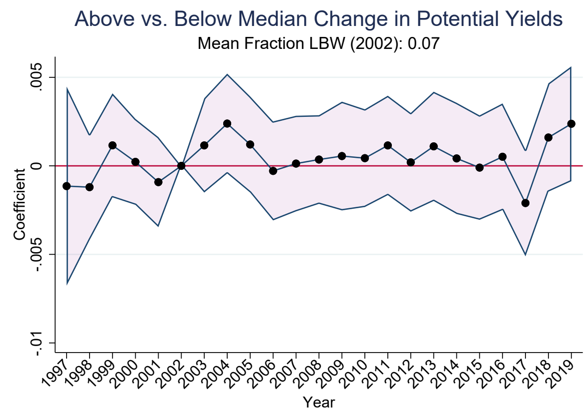

These sub-figures plots estimated coefficients from (11) where the dependent variable in the first figure is the infant mortality rate, defined as the number of deaths of children less than 1 years old per 1000 live births, and the second is fraction of children with low birth weight, where low birth 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.

Maize Technical Change

As mentioned in Section 3.2, the other major crop technologies adopted over this time period were those in maize. These included the increased use of machinery and of insect repellent GE maize. Bustos et al. (2016) explicitly analyze the economic consequences of maize technological change, finding that the expansions of maize technologies facilitated the growing of two seasons of maize in one year, effectively increasing the land endowment, and leading to a reallocation of resources into the agricultural sector. Multiple technologies created these changes and there is no discrete date to use to create a plausibly exogenous interaction of the timing of adoption with the geo-spatial potential maize measures. However, it may still be of interest to explore the potential gendered effects of these technologies. An increase in a second harvest of maize can increase the demand for tasks such as weeding and transplanting which are typically performed by women.[106]

Table 14 reports the main estimates including the coefficient on the measure of technical change in maize. The labor market results generally have the opposite effect as GE soy: There is an increase in the employment of both men and women in agriculture, specifically as employees for other farms. While positively signed, there is no statistically significant effect on the agricultural earnings of women, and ultimately, no change in fertility. These results do suggest however that understanding how technological change across different crops interacts with the gender-specific occupational sorting is vital for anticipating how future agricultural technologies may impact economic opportunities for men and women, as well as how they may impact family formation.

106. In a similar argument, Schultz (2001) posits that irrigation "permitted multiple cropping of the land in each year, raising the share of labor required for weeding and transplanting, tasks for which female labor may be more productive than men" (pg. 433). Moreover, note that Pryor (1985) lists maize as a 'plough-negative' crop, and thus we may expect there to be a relatively high proportion of female employment in maize cultivation (cf. Alesina et al., 2013).

A.14: Maize Technical Change

Panel A: Agricultural Earnings and Fertility			
	(1)	(2)	(3)
VARIABLES	Δ Female Earnings	Δ Male Earnings	Δ Fertility
ΔPot. Soy	-0.110**	0.0295	2.661**
	(0.0509)	(0.0198)	(1.050)
ΔPot. Maize	0.0161	-0.0111	-0.385
	(0.0230)	(0.0101)	(0.490)
R-squared	0.098	0.079	0.147

Panel B: Female Sectoral Reallocations			
	(1)	(2)	(3)
VARIABLES	Δ Female Share Agr	Δ Female Share Manu	Δ Female Share Ser
ΔPot. Soy	-0.00569	0.0117***	-0.00953**
	(0.00477)	(0.00324)	(0.00452)
ΔPot. Maize	0.00400*	-0.00398***	0.00148
	(0.00232)	(0.00146)	(0.00219)
R-squared	0.107	0.113	0.117
Observations	4,254	4,254	4,254
Controls	YES	YES	YES
State FE	YES	YES	YES

Robust standard errors clustered at the Microregion level are in parentheses

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

A.15: Maize Technical Change

Panel C: Male Sectoral Reallocations			
	(1)	(2)	(3)
VARIABLES	Δ Male Share Agr	Δ Male Share Manu	Δ Male Share Serv
ΔPot. Soy	-0.0275***	0.0200***	0.00160
	(0.00463)	(0.00423)	(0.00258)
ΔPot. Maize	0.00770***	-0.00311	-0.00207
	(0.00227)	(0.00202)	(0.00129)
R-squared	0.150	0.175	0.193

Panel D: Female Employment in Agriculture			
	(1)	(2)	(3)
VARIABLES	Δ Female Employees Agr	Δ Female Self Employed Agr	Δ Female Not Paid Agr
ΔPot. Soy	-0.0376***	-0.0147	0.0483***
	(0.0118)	(0.00905)	(0.0143)
ΔPot. Maize	0.0117**	0.00258	-0.0158**
	(0.00572)	(0.00426)	(0.00692)
R-squared	0.080	0.168	0.217

Panel E: Male Employment in Agriculture			
	(1)	(2)	(3)
VARIABLES	Δ Male Employees Agr	Δ Male Self Employed Agr	Δ Male Unpaid Agr
ΔPot. Soy	-0.0195***	-0.00568	0.0253***
	(0.00619)	(0.00589)	(0.00532)
ΔPot. Maize	0.00769**	-0.00198	-0.00555**
	(0.00314)	(0.00288)	(0.00262)
R-squared	0.109	0.085	0.105
Controls	YES	YES	YES
State FE	YES	YES	YES
Observations	4,254	4,254	4,254

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

Section 4: Robustness

A.16: Municipality Level Earnings: Sensitivity

Panel A: Female Agricultural Earnings			
	(1)	(2)	(3)
VARIABLES	Δ Female Earnings	Δ Female Earnings	Δ Female Earnings
ΔPot. Soy	-0.120***	-0.113**	-0.110**
	(0.0325)	(0.0504)	(0.0509)
Observations	4,254	4,254	4,254
R-squared	0.087	0.097	0.098
Controls	NO	Maize and Rural	ALL
State FE	YES	YES	YES
$\bar{Y}_{2000}$	218.22	218.22	218.22

Panel B: Female Agricultural Hours			
	(1)	(2)	(3)
VARIABLES	Δ Female Hours	Δ Female Hours	Δ Female Hours
ΔPot. Soy	-0.0535	-0.463	-0.433
	(0.232)	(0.408)	(0.411)
Observations	4,254	4,254	4,254
R-squared	0.021	0.025	0.026
Controls	NO	Maize and Rural	ALL
State FE	YES	YES	YES
$\bar{Y}_{2000}$	37.20	37.20	37.20

Panel C: Male Agricultural Earnings			
	(1)	(2)	(3)
VARIABLES	Δ Male Earnings	Δ Male Earnings	Δ Male Earnings
ΔPot. Soy	0.00391	0.0117	0.0295
	(0.0108)	(0.0207)	(0.0198)
Observations	4,254	4,254	4,254
R-squared	0.060	0.062	0.079
Controls	NO	Maize and Rural	ALL
State FE	YES	YES	YES
$\bar{Y}_{2000}$	508.81	508.81	508.81

Panel D: Male Agricultural Hours			
	(1)	(2)	(3)
VARIABLES	Δ Male Hours	Δ Male Hours	Δ Male Hours
ΔPot. Soy	-0.139	-0.0591	0.0338
	(0.113)	(0.204)	(0.203)
Observations	4,254	4,254	4,254
R-squared	0.040	0.054	0.059
Controls	NO	Maize and Rural	ALL
State FE	YES	YES	YES
$\bar{Y}_{2000}$	45.03	45.03	45.03

Panel E: Household Agricultural Earnings			
	(1)	(2)	(3)
VARIABLES	Δ Family Earnings	Δ Family Earnings	Δ Family Earnings
ΔPot. Soy	0.00629	0.0121	0.0274*
	(0.00890)	(0.0166)	(0.0158)
Observations	4,254	4,254	4,254
R-squared	0.041	0.054	0.072
Controls	NO	Maize and Rural	ALL
State FE	YES	YES	YES
$\bar{Y}_{2000}$	893.33	893.33	893.33

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

A.17: First Difference Estimates of the Effect of Potential Soy on Agricultural Earnings: Long Term Residents

	(1)	(2)	(3)	(4)
VARIABLES	Δ Female Earnings	Δ Male Earnings	Δ Family Earnings	Δ Birth in Reference Period
ΔPot. Soy	-0.158***	0.00379	0.0108	0.0418***
	(0.0414)	(0.0118)	(0.00673)	(0.00899)
Observations	4,254	4,254	4,254	4,254
R-squared	0.028	0.017	0.095	0.123

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 log of the average female municipality agricultural earnings. The second column does the same using the average male municipality earnings. The fourth column uses overall family earnings for households in agriculture. The first column uses the share of women who had a birth in the reference period. 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.

The Measurement of Soy Technical Change

In this section, I provide further validation of the measure of soy technological change. Specifically, I test whether the potential soy measure is picking up other factors that may be driving the main results besides GE soy, as the high technology regime for potential yields from FAO-GAEZ includes inputs such as increased mechanization. Recall that all specifications include controls for changes in potential maize, which are defined the same way (high technology potential yields minus the low potential yields), to capture the simultaneous expansion of maize technical change which included increased machinery. This would then control, to some extent, households finding ways to re-optimize labor and adopt machines generally over this time period that could be captured by the potential yields measure. However, this may not fully adjust for similar changes happening specifically in soy producing establishments. If the measure of potential soy is capturing other productivity enhancing other factors besides GE soy in soy producing regions, then it would also likely predict an increase in the area of non-GE soy as well. Recall that Table 3 showed that the potential soy measure does have predictive power over the actual adoption of genetically engineered soy. Appendix Table 19 expands this result by using the percent of harvested land that was harvested with non-GE soy as the dependent variable. Reassuringly, the potential soy yield has a negatively- signed but statistically insignificant effect on the share of non-GE soy.[107] Moreover, there is also no effect of potential maize on the expansion of GE soy.

Appendix Table 18 reports the estimates of equation (9) on employment outcomes from 1980 to 1991 to examine pre-trends in labor market outcomes. Reassuringly, the measure of potential soy has no statistically or economically significant effect on any of these outcomes.

107. Also using the total land used in agriculture as the dependent variable, there is a negatively signed but statistically insignificant coefficient - suggesting that the expansion on GE soy is not coming from an expansion into previously unfarmed land.

A.18: First Difference Estimates of Potential Soy on Employment Outcomes: 1980 and 1991 Census

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Female Economically Active	Δ Female Agricultural Occupation	ΔFemale Unpaid Agriculture	Δ Male Economically Active	Δ Male Agricultural Occupation	Δ Male Unpaid Agriculture
ΔPot. Soy	-0.00309	0.000651	-0.00134	-0.000860	0.00175	0.00197
	(0.00439)	(0.00749)	(0.0112)	(0.00274)	(0.00468)	(0.00421)
Observations	3,623	3,623	3,623	3,623	3,623	3,623
R-squared	0.383	0.287	0.048	0.342	0.533	0.076
Controls	ALL	ALL	ALL	ALL	ALL	ALL
State FE	YES	YES	YES	YES	YES	YES
$\bar{Y}_{2000}$	0.42	0.10	0.25	0.81	0.41	0.11

Robust standard errors clustered at the Microregion level are in parentheses

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

A.19: Agricultural Census GE and Non-GE Soy

Panel A: 2006

	(1)	(2)
VARIABLES	Share GE Soy (%)	Share Non-GE Soy (%)
ΔPot. Soy	1.162***	0.398
	(0.324)	(0.527)
ΔPot. Maize	0.128	-0.499**
	(0.163)	(0.213)
Observations	4,181	4,181
R-squared	0.354	0.153

Panel B: 2017

	(1)	(2)
VARIABLES	Share GE Soy (%)	Share Non-GE Soy (%)
ΔPot. Soy	3.713***	-0.656
	(0.837)	(0.648)
ΔPot. Maize	-0.529	-0.370
	(0.383)	(0.301)
Observations	4,192	4,192
R-squared	0.375	0.101

Robust standard errors clustered at the Microregion level are in parentheses

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

Inference and Unit of Analysis

In this section, I examine the robustness of the main results in regards to inference and the unit of analysis. All standard errors presented have allowed for the arbitrary correlation of errors across municipalities and over time within microregions. I now 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.[108] Online Appendix Table 21 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, I relax the assumption that the municipality approximates the local markets by aggregating the main variables to a larger level of observation, the Mesoregion. Appendix Table 23 reports the estimates at the level of the Mesoregion. Overall, the results are consistent with the main results: A reduction in the employee share of agricultural employment for men and women, and while the male earnings results are not statistically significant, the female earnings and overall fertility results are larger in magnitude than the municipality level estimates.

A.20: Additional Controls

	(1)	(2)	(3)	(4)
VARIABLES	Δ Log(Female Agr Earnings)	Δ Log(Male Agr Earnings)	Δ Log(Agr Family Earnings)	Δ Fertility
ΔPot. Soy	-0.105**	0.0268	0.0295*	2.739***
	(0.0503)	(0.0187)	(0.0152)	(1.051)
Observations	4,238	4,238	4,238	4,238
R-squared	0.103	0.089	0.076	0.153
Controls	YES	YES	YES	YES

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 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 from Column (1)-(4) are the log of female agricultural earnings,the log of male agricultural earnings, the log of overall family earnings for households in agriculture (taken from the household survey of the Demographic Census), and the birth rate per 1000 women aged 16-49. 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.

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

Inference, Weighting, and Aggregation

A.21: Inference: Alternative Standard Errors

	(1)	(2)	(3)	(4)
	Δ Female Agr Earnings	Δ Male Agr Earnings	Δ Family Earnings	Δ Fertility
ΔPot. Soy	-0.110**	0.0295	0.0274*	2.661**
	(0.0509)	(0.0198)	(0.0158)	(1.050)
Mesoregion	0.066	0.024	0.019	1.194
Conley SE, 50km	0.048	0.016	0.015	0.785
Conley SE, 200km	0.064	0.023	0.017	1.153
Observations	4,254	4,254	4,254	4,254
R-squared	0.098	0.079	0.072	0.147

Standard errors in parentheses
*** 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, showing sensitivity to different assumptions on standard errors. The dependent variables from Columns (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 16-49. 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.

A.22: Main Results: Weighted by Population

Panel A: Weighted

	(1) Δ Paid Agriculture^f	(2) Δ Log Agr. Female Earnings	(3) ΔLog Agr. Male Earnings	(4) ΔLog Agr. Relative Earnings	(5) ΔLog Family Earnings	(6) Δ Fertility
Pot. Soy	-0.00687*	-0.106***	0.0324*	-0.0953***	0.0191**	2.505***
	(0.00402)	(0.0391)	(0.0177)	(0.0277)	(0.00813)	(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

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 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 16-49. 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.

A.4: Robustness

Panel A Weighted by Female Population (16-49)

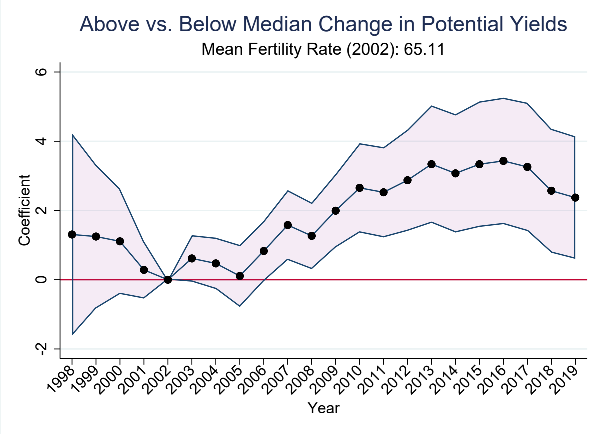

Panel B. Ln(Births)

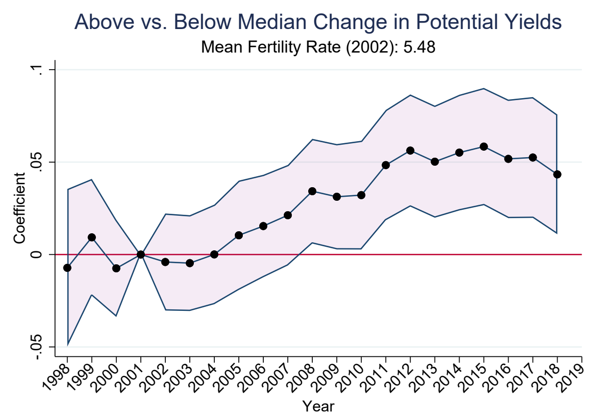

These sub-figures plots estimated coefficients from (11) where the dependent variable in the first figure is the birth rate, defined as the number of live births per 1000 women aged 16-49, and the second is the natural log of the number of births. The data run from 1997-2019 and have a balanced sample 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.

A.23: Main Results at the Mesoregion Level

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Δ Female Earnings	Δ Male Earnings	Δ Female Employee Earnings	Δ Female Self-Employed Earnings	Δ Male Employee Earnings	Δ Male Self-Employed Earnings
ΔPot. Soy	-0.259**	0.0220	0.00223	-0.114	0.0584	-0.114
	(0.109)	(0.0581)	(0.0704)	(0.141)	(0.0438)	(0.141)
Observations	136	136	136	136	136	136
R-squared	0.363	0.307	0.192	0.198	0.457	0.198

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Δ Female Employees	Δ Female Self-Employed	Δ Female Unpaid	Δ Male Employees	Δ Male Self-Employed	Δ Male Unpaid
ΔPot. Soy	-0.0769***	-0.00275	0.0770***	-0.0206*	-0.0149	0.0356***
	(0.0208)	(0.0174)	(0.0278)	(0.0115)	(0.0114)	(0.0130)
Observations	136	136	136	136	136	136
R-squared	0.485	0.485	0.539	0.551	0.477	0.442

VARIABLES	Δ Fertility
ΔPot. Soy	3.441*
	(1.922)
Observations	136
R-squared	0.530