Fathers and Children: Labor Market Opportunities and Fertility

Alexander Russov*

November 12, 2016

Abstract

This paper studies the causal impact of short-run shocks to men's and women's labor market opportunities on fertility. I evaluate how jobs in the formal sector, in manufacturing, and at export-assembly plants (maquiladoras) in Mexico shape childbirth, selection into fertility, and the timing of births using two complementary identification strategies. The first strategy exploits exogenous shocks to demand for male versus female labor using a region's industrial structure, and the second uses establishment-level data from the universe of maquiladoras to construct an instrumental variable based on large expansions and contractions in plant employment. Results show that positive shocks in the short run to men's employment have large, positive effects on fertility, whereas positive shocks to women's employment have small net impacts in the short run.

*Department of Economics, University of Michigan, email: arussov@umich.edu. I am grateful to Achyuta Adhvaryu, Manuela Angelucci, Hoyt Bleakley, John Bound, Charlie Brown, David Lam, and participants at the University of Michigan Summer Seminar, Economic Development Seminar, and PopPov 2016 conference for their comments. I acknowledge support from the Rackham Graduate School at the University of Michigan and the Michigan Institute for Teaching and Research in Economics. I also wish to thank Natalia Volkow for assistance with using the restricted-access EIME data at INEGI's microdata lab in Mexico City. The views expressed in this paper are those of the author and do not necessarily represent those of the Mexican National Institute of Statistics and Geography.

1 Introduction

This paper investigates how short-run changes in labor markets in Mexico for men and women affect decisions about fertility. In the last few decades, many developing countries have industrialized, and large numbers of women have entered the labor force for the first time. Much of the rise in women's employment is due to the growth of export-oriented manufacturing employment, especially in low-skill jobs. This study contributes to a recent, growing literature evaluating the consequences of this remarkable transition in women's labor force attachment on families' decisions about fertility and marriage. Unlike other work that focuses solely on positive shocks to women's labor market opportunities (Heath and Mubarak, 2015; Sivasankaran, 2014; Jensen, 2012), I study both expansions and contractions in employment for men and women.

I exploit multiple sources of data over different time periods to build two complementary instrumental variables identification strategies to study how labor market opportunities in the formal sector for both men and women affect fertility, selection into fertility, and the timing of births. The first approach generates a predicted measure of employment, building on an approach commonly used in labor and urban economics to isolate exogenous shocks to labor demand, to identify how these local labor demand shocks shift fertility. The second approach uses a restricted-access, establishment-level dataset on the universe of maquiladoras (export-assembly plants) to construct an instrument based on expansions and contractions at the factory-level.

Focusing on short-run shocks, I find that increases in labor market opportunities for men have large, positive impacts on fertility, whereas increases in women's labor market opportunities show small, negative impacts on fertility. I interpret these results in the context of a theoretical framework consistent with neoclassical theories of fertility dating back to Becker (1960): families choose whether to have a child in each period, and changes in employment for men and women generate income and substitution effects that alter the proportion of households choosing to have a child. The results indicate that income and substitution effects roughly offset each other in the short run for women.

This paper also presents new findings on how employment dynamics impact fertility. Distributed lag models provide evidence that the immediate effect of demand shocks for women's employment is negative, indicating that substitution effects dominate income effects in the near term. The net effect, however, indicates that short-run demand shocks for women are not associated with large changes in fertility. When I focus on different frequencies in employment variation, I find that high-frequency demand shocks to women's employment reduce fertility, whereas low-frequency movements are not associated with significant impacts on fertility. On the other hand, both high- and low-frequency movements in employment for men lead to increases in fertility. These findings contribute to the study of labor markets and fertility by showing that employment dynamics for both men and women meaningfully alter fertility decisions.

The relationship between female labor force participation and fertility remains a matter of debate (Engelhardt and Prskawetz, 2004; Kögel, 2004; Mishra, Nielsen, and Smyth, 2010). Lim (2009) notes that the cross-sectional relationship within countries has become much less steep across time, and in some regions there is no relationship at all.[1] If women's opportunity cost of time, measured by wages or hours worked, is a central mechanism driving fertility decisions, as in neoclassical models following Becker (1960), then why is the relationship not robust across time and space?[2]

That theory does not provide a firm prediction about how changes in labor markets should affect fertility may explain some of the differing results across studies. The empirical analysis in this paper is informed by two points. First, labor demand shocks may lead to heterogeneous effects that differ across households, depending on initial labor force participation, labor market frictions (Da Rocha and Fuster, 2006), availability of childcare arrangements (Del Boca, 2002), other policies that alter the incentives to have children, or the quality of the work itself.[3]

[1] She shows that there is no clear pattern in Asia-Pacific countries, while both female labor force participation and fertility fell in many countries in the Middle East and North Africa in the 1990s.

[2] A neoclassical labor supply model predicts income and substitution effects resulting from increased wages. There is no theoretical reason to assume that substitution effects will dominate income effects for women, though existing research focusing on both men's and women's labor market opportunities (e.g. Schultz, 1985; Schaller, 2016) supports the view that substitution effects dominate. Typically, it is thought that at low wages, increases in wages lead to more hours worked as the substitution effect of higher wages dominates the income effect. This results in less time for leisure and raises the cost of having a child. At higher wages, increases in the wage may lead to reduced hours of work if the income effect dominates the substitution effect. Of course, as noted by Ahn and Mira (2002), in reality most individuals do not pick their optimal hours of work. Imagine, for simplicity, that work options are binary: an individual may either accept a full-time job or not. If she is already working, then an exogenous shock to her wage, say through a positive labor demand shock, induces only an income effect. That implies that an increase in a woman's wage would lead to increased fertility, assuming children are normal goods. Lindo (2010) and Black et al. (2013) provide evidence that children are indeed normal goods.

[3] Employment is not merely associated with changes in wages and the cost of time. Lim (2009) argues that

Second, the interaction between labor markets for men and women is critical. The bulk of the research on fertility focuses exclusively on either male wages or female wages, aggregate unemployment, net job changes for women only, or devises a single ad hoc measure, such as the ratio of male wages to female wages or the female share of exports (e.g. Do, Levchenko, and Raddatz, 2016), even though the motivation for such work relies on how wage changes separately affect income and the cost of time for men and women. Approaches that cannot distinguish male labor demand shocks from female labor demand shocks potentially suffer serious omitted variables biases when the fertility decision depends jointly on male and female income and time.[4] In particular, the relationship between fertility and the woman's wage may depend on whether one conditions on her partner's wage or not, and the same applies to the man.[5]

Furthermore, existing work is not always clear regarding what parameter is being identified, either from a theoretical or policy-based perspective. For instance, increases in wages may induce increases or decreases in hours worked among those already working (especially if the sample comes from a developed country), leading to an ambiguous connection between wages and female time use. Unemployment rates suffer from well-known issues relating to indeterminate changes in the numerator and denominator (since both employment and labor force participation may change across business cycles). Wages or earnings, especially in a setting like Mexico where labor

"increases in labor force participation have not been matched by improvements in job quality and that the kinds of jobs women are engaged in and their working conditions have not led to their true socio-economic empowerment, have not provided adequately satisfying alternatives to childbearing or have not involved serious incompatibility between paid and unpaid work." She suggests a number of mechanisms linked to women's employment that should lower fertility, such as whether the quality of the work enhances women's status, thereby increasing their independence and bargaining power within the household. Other research (Ñopo, 2012; World Bank, 2011) documents that many women's jobs in Latin America are low-wage and highly segregated from men's jobs, implying that employment opportunities may not sufficiently lift women's status to change fertility decisions. The findings in this paper show that while labor demand shocks to men's labor increase wages, labor demand shocks to women's labor actually lead wages to decline. Part of this result, however, may be due to changes in the composition of the labor force.

[4]Had I used a measure of the ratio of men's to women's employment, I would have incorrectly found that short-run growth in women's employment decreases fertility, when the effect is driven by short-run growth in men's employment increasing fertility.

[5]Summarizing a major strand of the literature, Jones et al. (2010) note that the correlation between fertility and the wife's wage is usually negative, whether one conditions on the husband's wage or not; the unconditional correlation between the husband's wage and fertility is negative; and lastly, the correlation between the husband's wage and fertility, conditioning on the wife's wage, is either positive or negative, depending on the study.

force participation among women remains low, can suffer from severe composition bias (Solon, Barsky, and Parker, 1994). To address these issues, my analysis focuses on changes in net formal sector employment for men and women. Thus, I identify the net impact of individuals being induced into formal sector employment on fertility.

This paper also contributes to a large literature (e.g. Butz and Ward, 1979; MacDonald, 1983, Macunovich, 1995; Currie and Schwandt, 2014) relating fertility rates to the business cycle, and whether fertility is pro- or countercyclical remains an open question in the literature. Ahn and Mira (2002) note that in their sample of OECD countries, the relationship abruptly switched from being negative to positive. They link the change to a new equilibrium of high unemployment and higher female labor force participation in the OECD. Most of this work focuses on wealthy nations, and one of the contributions of this study is to fill the gap by focusing on a developing country. Moreover, much of this literature focuses on time series regressions that take the country as a unit of observation, but these studies do not identify the impacts of local labor market demand shocks on fertility; for instance, nation-wide changes in policy on employer-provided childcare can induce changes in work and fertility that are unrelated to demand shocks.

Finally, because the time and resource costs of children extend across time[6], it is important to consider how a dynamic theory alters the predictions of the static model. Consider a woman who loses her job or whose hours of work decrease. In the near term, the substitution effect may dominate the income effect and lead her to increase time in child care, potentially increasing fertility. However, if the job loss or wage reduction is due to a widespread negative economic shock, such as a recession, it may lead her to change her expectations about her future employment prospects. If she now anticipates permanently lower income in the future, then the job loss could lead the income effect to dominate the substitution effect and hence lead to lower fertility. A more nuanced alternative is that she may increase fertility in the near term, when the opportunity cost of her time is lower, but decrease fertility in the long-run.[7] The findings in this paper provide support for the latter view. In sum, changes in economic conditions can also

[6]Becker (1960) originally compared children to durable goods. Álvarez-Parra, Brandao-Marques, and Toledo (2013) show that spending on durable goods is particularly volatile in developing countries, including Mexico, which is consistent with my finding that fertility strongly responds to the business cycle.

[7]The presence of liquidity constraints or uncertainty about future outcomes, however, may prevent households from fully optimizing across time, implying that even shocks of a short duration matter for households.

alter expectations about future economic possibilities, which may also affect how families space births across time.[8]

2 Theoretical framework

In this section I describe a simple conceptual framework to facilitate the interpretation of the empirical results. The first neoclassical static models of fertility (Becker, 1960; Mincer, 1963) incorporated fertility choice into traditional models of consumer demand, and subsequent work (Becker, 1965; Willis, 1973; Becker and Lewis, 1973; Gronau, 1977) formalized and extended the framework to study the quantity-quality tradeoff, incorporate household production, and establish tradeoffs between human capital investments, working, and having children. Later, the theory was extended into a dynamic setting, focusing on the rich choice set influencing the timing of births.[9]

2.1 General utility maximization problem

A general formulation of the household's utility maximization problem takes the following form:

$$E\left[\sum_{t=a}^{T}\delta^a U_{i,t}(c_{i,t}, n_{i,t}, \epsilon_{i,t})\right] \tag{1}$$

If there is no saving or borrowing, the money and time constraints are as follows:

$$c_{i,t} + \gamma_{i,t} n_{i,t} \leq Y_{i,t} + w_{i,t} L_{i,t} \tag{2}$$

$$\theta_{i,t} n_{i,t} + L_{i,t} \leq 1 \tag{3}$$

[8] An additional line of thinking, dating back at least to Leibenstein (1957), notes that intergenerational transfers play a major role in encouraging fertility in developing countries. Although I do not pursue the study of the intergenerational transmission of fertility here, it is worth keeping in mind that in countries like Mexico, where some individuals belong to a formal sector that provides social security, and others belong to an informal sector without similar protections, the implications of increased jobs and changes in wages may differ in regions with few formal sector jobs.

[9] Hotz, Klerman, and Willis (1997) provide an overview of many of these models, and Schultz (1997) describes features particular to low-income countries.

In equation (1), each household i maximizes the expected remaining lifetime utility of consumption, denoted by $c_{i,t}$, children, denoted by $n_{i,t}$, and other factors, denoted by $\epsilon_{i,t}$, across time periods denoted by t. The budget constraint in equation (2) requires consumption and the goods-cost of children (where $\gamma_{i,t}$ is the average goods-cost per child) to equal the household's earnings, which are equal to the husband's income, denoted by $Y_{i,t}$, and the wife's earnings, which are equal to her wage, denoted by $w_{i,t}$, multiplied by the proportion of time she spends working, denoted by $L_{i,t}$. Consistent with most models of fertility, I assume only the mother plays a role in taking care of the child. Thus, the father's income enters the problem purely as earnings. The time budget constraint in equation (3) requires the time the mother devotes to caring for the child, where the time-cost of the child is denoted by $\theta_{i,t}$, and working not to exceed the total amount of time available, which is standardized to equal 1.[10]

The time and goods costs of children are likely to differ across couples, such that $\gamma_{i,t}$ and $\theta_{i,t}$ are better conceived of as functions of all variables that could possibly affect how time and money-intensive raising children is for a particular family; in this case, we can write $\gamma_{i,t} = f_{i,t}(Y_{i,t}, w_{i,t}, \epsilon_{\gamma,i,t})$ and $\theta_{i,t} = g_{i,t}(Y_{i,t}, w_{i,t}, \epsilon_{\theta,i,t})$, where I assume that the time and goods costs depend on the earnings of the father, the wage of the mother, and all other factors, denoted by $\epsilon_{\gamma,i,t}$ and $\epsilon_{\theta,i,t}$. For instance, if the mother prefers to stay at home to take care of children, $\theta_{i,t}$ may be high, whereas if she can find relatives to take care of children, $\theta_{i,t}$ may be low. If the mother prefers to hire a nanny or pay for childcare rather than spending time with the child, $\gamma_{i,t}$ and $\theta_{i,t}$ will be inversely related, as she can find substitutes for her time. The availability of these substitutes, in turn, depends on the income of the household as well as other factors that influence expenditures on children and time costs of children.

2.2 A model of fertility

Equations (1) through (3) describe a general framework for how households optimize consumption and fertility across time, including both tempo (timing of births) and quantum (total number of births) effects. To illustrate how changes in income for men and changes in labor force participation and wages for women affect the decision to have a child and motivate the reduced-form analysis in this paper, I adopt a special form of the above framework, and assume

[10]The constraint is binding as long as the household has a child or the woman works.

that a unitary household[11] chooses consumption, whether to have a child, and whether the mother works to maximize the following in a given time period:

$$u(c) + \beta n - d(LFP) \tag{4}$$

I drop the household and time subscripts for notational convenience, but it should be understood that the form of this function can vary across households and time. The function u is assumed to be strictly concave, and d represents the disutility to the household from the woman's work. I assume disutility arises from labor force participation and not wages or hours of work, so that d can be written as the product of some constant and an indicator function, i.e. $d(LFP) = d_{i,t}1\{w > 0\}$. [12] Substituting (3) into (2), the only constraint[13] is

$$c + \gamma n \leq Y + (1 - \theta n)w \tag{5}$$

To focus on how short-run changes in labor market opportunities for men and women affect current fertility decisions, I assume that households choose whether to have one or no children in each period, and $\theta \in [0, 1]$. The simple nature of this binary problem allows us to focus on the connection between employment, wages, and fertility across time.[14] In particular, the analysis below addresses the following question: how do changes in men's and women's earnings via demand shocks from one period to the next affect the optimal choice n^*?

[11] Assuming a unitary household is done for mathematical convenience. One can assume a particular bargaining rule, or use the more general collective framework of Chiappori (1992) and Blundell, Chiappori, and Meghir (2005). If the child is thought of as a public good, then as long as parents have the same preferences the weights reflecting bargaining power only determine the relative allocation of consumption, not children. However, if they have different preferences over consumption and children, then the distribution of bargaining power will affect optimal choices of both. Bargaining over fertility is discussed in Ashraf, Field, and Lee (2014), Kemnitz and Thum (2015), Rasul (2008), and Doepke and Kindermann (2016).

[12] The choice of the disutility function is in the spirit of Basu (2006), who models multiple equilibria in the context of female labor force participation.

[13] The time constraint holds with equality only if the optimal choice involves having a child, working, or both. The woman does not work and the household chooses not to have the child if the disutility of working is high enough.

[14] The set-up assumes that households face liquidity constraints or uncertainty about the future, which would lead households to potentially alter fertility decisions in response to short-run labor market shocks.

Case 1: An Increase in Y

For any $w \geq 0$ and $Y \geq 0$, the household chooses to have a child in a given period if the utility from having a child exceeds the utility from not having a child. If $n^* = 1$, consumption is given by $c^* = Y + (1-\theta)w - \gamma$, and if $n^* = 0$, consumption is given by $c^* = Y + w$. Thus, the household chooses not to have a child if

$$u(Y + (1-\theta)w - \gamma) + \beta < u(Y + w)$$

$$\implies \beta < u(Y + w) - u[Y + (1-\theta)w - \gamma] \tag{6}$$

Recall that values of β differ across households and time, allowing us to think of these values as being drawn from some probability density function p: $\beta_{i,t} \sim p(\beta)$. Given the functional form of u and the available labor market opportunities, households that draw a small value of β from p will choose not to have a child. For a household on the margin between having and not having a child the following holds for some $\beta = \bar{\beta}$:

$$\bar{\beta} = u(Y + w) - u[Y + (1-\theta)w - \gamma] \tag{7}$$

Suppose the man's earnings increase in a given period. Strict concavity of u implies that now a smaller value of $\bar{\beta}$ is required to make (7) hold. Thus, a smaller proportion of households now chooses not to have a child—all that satisfy $\beta_{i,t} < \bar{\beta}$.[15] To summarize, β represents the compensation (in utils) the household requires to choose to have a child in any period, and positive (negative) changes in men's earnings due to labor demand shocks require a lower (higher) amount of compensation. Since a higher proportion of households chooses to have a child with increases in Y, it follows that a child is a normal good in this framework.

Case 2: An increase in w

Because this paper focuses on changes in employment opportunities, I consider two cases: $w = 0$ (the woman does not work) and $w > 0$. Some types of labor demand shocks may induce women to enter the labor force, increasing the proportion of women with positive wages, while

[15] More formally, $\frac{\partial \bar{\beta}}{\partial Y} < 0$ implies $\frac{\partial P(\bar{\beta})}{\partial Y} < 0$ since P is the probability distribution function associated with p.

others may mainly affect the wages of women already working.[16] Continuing with the binary setup introduced earlier, I consider the choice between having and not having a child, separately for women who work and those who do not. The table below lists the utility of the household in each case:

	$w = 0$	$w > 0$
$n = 0$	$u(Y)$	$u(Y + w) - d(LFP)$
$n = 1$	$u(Y - \gamma) + \beta$	$u(Y + (1 - \theta)w - \gamma) + \beta - d(LFP)$

Consider a demand shock to women already working (that is, the righthand column) that raises their wages and a household on the margin between having and not having a child. By revealed preference, the household never switches into either cell in the lefthand column, and since $d(LFP)$ is the same for all positive wages, the problem is the same as in (7), but now strict concavity of u is not sufficient to establish whether $\bar{\beta}$ must rise or fall for (7) to hold. To see this, consider two separate cases. If $\theta = 0$, then $\bar{\beta}$ must decline for (7) to hold. If $\theta = 1$, then $\bar{\beta}$ must rise for (7) to hold. For intermediate values of θ, the relationship is indeterminate. Intuitively, heterogeneity in θ—the variation across households in terms of how time-intensive raising a child is—determines whether income or substution effects dominate for increases in the woman's wage.

Now consider a demand shock to women who are not working (the lefthand column). If the wage change is not sufficient to counteract the disutility of labor force entry, then the optimal choice of n^* is not affected. If the wage change does lead to entry into the labor force, then the same logic as above holds: some households may choose to have a child, while others may choose not to have a child. Heterogeneity in values of θ, γ, $d(LFP)$, and β gives rise to predictable patterns: households that draw a higher value of β, a taste parameter for children, will be more likely to have a child; women in households that draw a higher value of $d(LFP)$ will be less likely to enter the labor force; and households that draw higher values of θ and γ, which increase the time and monetary costs of having a child, will be less likely to have a child. Whether shocks to the woman's wage lead to aggregate changes in fertility, then, is an empirical matter.

[16]If wages for low-skilled women increase, then average wages for women may decrease, even though wages have weakly increased for all women. Although the model of fertility relies on changes in wages for women, potential changes in the composition of the labor force imply that aggregate changes in employment are a better proxy than wages for changes in labor market opportunities.

3 Data, Background, and Trends in Employment and Fertility

3.1 Data description

The fertility data in this paper come from the Mexican National Institute of Statistics and Geography (INEGI). These vital statistics data contain individual-level information on all births, including parents' ages, education, type of union, and the municipality they live in. I use the month of birth variable and lag it by nine months to proxy for the year of conception.

Demographic data on municipalities come from the 2005 and 2010 Mexican censuses, and demographic information for intercensal years is linearly interpolated. Employment data come from two sources. The main analysis uses data from the Mexican Social Security Institute (IMSS), which administers the provision of health care, pensions, and social security. All employees in the formal private sector are obligated to register with the IMSS, so these administrative, job-level data contain employment information on the universe of jobs in the formal private sector. I also use data from the Survey of Occupation and Employment (ENOE), which tracks Mexico's labor force and provides detailed information on the characteristics of employment. Unlike the IMSS data, it has the advantage of providing information on employment in all sectors of the economy, but as a survey it contains only a small proportion of formal sector employment and lacks information on some municipalities entirely. The main analysis focuses on years 2005 to 2013, which is the set of years for which IMSS and fertility data are available.

Data on maquiladora line-employment come from the Maquiladora Export Industry Dataset. INEGI collected these data at the monthly level from all export-assembly plants in Mexico from 1990 to 2006. These data contain plant identifiers[17] as well as a variety of information on inputs, expenditures, sales, and value-added. The period of study encompasses major changes in Mexico's exposure to trade, including the signing of NAFTA, the peso crisis, and China's entry into the WTO.

3.2 Background on fertility and employment

Figure 1 shows the tremendous variation in general fertility rates[18] across municipalities across Mexico in 2010. These, unsurprisingly, are highly correlated with local incomes, stocks of human capital, urbanization rates, and proportion of speakers of indigenous languages. Mexico's total

[17] Data with plant identifiers must be accessed on-site at INEGI's microdata lab in Mexico City.

[18] I define the general fertility rate as the number of births per 1000 women of ages 15-44.

fertility rate peaked at above seven children per woman in the early 1960's and then entered a period of sharp, continuous decline; it is currently estimated at being just over two children per woman. General fertility rates are also declining: Figure 2 shows the overall trajectory in general fertility rates for the country since 2006, and Table 1 lists fertility rates for federal entities in 2005 and 2013. The most recent trend continues to be downward, but the overall trend in Figure 2 masks substantial variation across the country. For instance, splitting the sample of municipalities into quartiles by levels of urbanization indicates that mostly rural municipalities actually experienced a bump in fertility in 2009, the year of the U.S. financial crisis spilling over into Mexico.

While fertility was steadily trending downward prior to 2009, the formal sector employment to population ratio was increasing for both men and women, as indicated in Figure 3. Formal sector employment then contracted sharply in 2009 and has been slowly recovering since. Manufacturing employment in the formal sector has experienced the same trends, but witnessed a much sharper decline in 2009. As has been the case in the U.S. in recent recessions, the men's employment to population ratio suffered a substantially larger decrease in 2009, though within manufacturing women experienced a slightly sharper decline. The overall trends for men and women in employment, though, are similar, which makes connecting aggregate changes in fertility to changes in either men's or women's labor market opportunities (without conditioning on the other gender's job prospects) problematic.

3.3 Employment and fertility correlations

Consider, as a benchmark, the case where fertility is a function of overall job opportunities. I first estimate the following regression:

$$y_{m,t} = \lambda + \beta E_{m,t} + \boldsymbol{\alpha} f(X)_{m,t} + \gamma_m + \delta_t + \theta Trend_s + \varepsilon_{m,t} \tag{8}$$

The outcome is the natural logarithm of the fertility rate[19], measured as number of births per 1000 women aged 15-44.[20] $E_{m,t}$ measures the natural logarithm of total formal sector employment in municipality m in year t for individuals aged 15-44. I also include municipality

[19]Results are similar if using the raw fertility rate. Using natural logarithms on both the lefthand and righthand sides facilitates interprtation of the coefficients of interest as elasticities.

[20]This is equivalent to simply using the log of the number of births when the log of population of women aged 15-44 is used as a control.

fixed effects, which control for time-invariant unobservable characteristics specific to each locality; year fixed effects, which control for annual shocks to fertility; and linear state time trends, which control for smoothly evolving determinants of fertility that vary across states (e.g. if regions with high fertility are converging to the fertility rates of regions with low fertility). [21] Standard errors are clustered at the municipality level to account for serial correlation within municipalities (Bertrand et al., 2004).

I do not calculate an unemployment rate as is commonly done in the literature (e.g. Dehijia and Lleras-Muney, 2004; Örsal and Goldstein, 2010; Schaller, 2016) since I only have administrative data on the formal sector, which accounts for fewer than half of all jobs in Mexico. To account for the potentially complex relationship between formal and informal sector jobs, I flexibly control for the natural logarithm of the population of men and women aged 15-44 in the function $f(X)$. These regressions are weighted by the population of women aged 15-44 in each municipality, averaged across years. I limit the sample to those municipalities with employment in manufacturing to make results comparable between samples using all employment or only employment in manufacturing.

Column 1 in Table 3 show that fertility is procyclical, and whether changes in population are accounted for linearly or in a more flexible way matters little. Following the framework established earlier, I examine whether manufacturing jobs have a similar effect on employment and re-estimate equation (1) using only formal sector jobs in manufacturing. The results are qualitatively similar using both measures of employment, although the magnitudes are much smaller when focusing on the subset of the formal sector in manufacturing, which is in line with manufacturing employment making up a smaller fraction of the employment stock in municipalities.[22]

Labor markets in Mexico are segmented by sex.[23] For instance, production in garments, toys,

[21] Results are unchanged when using quadratic time trends, and results using state-by-year fixed effects, which are shown in the robustness section, are similar.

[22] Strictly speaking, this interpretation depends on the omitted part of employment being orthogonal to manufacturing, which is unlikely to hold due to spillovers among sectors, but is a rough approximation of effects. In a later section, I show instrumental variables results as well as results using only maquiladoras, which relies on expansions and contractions that are arguably exogenous to shifts in other sectors.

[23] Labor market segmentation by sex is common across all countries, even those with high degrees of male-female equality in other spheres, but developing countries in Latin America are especially likely to show segmentation (World Bank, 2011).

musical instruments, perfumes, and cosmetics are the most female-centric sectors within manufacturing, while alcohol, automotive, and concrete manufacturing are the most male-intensive. As a result, opportunities for men and women differ widely across municipalities and time, depending on the share of industries in each location and the growth rate in employment across time. To test whether men's and women's employment have differential impacts on fertility, I estimate the following specification:

$$y_{m,t} = \beta_{male} E_{male,m,t} + \beta_{female} E_{female,m,t} + \boldsymbol{\alpha} f(X)_{m,t} + \gamma_m + \delta_t + \theta Trend_s + \varepsilon_{m,t} \quad (9)$$

Results are in Table 3. The estimates for female employment are negative but small and statistically insignificant, while the estimates for men are larger and statistically significant at the 1% level. Note that the magnitudes can be interpreted as elasticities. Controlling for female employment and changes in population, the results imply that a 10% increase in formal sector employment for men is associated with a 0.3% increase in fertility rates in that municipality.

These regressions omit the informal sector, which may have differential impacts on fertility. Employment in the informal sector tends to be countercyclical, and negative shocks to the formal sector historically have not led to significantly higher unemployment rates in Mexico. Instead, the informal sector has served as a "safety valve" for individuals on the margin of losing formal sector employment.[24] Since this implies that changes in employment in the informal sector are themselves outcomes of changes in employment in the formal sector, estimates that include both informal and formal sector employment should be interpreted with caution.

Under the theory outlined here, employment in the informal sector should have a smaller impact than the formal sector on fertility. Even if informal sector jobs do not pay less than the formal sector[25], they do not typically provide childcare, social security, and health benefits, so the expected income effect for men is smaller. Women are less likely to have strong attachments to the labor force in the informal sector, so the expected negative effect (if substitution effects dominate income effects) on fertility from growth in the informal sector is likely to be smaller as well. Table 3 shows results combining the formal sector data in the IMSS and data on the

[24] Early models of labor market segmentation construed the informal market as a separate market that serves as an alternative to individuals who cannot enter the formal sector, but more recent work has emphasized that many workers move voluntarily between the formal and informal sectors (Maloney, 2004). Time series of employment indicate that, perhaps due to the stickiness of wages in the formal sector, employment in the informal sector expands during economic crises (Binelli and Attanasio, 2010).

[25] Marcouiller et al. (1997) discuss the wage gap between the formal and informal sectors in Mexico.

informal sector from the ENOE labor force survey, disaggregated by sex. They are consistent with the theory: estimated coefficients for both male and female employment in the informal sector are close to zero. Furthermore, the inclusion of the informal sector does not alter the main results using employment only the formal sector.[26]

4 Identifying the causal impact of local labor market shocks

4.1 A measure of predicted employment

The analysis so far has shown that fertility is positively correlated with employment in the formal sector, and this relationship is driven entirely by male employment. Such a relationship may not be causal, however. Using local aggregate employment, as opposed to own-employment, alleviates a major concern about reverse causality: a parent may change his or her own labor force participation in anticipation of or as a consequence of having a child.

Nevertheless, using aggregate measures of employment may still not identify the causal impact of local labor demand shocks on fertility. Potential sources of bias for the fixed effects estimator include simultaneity, omitted variables, and measurement error. Simultaneity may arise if women decrease childbearing, leading some to enter the labor force differentially across municipalities and time in a manner not fully accounted for by the fixed effects employed here. Such a relationship would lead to a downward bias in estimates on women's labor demand.

Omitted factors correlated with labor demand and fertility may also bias estimates. For instance, La Ferrara, Chong, and Duryea (2012) find that telenovelas decrease fertility in Brazil. This points to the potential importance of social norms in shaping households' decisions about women working outside the home and fertility preference. Furthermore, linearly interpolating the population between census years may smooth temporary shocks to population. Such shocks are likely to be positively correlated with employment and the number of births, leading to an upward bias in the ordinary least squares estimator for both male and female employment.[27]

Finally, measurement error in the variable measuring formal sector employment can induce biases in estimates of both men's and women's employment. The administrative data cover the

[26]Note the labor force survey does not cover all municipalities in Mexico and only goes back to 2006 in its current version, so the sample of included municipalities and years is smaller.

[27]Replacing the census population figure with the population estimate from an average of the four quarters in the subsample of municipalities and years in the ENOE does not alter the results here.

universe of formal sector employment, so in that sense they are measured without error. However, some of the movements in and out of the formal sector may represent firms' selectively registering with the formal sector in some periods and not in others. This would lead to attenuation bias in both coefficients.

To deal with these sources of bias and isolate the impact of changes in local labor demand, I employ two identification strategies in this paper. The first strategy builds on the approach originally employed by Bartik (1991) and used in Blanchard and Katz (1992), Bound and Holzer (2000), and others in urban and labor economics in constructing an instrumental variable predicting labor demand using a "shift-share index". I define

$$Log\ (predicted\ employment)_{m,g,t} = log \sum_{ind} \frac{Emp_{m,g,ind,t=0}}{Emp_{g,ind,t=0}} (Emp_{g,ind,t} - Emp_{m,g,ind,t})$$

to predict the log of employment for group g (men or women) in each municipality m at each time period t. The numerator in the fraction is equal to employment of group g in municipality m in industry ind at time 0, that is, the baseline period. The denominator in the fraction is equal to national employment of group g in industry ind in the baseline period. If one ignores the last of the two terms in parentheses, then in the baseline period this expression is equal to actual employment. In subsequent years it deviates from actual employment because the mix of industries in each municipality is kept constant to address endogenous changes in the industrial mix resulting from local changes in labor supply. The instrument predicts local employment by weighting national employment in each industry with the proportion of employment in that industry located in the municipality in the first period and summing over all industries, separately for men and women. As is common in the literature, I subtract local employment from national employment in parentheses so that the predicted labor market outcome excludes actual local employment. Otherwise, part of the association between the instrument and actual employment would be mechanical.[28]

This methodology relies on a municipality's industrial mix in the baseline period predicting outcomes for local workers in subsequent periods. That is, if one municipality has a large employment share in sectors that employ women, such as textile manufacturing, and the employment of women in textile manufacturing increases across the country, we would expect local

[28] In practice, local employment in each municipality is small enough that it makes no difference for the outcomes in this paper whether it is included in the expression or not, although the power of the instrument is of course stronger when own-employment is included.

employment of women to increase. Assuming workers in one sector are comparable across the country, such that a positive (negative) shock nationwide translates into a positive (negative) shock locally, the IV should predict actual local employment.

To satisfy the exclusion restriction, the predicted employment measure must not be correlated with local labor supply shocks. This requires that national changes in employment in a given industry are not due to changes within a single municipality. Mexico consists of over 2,000 municipalities, with the largest accounting for a little over 1% of the population, so this is much less of a concern than in similar studies using each state as the local labor market in the U.S. A more subtle point, made in Cosman (2014), is that localities with a similar mix in sectors may be related in other ways, leading the instrument to pick up differences associated with a particular industrial mix rather than differences in employment in that industry.[29] He concludes that Bartik-style instruments do indeed predict local changes in employment from national shocks to industries.

In principle, the shift-share index can be created for any set of industries, but the argument behind it relies on local changes in employment being sensitive to national trends. Local industries, such as those in services, are less likely to respond to national changes in the same industry than sectors that are traded nationally or globally. A demand shock to internationally traded products, in particular, is likely due to exogenously determined factors that lead to a push in local employment in sectors making those products. Hence, I create two measures of predicted employment: one set for men and women employed in the formal sector, as well as a second set focusing only on manufacturing employment in the formal sector. Focusing on manufacturing provides other advantages. First, different types of manufacturing in Mexico, as in the U.S., are centered on particular regions, leading to spatial variation in how susceptible places are to exogenous demand shocks. Second, many industries within manufacturing have been in decline in Mexico during this period due to competition from low-cost Asian producers (especially after China's entry into the WTO), but the shocks have not been felt equally across all industries, leading to another source of spatial variation within regions containing manufacturing employment. Third, to the extent that different types of employment differentially impact fertility decisions (Lim, 2009), the analysis answers a well-defined question: what impact

[29] Using Monte Carlo trials, Cosman investigates whether such a correlation would lead to a spuriously strong first stage and finds that the coefficient on the instrument would actually be negative.

do the expansion and contraction of manufacturing jobs have on fertility? Fourth, focusing on manufacturing makes the results comparable to a growing literature in economics evaluating how the liberalization of trade regimes, leading to large growth in low-skill factory jobs for young women in developing countries, has affected these women's lives.

4.2 Results using instrumental variables

Table 4 shows the first stage results, as well as the instrumented results and the reduced-form for equation 1 using all formal sector employment. Predicted employment is highly correlated with actual employment, and the first stage appears strong. The instrumented regressions indicate that the elasticity of fertility with respect to employment is slightly below 0.2. This coefficient is much larger in magnitude than the OLS coefficient, a point to which I return below.

The theory predicts that exogenous shocks to male and female labor demand differentially affect fertility decisions. If employment for men mainly translates into income effects, then we may expect that positive shocks to men's labor demand should increase fertility. Positive shocks to women's labor demand may increase the opportunity cost of having a child, while also leading to positive income effects, so the net impact of changes in women's labor market opportunities is a priori indeterminate. To examine the theory, I return to equation (9) and instrument employment for men and women with predicted employment for each group. First stage results using full formal sector employment and only manufacturing employment are presented in Table 5. When focusing on all formal sector employment, the IV for women is strongly correlated with actual female employment, while both IVs have some predictive power for male employment, although the IV for male employment is greater in magnitude and slightly more precisely estimated. These results are not surprising: since labor markets are not perfectly segmented, labor market shocks for one gender are correlated with labor market shocks for the other gender. When including only manufacturing employment, only the male IV predicts male employment and only the female IV predicts female employment.

The results of OLS, IV, and reduced form regressions for all formal sector employment and for manufacturing employment only are in Table 6. The IV coefficients indicate that a 1% increase in formal sector employment for men raises fertility rates by about 0.3%, depending on the specification, while an increase in female employment has a small and statistically insignificant impact on fertility rates. The impacts from manufacturing employment are roughly half of those including all formal sector employment. To put these numbers in perspective, Schaller (2016) finds that a 1% increase in unemployment, using a state-level analysis in the U.S., decreases birth rates by 2.6%. When she disaggregates by gender, she finds that decreases in male unemployment raise fertility, whereas decreases in female unemployment lower fertility, with stronger effects for men. My results are broadly consistent with hers, although the use of unemployment rates in her study makes the magnitudes not directly comparable with mine.

4.3 Reconciling the IV and OLS estimates

The OLS and IV results presented so far differ markedly in magnitude, which bears investigating. It seems unlikely that a local average treatment effect is inducing such substantial heterogeneity in responses. It is also difficult to think of an omitted variable that is more correlated with the instruments than the raw employment numbers leading to results of such magnitude.

If the Bartik-style construction of predicted employment is weakly correlated with actual employment, then the instrumental variables estimators can be very inconsistent. As documented in Bound et al. (1995), the IV estimates are biased in the direction of OLS in finite samples if the instruments satisfy the exclusion restriction. If the instruments do not fully satisfy the exclusion restriction, the degree and direction of inconsistency depends on the correlation between the instruments and the error term in the "structural" equation. It cannot be tested directly if the composition of industries within a municipality is correlated with an omitted variable that also affects fertility, but it is difficult to argue that the measure of predicted employment used here leads to a more inconsistent estimator than actual employment in establishing demand shocks. In that sense, the reduced form results can be construed as a bound on the results. Moreover, if the instruments do not satisfy the exclusion restriction, then they are likely to be biased in the same direction. Yet the estimated coefficients on male employment become much more positive and the estimated coefficients on female employment become much more negative.

The most plausible explanation for the discrepancy in magnitudes appears to be measurement error in the employment data. Although these data are at the administrative level, they only comprise the formal sector. Firms, especially smaller ones, can move in and out of the formal sector in Mexico. Some component of what may appear to be job gains or losses can simply be the result of how firms choose to classify themselves. It is well-known that panel data methods amplify the effects of measurement error. The higher the correlation between employment levels in adjacent time periods, the more inconsistent the fixed-effects estimator of the effects of employment is likely to be.

To investigate this more formally, consider a stripped-down first-differenced version of equation (9) regressing changes in fertility on changes in log employment:

$$\triangle y_{m,t} = \beta \triangle Emp_{m,t} + \triangle \varepsilon_{m,t}$$

It follows that

$$plim\ \hat{\beta} = \beta - \frac{\beta \sigma_v^2}{(1-\rho)\sigma_{emp}^2 + \sigma_v^2}$$

where σ_v^2 is the variance of the measurement error (for a derivation see Cameron and Trivedi, 2009) and ρ is the correlation between employment in adjacent time periods. The correlation between the logarithm of employment in adjacent years is close to one for some pairs of years. The equation above shows that such high serial correlation leads to a strong attenuation bias for the estimated coefficient on employment. A back-of-the-envelope calculation using the fixed effects estimates and the IV estimates in Table 4 and the variance in employment in adjacent time periods in the data (9.24) leads to an estimated variance in measurement error of 0.54. This appears reasonable given the magnitude of the variance in reported employment and is consistent with the instrumental variables estimates increasing by such large magnitudes.

4.4 An approach robust to weak instruments

In this section I propose an alternative set of results that is robust to potentially weak instruments. As noted previously, the large difference in magnitudes between the OLS and IV estimators is unlikely to be due to weak instruments. However, the values of the Kleibergen Paap Wald F statistic, especially for the estimates using all formal sector employment separately for men and women, may be a concern. A traditional rule of thumb from Staiger and Stock (1997) rejects the hypothesis of weak instruments if the first stage F statistic is above 10. That rule of thumb was revised in Stock and Yogo (2005), who formalize the arguments in Staiger and Stock (1997) and provide two criteria for establishing the presence of weak instruments. First, an instrument can be construed as weak if the relative bias of the IV estimator exceeds some level (say 10%) of the bias in the OLS estimator (bias test), and second, an instrument can be thought of as weak if the size of the Wald test exceeds a particular threshold (size test). Stock and Yogo suggest a test statistic that is equivalent to the first stage F statistic (if there is one endogenous regressor) or to the Cragg-Donald F statistic (if there is more than one endogenous regressor) and provide critical values for the F statistic based on the number of endogenous regressors and instruments, the maximum bias allowed (if using the bias test), and the estimation method used.

However, the critical values provided by Stock and Yogo crucially depend on the assumption of conditional homoscedasticity. In models containing heteroscedasticity, serial correlation, or clustering, as is the case in this paper, the critical values are no longer valid. Instead, I follow Kleibergen and Paap (2006), who suggest an F statistic that is robust to the presence of non-independent and identically distributed standard errors. To the best of my knowledge, however, the econometrics literature has not generated a formal test for the presence of weak instruments when errors are not i.i.d. and there are multiple endogenous regressors (see Montiel Olea and Pflueger, 2013, for a recent contribution when there is only one endogenous regressor).

An alternative to testing for weak instruments involves the construction of confidence sets that are robust to weak instruments; this approach exploits a duality to hypothesis testing. Given a test of $\beta = \beta_0$, one can create a confidence set for all values of β_0 for which the hypothesis is not rejected. Several tests have been proposed, including the conditional likelihood-ratio test (Moreira, 2003), the Lagrange multiplier test (Kleibergen, 2002; Moreira, 2002), and the Anderson-Rubin test (Anderson and Rubin, 1949). I present results from the latter because, unlike many of the other tests, it is generalizable to the case of more than one endogenous regressor, uses standard F critical values, and is regression-based, making the implementation straightforward.

Analysis using formal sector employment separated by gender has the smallest values of the Kleibergen-Paap Wald F statistic. This is likely due to the non-tradable sector being included in employment measures, as well as collinearity between male and female measures of employment. Thus, I redo the instrumental variables analysis in Table 6 (that is, of equation 9) by inverting the Anderson-Rubin test to create confidence sets for coefficients on men's and women's employment (see Baum et al., 2007).[30] The Anderson-Rubin test jointly tests $\beta = \beta_0$ and the exogeneity of the instruments. That is, rejection would occur if the null hypothesis is false or if the instruments are endogenous. The construction of the confidence set is done via grid testing. Figures 4, 5, and

[30]The procedure is implemented using the command weakiv in Stata provided by Findlay et al. (2013).

6 show the resulting confidence sets (i.e. coefficients within the acceptance region), as well as the rejection surface, for the effect of employment on general fertility rates. It makes little difference whether population is controlled for via a linear, quadratic, or cubic polynomial. Consistent with the earlier results assuming identification, positive (negative) shocks to employment for men lead to positive (negative) changes in fertility rates; that is, for any potential effect of female employment, the effect of male employment is positive. On the other hand, both positive and negative effects are consistent with female employment, although the bulk of the confidence set falls within negative parameter values for female employment.

5 Timing of births

Fertility choice is a dynamic process. In theory, it is possible for changes in wages and employment probabilities to affect aggregate fertility as well as the timing of births. Since my data tie yearly birth records to employment data, I cannot address directly the question of how long-lasting any effects from yearly variations in employment conditions may be. However, I can exploit additional information in birth certificates to indicate whether couples are substituting across years, and whether differential effects across the lifecycle are present.

In Table 7, I show results on fertility rates that separately consider only first births, second births, and births of third or higher parity.[31] IV results focusing only on manufacturing employment indicate that only births of third or higher parity have a statistically significant response to changes in employment for men. IV results using the full set of formal sector employment show progressively larger impacts for men and women for higher parities. If couples were timing births earlier in response to increased male employment, then we might expect women with no previous births to increase fertility, but both sets of results indicate that women with no previous births are the least responsive to changes in employment.

To further investigate whether couples time fertility across years, I augment equation (9) with a one year, two year, and three year lag in male and female employment. Table 8 focuses on the reduced form results using all formal sector employment, with each column introducing an additional lag.[32] Including lags of male and female employment substantially increases the

[31] I divide the total number of births in each parity group by the population of women aged 15-44.

[32] The sample size shrinks with every additional lag. Results are very similar if restricting to a consistent set of years across column. I only show reduced form results, not IV, because all specifications are restricted to contain

magnitude of both current-year coefficients, implying that a 10% positive shock to the men's labor demand index increases fertility rates by slightly over 3%. The coefficent on predicted female employment becomes negative and statistically significant from zero once a set of three lags is included, implying a 2% decreases in fertility rates from a 10% positive shock to the women's labor demand index. The net effect for women (the sum of the lags), however, is close to zero and statistically insignificant from zero.

These estimates provide evidence in favor of two points. First, current-period labor market structure has the biggest impact on fertility decisions. This may be due to liquidity constraints that prevent individuals from optimizing across years, to uncertain expectations about future earnings, or simply the salience (in behavioral terms) of current conditions. Second, substitution effects appear strongest in the current period for women, which is in line with theory: a short-run negative shock to employment may reduce the opportunity costs of having a child, while a short-run positive shock may induce some women to postpone fertility. Nevertheless, these effects appear transitory.

5.1 Time series properties of employment

To gain a better understanding of the time series properties of employment for men and women and explore the effect of employment dynamics, I follow the methodology in Baker, Benjamin, and Stanger (1999), who filter the minimum wage to study how high- and low-frequency cycles in the minimum wage affect the employment-to-population ratio. I decompose the natural log of employment as follows:

$$Emp_{m,g,t} = \tfrac{1}{2}(Emp_{m,g,t} - Emp_{m,g,t-1}) + \tfrac{1}{2}(Emp_{m,g,t} + Emp_{m,g,t-1}).$$

The first term in parentheses focuses on sharp, high-frequency changes between years to employment, whereas the second term in parentheses, a moving average, emphasizes slower-moving cycles in employment. In addition to using this filter, I also apply spectral analysis, which provides a set of tools to formally decompose a time series into components of varying periodicities (Iacobucci, 2005; Pollock, 1999). I generate a finite Fourier transform of the log of employment by decomposing it into a sum of orthogonal, sinusoidal functions, where

at most two instruments to avoid multicollinearity problems with multiple instruments. Additional lags are not statistically significant from zero.

$$Emp_{m,g,t} = \sum_{j=0}^{J} \left[\alpha_j cos\Big(\frac{2\pi jt}{T}\Big) + \beta_j sin\Big(\frac{2\pi jt}{T}\Big) \right].$$

The decomposition[33] in the frequency domain allows me to filter the data and determine whether the observed impacts on fertility rates are driven by slow-moving, low-frequency cycles or fast-moving, high-frequency cycles in employment.

Results for the log of predicted employment using the first filter are in Table 9. They indicate that cycles in employment for men at both frequencies increase fertility rates, although the impact appears to be strongest at the high frequency. For women, the effect is close to zero at the low frequency, but becomes substantial, negative, and statistically significant from zero at the high frequency. The frequency domain sheds light on which components of employment—that is, slow, secular changes, versus shocks associated with sharp responses to exogenous shifts in demand—drive fertility decisions. Because men's employment is associated mainly with income effects for the family, both slow and fast shifts in employment for men should have a positive impact on fertility, and indeed the results are consistent with the theory. For women, it appears that fast-moving, unexpected changes induce larger substitution effects, at least in the short run. These heterogeneous effects are masked in the earlier regression results, where both long and short cycles are combined into a single measure.

6 How do changes in employment affect wages?

The results in this paper indicate that men's employment has a robust, positive impact on fertility, whereas women's employment has weaker, negative impacts. In a wealthy country with high female labor force participation, we might expect that increases in women's employment opportunities generate large income effects that may counteract substitution effects from the increased opportunity cost of time, but this is unlikely to be the case in Mexico. As theory implies that changes in employment should lead to changes in wages, I directly evaluate the relationship between employment wages, wages, and sectoral gender intensity in this section.

First, changes in relative employment opportunities for women may increase their bargaining power. Research indicates that men have higher fertility preferences than women (Westoff and Bankole, 2002), so if increases in relative labor market opportunities for women lead to shifts

[33]The coefficients α_j and β_j are estimated separately for each municipality. Since I have nine years of data, T=9, and J is set to be the largest integer less than or equal to $T/2$.

in favor of their preferences, we should also expect to see declines in fertility. Although direct changes in wages and employment status can bring about a change in personal bargaining power, women who observe no change in job status or quality (that is, those women who are unaffected on the extensive margin studied in this paper) may still see enhanced bargaining power when their labor market opportunities improve: what matters is their outside option. These inframarginal changes work in the same direction as changes on the margin, so it seems puzzling that results are relatively weak for women's employment.

Do positive shocks to women's employment increase their earnings? To address this question, I evaluate how changes in labor markets in Mexico affect earnings. Since the instruments constructed in this paper exploit compositional differences across sectors, I focus my analysis on differences among sectors. As a starting point, I calculate mean earnings for men and women in each 4-digit sector in the IMSS data and plot the ratio of mean earnings for men to mean earnings for women against the proportion of men in each sector, along with the line of best fit, in Figure 7. The relative size of each circle indicates the number of individuals in each sector. The graph suggests that sectors with relatively more men have less gender-related earnings inequality. I evaluate this claim by regressing the ratio of earnings on the proportion of men in each sector. Table 10 shows that the relationship is strongly negative when pooling all years and comparing across sectors (column 1).

Some authors have argued that sectors that employ mainly women pay lower wages.[34] When I separately estimate how the proportion of men in each sector is correlated with earnings for men and women, however, I find that women's earnings have a small, albeit positive, correlation with how male-dominated the sector is (column 2). It does not seem to be the case that women working in female-centric industries are earning substantially lower earnings. Men's earnings, on the other hand, are strongly negatively correlated with how many men there are in the sector (column 3).

Suppose, for simplicity, that there are two types of jobs within each sector: a high-paying, high-skill type and a low-paying, low-skill type. If men work in both types of jobs, while women are only employed in the latter, then sectors with few men should have disproportionately more men in the high-paying sectors, leading to the type of wage inequality observed in the IMSS

[34]For an extensive discussion of forces shaping gender differences in employment in developing countries, see the World Bank Development Report (2011).

data. Alternately, it is possible that both women and men are distributed in high-skill jobs, but women earn lower relative wages in higher-paying occupations. To investigate these possibilities, I turn to the ENOE, a labor force survey in Mexico similar to the CPS in the U.S., to investigate the relationship between wages, gender composition of industries, and skill intensity.[35] Columns 4-6 in Table 10 reproduce the same results as in the IMSS data: wage inequality is negatively correlated with the proportion of men in each sector, and this is mainly due to men earning less on average in sectors with more men.

Table 11 shows results for regressions of either wage inequality, log of male wages, or log of female wages in each sector against the proportion of men in each sector and the proportion of each gender in high-skill occupations. I define individuals as being in high-skill occupations if they are employed as professionals or managers, which are the two highest paid occupational classifications.[36] Once the proportion of men and women in high-paying jobs is accounted for, the degree of male bias in each sector's workforce ceases to predict either wage inequality or wages for men or women.

To exploit differences across time and space in male versus female specialization, I return to equation (9) and use the IMSS data to evaluate how employment for men and women affects their earnings. Table 12 shows the impact of raw employment itself and the reduced form. Earnings for both men and women are positively correlated with expansions in male employment and negatively correlated with expansions in female employment. The predicted demand measure, which isolates purely demand factors, shows even larger impacts for women, with an elasticity of 0.19 for men's demand shocks and -0.21 for women's demand shocks. Men's earnings are unaffected by changes in demand for male or female labor. If all jobs were identical, a positive shock to demand would be expected to raise wages, but because women occupy lower wage positions, positive shocks to their demand actually lower their wages, conditional on men's demand shocks. In other words, separate demand shocks for women and men likely lead to compositional changes in the labor force. For instance, if low-skill women are induced into the labor force through positive shocks while high-skill women's labor market opportunities remain

[35] The ENOE contains a different classification of industries from the IMSS and consists of a smaller sample, so the results are not directly comparable. I limit the analysis to those sectors with at least 10 individuals of each gender in all years. Varying the cell size limit slightly does not affect the results qualitatively.

[36] The other choices in the survey are as follows: educational workers, clerks, industrial workers, merchants, transport operators, workers in personal services, protection and surveillance workers, and farmworkers.

unaffected, then average female wages may decline, even though wages for all women have weakly increased.

The difference in results for men's wages when focusing on supply and demand shocks versus demand shocks only can be reconciled if male supply is increasing more among educated groups, yet demand is more pronounced in lower-wage positions. This is consistent with the results in Campos-Vázquez (2013), who evaluates trends in wage inequality in Mexico following the passage of NAFTA. He finds that the supply of college-educated workers has grown rapidly, but high-skill occupations have not expanded enough to fill the new demand for these positions, leading to wage compression at the top.

Even if positive demand shocks for men do not translate into higher wages *among formal sector jobs*, formal sector jobs provide substantial health, social security, and childcare benefits that are not available in the informal sector. Second, increased labor force participation on the extensive margin translates into pure income effects. Thus, regardless of the effect on wages within the formal sector, increases (decreases) in male employment should lead to higher (lower) fertility, which is consistent with the results in this paper.

7 The impact of maquiladora employment on fertility

7.1 Context behind the expansion of maquiladoras

The previous section established how men's and women's formal employment impacts fertility decisions. In the following sections, I describe the history and context behind a particular type of formal sector employment in Mexico's export-assembly plants, introduce a different dataset and estimation strategy, and show results consistent with the earlier set of findings.

After the termination of an agreement in which Mexico sent seasonal farm laborers to the U.S., Mexico faced the prospect of a large pool of unemployed people living in the north.[37] To generate incentives for firms to locate in the northern border region, in 1965 the Mexican government introduced a plan called the Border Industrialization Program, which allowed full foreign ownership of establishments in Mexico. Although these establishments, called maquiladoras,

[37] This was known as the Bracero Program. The combination of demand for low-wage agricultural workers in the U.S and excess supply of labor in Mexico allowed the program to operate until 1964, when pressure from U.S. labor unions led to the termination of the program.

were highly regulated initially, over time these regulations were relaxed to attract more foreign investment.[38]

In particular, Mexico's poor economic performance in the 1980s led to a series of economic reforms and the liberalization of trade conditions, which led to large-scale growth in maquiladora employment starting with the late 1980s. Figure 9 shows the growth of employment for each gender over the time period studied. Although aggregate employment for men and women closely tracked each other, the graph masks how gender-segmented factories are: for instance, factories specializing in textiles employ mainly women, while factories manufacturing electronics employ mainly men. This segmentation applies not only at the establishment-level, but also at the regional level, as shown in figure 10, which illustrates the average female share of maquiladora employment across Mexican municipalities.

7.2 How maquiladoras affect fertility

The growth in maquiladora line-employment happened at both the extensive and intensive level: new plants opened, and plants that continued to operate increased in size. Figure 11 shows the density of employment at the establishment-level in 1990, 2000, and 2006. Although there is a clear clear shift toward bigger sizes over time, typical sizes of maquiladoras remained small, with median employment below 100 people.

To study the impact of the rapid growth in employment for both men and women across sectors and regions on fertility decisions, I focus on establishment-level changes in employment from 1990-2006. I instrument for net new jobs for women (men) in export assembly plants within each municipality with net new jobs for women (men) in these plants that are solely due to large single-firm expansions/openings and contractions/closings (i.e. a change of least 50 individuals in a given year). As figure 11 indicates, maquiladoras are quite small, so these are large changes relative to the size of the establishment.[39].

For the exclusion restriction to hold, I require that firms do not respond to fertility decisions with large expansions or contractions, conditional on the fixed effects and controls for population in the estimating equation. This seems especially plausible in my context, as I focus only on maquiladoras. It is likely that large changes in employment at these plants involve high fixed

[38]For instance, maquiladoras were required to be within twenty kilometers of the border and all output had to be exported. The bulk of employment has stayed near the border, however, as shown in figure 8.

[39]Similar identification approaches have been employed by Atkin (2012) and Ananat et al. (2011)

costs and are due to shocks in external demand from other countries (Atkin, 2012).

I create the employment variables as follows. I aggregate male and female employment separately across all establishments within a municipality to create the main independent variable for each gender.[40] To standardize the variable, I divide employment by the 1990 population of working-age men or women.[41] I use the 1990 baseline year in the data for two reasons: the main one is that the denominator may vary along with other conditions in the municipality, and since maquiladora employment is very small relative to all employment in most places, this can lead to a severe bias in the variable. (For instance, if maquiladora employment increases for men, but population shifts lead to a large enough increase in the denominator, then the term may decrease even when maquiladora employment goes up.) Second, I interpolate the population between census years, which may introduce an additional source of bias. Hence, using the 1990 population creates a standardized measure to track how changes in maquiladora employment for each gender relative to the baseline population of that gender affect fertility and marriage outcomes.

To construct the instrumental variable, I first difference employment for each gender at each establishment across years. Keeping only the sample of establishments that contains a change of at least 50 individuals from the previous year, I then aggregate employment to the municipality level. Finally, I standardize by the 1990 population of that group, as with the main independent variable.

Table 13 shows the first stage results for the full sample of municipalities in columns 1 and 2. Since a large proportion of municipalities has only one or two small factories, which are unlikely to have a large impact on fertility if the population is large, I also limit the sample to those municipalities that have at least 1% employment of either gender in at least one year, which I call the restricted sample.[42] The results for this sample are in columns 3 and 4. In all cases, the Kleibergen-Paap Wald F statistic is above 50, and both the instruments for male and female

[40] I calculate average yearly employment from the monthly observations. Establishments can enter or exit the data in any month, and some establishments do not contain data for some months between entry and exit (i.e. due to temporary shut-downs or not answering the survey, which was required). For the set of plants missing months, I tried three techniques: calculating average yearly employment only for those months in the data, imputing employment based on previous months, or simply imputing zero. Results are similar for all approaches.

[41] The denominator is the population of men or women aged 15-44 in that municipality, linearly interpolated from the decennial census.

[42] Results are similar if restricting to 3% or 5% employment, although the sample size is substantially reduced.

employment are strongly correlated with actual employment.

Table 14 shows the OLS, IV, and reduced form results for the full sample, as well as the restricted sample. The reduced form results indicate that men's employment has a large, positive impact on fertility, while the impact of women's employment is negative, albeit imprecisely estimated. The results imply that if, say, expansions or openings of new maquiladoras lead to an additional 5% of the population to work in maquiladoras, this would translate into a 0.01 increase in the log of the fertility rate. Although this may seem small, it is important to keep in mind that maquiladoras make up a small proportion of the labor force in most municipalities.

8 Migration, robustness, and alternative specifications

8.1 Migration

If individuals migrate in or out of municipalities experiencing changes in labor markets, then resulting changes in fertility may be due to changes in the composition of individuals living there and not to actual changes in behavior. There are three types of migration that may affect the results. The first type concerns local migration: individuals may live in one municipality and work in another. Since I link aggregate employment to aggregate fertility in the same municipality, we can think of some births in the data as being assigned to the wrong municipality. If men and women are equally likely to commute to other municipalities for work, this can lead to measurement error that should bias the estimators for both groups' employment toward zero. If men are more likely to commute to different municipalities, as seems likely the case in this setting (unfortunately, I cannot test this directly), then the male estimator should be more attenuated. Since I find larger results for men rather than women, it appears unlikely that this type of bias is driving the results.

A second type of migration concerns moves across the country. As Mexico has industrialized, individuals living in poor rural communities in the south have moved to the north to work in maquiladoras and related enterprises. Suppose some young women migrate to the north to work in factories and then return home to their rural communities, which have little formal sector employment, to have children. This should bias me in favor of finding a larger negative impact of female employment, yet I find no statistically significant impact.

Another possibility concerns selective migration into municipalities that undergo demand shocks. This type of migration can result from either population movements within Mexico, or cross-country migration (such as Mexican migrants returning from the U.S.). To probe this further, I re-run equation (9), except I exclude population controls from the right-hand side and instead use them as the outcome variables. Instrumenting employment for men and women with predicted employment, I find that labor demand shocks for men and women are not associated with changes in population for either men or women (Table 15). Although it is possible that in-migration of a selected sample is exactly balanced by out-migration, these results suggest that migrants are not choosing municipalities based on labor demand shocks, at least in the short run. This is reassuring: changes in raw employment are due to shifts in supply and demand, whereas the Bartik-style instrument should be isolating only changes in labor demand.

8.2 Robustness

This paper uses an instrument in the tradition of Bartik (1991) to isolate an exogenous predictor of employment. To test whether the results hold using a traditional Bartik-style instrumental variable, I create

$$Bartik\ instrument = \sum_{ind} \frac{Emp_{m,g,ind,t=0}}{Emp_{m,g,t=0}}\ log(Emp_{g,ind,t} - Emp_{m,g,ind,t}).$$

The definition of the terms is the same as before, except the denominator in the fraction is now equal to employment in municipality m for group g (men or women) in the baseline time period. This instrument is often constructed in first-differenced form and used to predict employment growth. Two municipalities with exactly the same industrial composition would have the same values of this measure (ignoring the subtraction of own-employment in parentheses), as they would be expected to have similar levels of growth or decline in employment; in other words, it is invariant to the municipality's population. However, including municipal fixed effects means the Bartik instrument effectively isolates the same type of variation as the instrument used in the paper, and results replicating Table 6 (shown in Table 16) using the instrument are qualitatively similar, though formal sector estimates are slightly larger in magnitude.

Finally, I replace linear state trends with state-by-year fixed effects and again reproduce the main analysis in Table 17. Such a specification flexibly controls for any unobservable shocks specific to states in any particular year. Of course, municipalities located close to each other are more likely to have similar labor market structure and thus to face similar types of labor market shocks. There is no reason to believe that identification based on the remaining variation across municipalities leads to more consistent estimators, given that state-by-year fixed effects may absorb "too much" labor market variation[43]. In practice, the choice of linear, quadratic, or state-by-year fixed effects matters little: for instance, if including year fixed effects, municipality fixed effects, a cubic in log population for men and women, and year-by-state fixed effects, the impact of male formal sector employment decreases slightly from 0.32 to 0.30, and standard errors increase as well, but the results for men remain statistically significant at the 5% level, and results for women remain small and statistically insignificant. It does not appear that nonlinear, unobservable deviations correlated with labor market changes within states are driving the earlier results using all formal sector employment. Specifications using only manufacturing employment result in very similar estimated impacts for job opportunities for men, but standard errors increase and the coefficients are statistically significant only at the 10% level.

9 Conclusion

The question of how labor market opportunities shape decisions about the family has long interested economists. In some developed countries, fertility rates are arguably too low, and policymakers have invested large sums in relaxing perceived constraints to having children, such as providing flexible working arrangements for young mothers, daycare, or simply lump sum payments. On the other hand, in many developing countries fertility rates remain stubbornly high, making countries that have gone through a large transformation in family structure potentially useful guides for their own experiences.

Mexico has experienced a dramatic fall in fertility, as well as a steady increase in labor market opportunities for young women, driven in part by expansion in trade-oriented manufacturing jobs. In more recent years, sectors traditionally employing young women, such as textile manufacturing, have become less competitive as production has shifted to China and other Asian economies with lower labor costs. This paper evaluates how expansions and contractions in employment that vary across municipalities in Mexico and over time affect fertility. Because fertility and employment are joint household decisions, I focus on aggregate changes in demand conditions for both men and women, using a measure of predicted employment that exploits

[43] An analogous point is made in Bound and Solon (1999), who show that identifying the returns to schooling based on twin-comparisons may not lead to more consistent estimation.

labor market segmentation by sex and the industrial structure of each local labor market, to isolate exogenous demand shocks for each group. The findings are robust to an alternative identification strategy based on the expansion and contraction of maquiladoras during the 1990s and early 2000s.

That women's employment does not appear to have a significant net effect on fertility in the short run may indicate that other broader, long-run social changes play an even greater role in explaining variation in fertility rates across time and space. In particular, economists and demographers have documented that increases in educational attainment or health (leading to a preference for child quality over quantity), urbanization, and better access to and knowledge of contraceptives have all been associated with declines in fertility.[44] By focusing on short-run shocks to employment within a single country, which arguably keeps these longer-horizon variables fixed, I am able to identify the causal impact of shocks to labor market opportunities on fertility.

The findings in this paper show that employment dynamics for both men and women matter: positive demand shocks to men's labor lead to positive changes in fertility, driven by both high- and low-frequency trends in employment. Current-period positive demand shocks shocks to women's labor reduce fertility, indicating that substitution effects dominate income effects in the short run, but the net effect is negligible. When focusing on the effects of the business cycle, I show that couples do appear to time births, consistent with dynamic models of fertility, but current economic conditions play the biggest role in determining fertility, indicating the presence of liquidity constraints or uncertainty about future employment expectations.

References

Ahn, Namkee, and Pedro Mira. 2002. "A Note on the Changing Relationship Between Fertility and Female Employment Rates in Developed Countries." *Journal of Population Economics*, 15(4): 667–682.

Ananat, Elizabeth Oltmans, Anna Gassman-Pines, Dania V. Francis, and Christina M. Gibson-Davis. 2011. "Children Left Behind: The Effects of Statewide Job

[44]For a discussion of theories of the fertility transition, see Mason (1997) and Guinnane (2011) for the perspectives of a demographer and an economist, respectively, and the references therein.

Loss on Student Achievement." NBER Working Papers 17104, National Bureau of Economic Research, Inc.

Anderson, T. W., and Herman Rubin. 1949. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *Ann. Math. Statist.*, 20(1): 46–63.

Ashraf, Nava, Erica Field, and Jean Lee. 2014. "Household Bargaining and Excess Fertility: An Experimental Study in Zambia." *American Economic Review*, 104(7): 2210–37.

Atkin, David. 2012. "Endogenous Skill Acquisition and Export Manufacturing in Mexico." NBER Working Papers 18266, National Bureau of Economic Research, Inc.

Baker, Michael, Dwayne Benjamin, and Shuchita Stanger. 1999. "The Highs and Lows of the Minimum Wage Effect: A Time-Series Cross-Section Study of the Canadian Law." *Journal of Labor Economics*, 17(2): 318–50.

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?*.: W.E. Upjohn Institute for Employment Research.

Basu, Kaushik. 2006. "Gender and Say: A Model of Household Behaviour with Endogenously Determined Balance of Power." *The Economic Journal*, 116(511): 558–580.

Baum, Christopher, Mark Schaffer, and Steven Stillman. 2007. "Enhanced Routines for Instrumental Variables/Generalized Method of Moments Estimation and Testing." *Stata Journal*, 7(4): 465–506.

Becker, Gary S. 1960. "An Economic Analysis of Fertility." In *Demographic and Economic Change in Developed Countries.*: National Bureau of Economic Research, Inc, 209–240.

Becker, Gary S. 1965. "A Theory of the Allocation of Time." *The Economic Journal*, 75(299): 493–517.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-In-Differences Estimates?" *The Quarterly Journal of Economics*, 119(1): 249–275.

Binelli, Chiara, and Orazio Attanasio. 2010. "Mexico in the 1990s: the Main Cross-Sectional Facts." *Review of Economic Dynamics*, 13(1): 238–264.

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.

Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, 23(1): 1–76.

Blundell, Richard, Pierre-Andre Chiappori, and Costas Meghir. 2005. "Collective Labor Supply with Children." *Journal of Political Economy*, 113(6): 1277–1306.

Bound, John, and Harry J. Holzer. 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s." *Journal of Labor Economics*, 18(1): 20–54.

Bound, John, David A. Jaeger, and Regina M. Baker. 1995. "Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogeneous Explanatory Variable is Weak." *Journal of the American Statistical Association*, 90(430): 443–450.

Butz, William P., and Michael P. Ward. 1979. "The Emergence of Countercyclical U.S. Fertility." *American Economic Review*, 69(3): 318–28.

Campos-Vázquez, Raymundo M. 2013. "Why Did Wage Inequality Decrease in Mexico After NAFTA?" *Economía Mexicana NUEVA ÉPOCA*, 0(2): 245–278.

Chiappori, Pierre-Andre. 1992. "Collective Labor Supply and Welfare." *Journal of Political Economy*, 100(3): 437–67.

Currie, Janet, and Hannes Schwandt. 2014. "Short- and Long-Term Effects of Unemployment on Fertility." *Proceedings of the National Academy of Sciences*, 111(41): 14734–14739.

Dehejia, Rajeev, and Adriana Lleras-Muney. 2004. "Booms, Busts, and Babies' Health." *The Quarterly Journal of Economics*, 119(3): 1091–1130.

Del Boca, Daniela. 2002. "The Effect of Child Care and Part Time Opportunities on Participation and Fertility Decisions in Italy." *Journal of Population Economics*, 15(3): 549–573.

Do, Quy-Toan, Andrei A. Levchenko, and Claudio Raddatz. 2016. "Comparative Advantage, International Trade, and Fertility." *Journal of Development Economics*, 119(C): 48–66.

Doepke, Matthias, and Fabian Kindermann. 2016. "Bargaining over Babies: Theory, Evidence, and Policy Implications." Working Paper 22072, National Bureau of Economic Research.

Gronau, Reuben. 1977. "Leisure, Home Production, and Work–the Theory of the Allocation of Time Revisited." *Journal of Political Economy*, 85(6): 1099–1123.

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(C): 1–15.

Henriette Engelhardt, Alexia Prskawetz. 2004. "On the Changing Correlation between Fertility and Female Employment over Space and Time." *European Journal of Population*, 20(1): 35–62.

Hotz, Joseph V., Jacob Alex Klerman, and Robert J. Willis. 1997. "The Economics of Fertility in Developed Countries." In *Handbook of Population and Family Economics*. eds. by Mark Rosenzweig, and Oded Stark, 1A: Elsevier, 275–347.

Iacobucci, Alessandra. 2005. "Spectral Analysis for Economic Time Series." In *New Tools of Economic Dynamics*. eds. by Jacek Leskow, Lionello F. Punzo, and Martín Puchet Anyul, Berlin: Springer, 203–219.

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 Michèle Tertilt. 2010. "Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?" In *Demography and the Economy.*: National Bureau of Economic Research, Inc, 43–100.

José María Da Rocha, Luisa Fuster. 2006. "Why Are Fertility Rates and Female Employment Ratios Positively Correlated across O.E.C.D. Countries?" *International Economic Review*, 47(4): 1187–1222.

Karaman Örsal, Deniz, and Joshua R. Goldstein. 2010. "The Increasing Importance of Economic Conditions on Fertility." MPIDR Working Papers WP-2010-014, Max Planck Institute for Demographic Research, Rostock, Germany.

Kemnitz, Alexander, and Marcel Thum. 2015. "Gender Power, Fertility, and Family Policy." *Scandinavian Journal of Economics*, 117(1): 220–247.

Kleibergen, Frank. 2002. "Pivotal Statistics for Testing Structural Parameters in Instrumental Variables Regression." *Econometrica*, 70(5): 1781–1803.

Kleibergen, Frank, and Richard Paap. 2006. "Generalized Reduced Rank Tests Using the Singular Value Decomposition." *Journal of Econometrics*, 133(1): 97–126.

Kögel, Tomas. 2004. "Did the Association between Fertility and Female Employment within OECD Countries Really Change Its Sign?" *Journal of Population Economics*, 17(1): 45–65.

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.

Leibenstein, Harvey. 1957. *Economic Backwardness and Economic Growth: Studies in the Theory of Economic Development.*: Wiley.

Lim, Lin Lean. 2009. *Female Labour-Force Participation.* 195–212: United Nations.

Lindo, Jason M. 2010. "Are Children Really Inferior Goods?: Evidence from Displacement-Driven Income Shocks." *Journal of Human Resources*, 45(2): 301–327.

Macunovich, Diane J. 1995. "The Butz-Ward Fertility Model in the Light of More Recent Data." *Journal of Human Resources*, 30(2): 229–255.

Maloney, William F. 2004. "Informality Revisited." *World Development*, 32(7): 1159 – 1178.

Marcouiller, Douglas, Veronica Ruiz de Castilla, and Christopher Woodruff. 1997. "Formal Measures of the Informal-Sector Wage Gap in Mexico, El Salvador, and Peru." *Economic Development and Cultural Change*, 45(2): 367–392.

Mason, Karen. 1997. "Explaining Fertility Transitions." *Demography*, 34(4): 443–454.

McDonald, John. 1983. "The Emergence of Countercyclical US Fertility: A Reassessment of the Evidence." *Journal of Macroeconomics*, 5(4): 421–436.

Mincer, Jacob. 1963. "Market Prices, Opportunity Costs, and Income Effects." *Measurement in Economics* 67–82.

Mishra, Vinod, Ingrid Nielsen, and Russell Smyth. 2010. "On the Relationship Between Female Labour Force Participation and Fertility in G7 Countries: Evidence from Panel Cointegration and Granger Causality." *Empirical Economics*, 38(2): 361–372.

Moreira, Marcelo J. 2002. "Tests with Correct Size in the Simultaneous Equations Model." Ph.D. dissertation, University of California, Berkeley.

Moreira, Marcelo J. 2003. "A Conditional Likelihood Ratio Test for Structural Models." *Econometrica*, 71(4): 1027–1048.

Ñopo, Hugo. 2012. *New Century, Old Disparities: Gender and Ethnic Earnings Gaps in Latin America and the Caribbean.*: The World Bank.

Olea, Jose Luis Montiel, and Carolin Pflueger. 2013. "A Robust Test for Weak Instruments." *Journal of Business & Economic Statistics*, 31(3): 358–369.

Álvarez Parra, Fernando, Luis Brandao-Marques, and Manuel Toledo. 2013. "Durable Goods, Financial Frictions, and Business Cycles in Emerging Economies." *Journal of Monetary Economics*, 60(6): 720 – 736.

Pollock, D.S.G. 1999. "The Discrete Fourier Transform." In *Handbook of Time Series Analysis, Signal Processing, and Dynamics.* ed. by D.S.G. Pollock, London: Academic Press, 399 – 425.

Rasul, Imran. 2008. "Household Bargaining over Fertility: Theory and Evidence from Malaysia." *Journal of Development Economics*, 86(2): 215–241.

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." In *Handbook of Population and Family Economics*. eds. by Mark Rosenzweig, and Oded Stark, 1A: Elsevier, 349–430.

Sivasankaran, Anitha. 2014. "Work and Women's Marriage, Fertility and Empowerment: Evidence from Textile Mill Employment in India."

Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias." *The Quarterly Journal of Economics*, 109(1): 1–25.

Staiger, Doug, and James Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica*, 65(3): 557–586.

Stock, James, and Motohiro Yogo. 2005. "Testing for Weak Instruments in Linear IV Regression." In *Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg.*: Cambridge University Press.

Westoff, Charles, and Akinrinola Bankole. 2002. *Reproductive Preferences in Developing Countries at the Turn of the Century.* DHS Comparative Reports: ORC Macro.

Willis, Robert J. 1973. "A New Approach to the Economic Theory of Fertility Behavior." *Journal of Political Economy*, 81(2): S14–64.

World Bank. 2011. *World Development Report 2012: Gender Equality and Development.*: The World Bank.

Figure 1: Fertility rates across Mexico

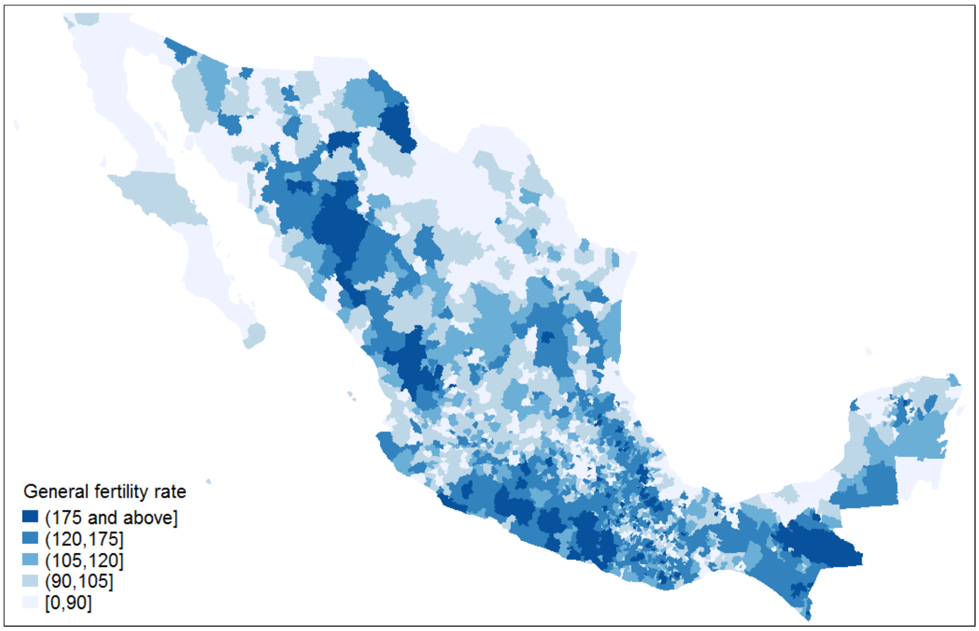

Notes: The general fertility rate is calculated as the number of births per 1000 women aged 15-44. Data are calculated from Mexico's National Institute of Statistics and Geography natality and census statistics in 2010.

Figure 2: Fertility rate over time

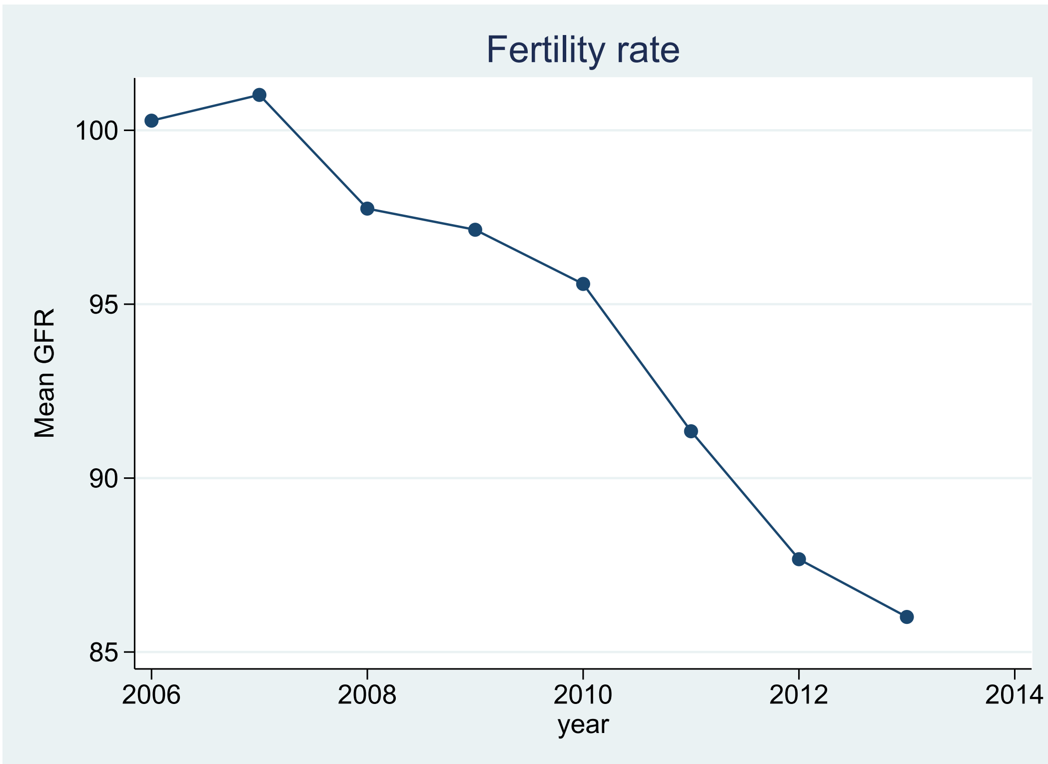

Notes: The general fertility rate is calculated as the number of births per 1000 women aged 15-44. Data are calculated from Mexico's National Institute of Statistics and Geography natality and census statistics.

Figure 3: Standardized proportion of men and women in employment

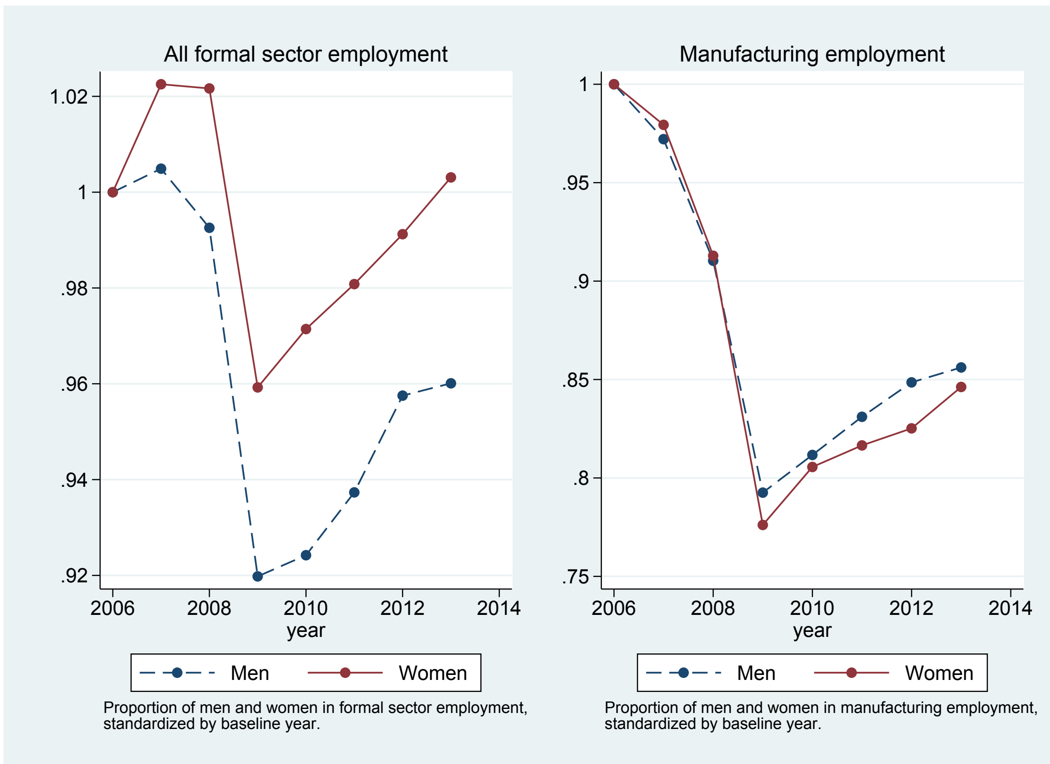

Notes: Trends in formal sector employment, standardized by employment in the first year, are shown. Employment data come from the Mexican Social Security Institute.

Figure 4: Confidence set with linear population controls

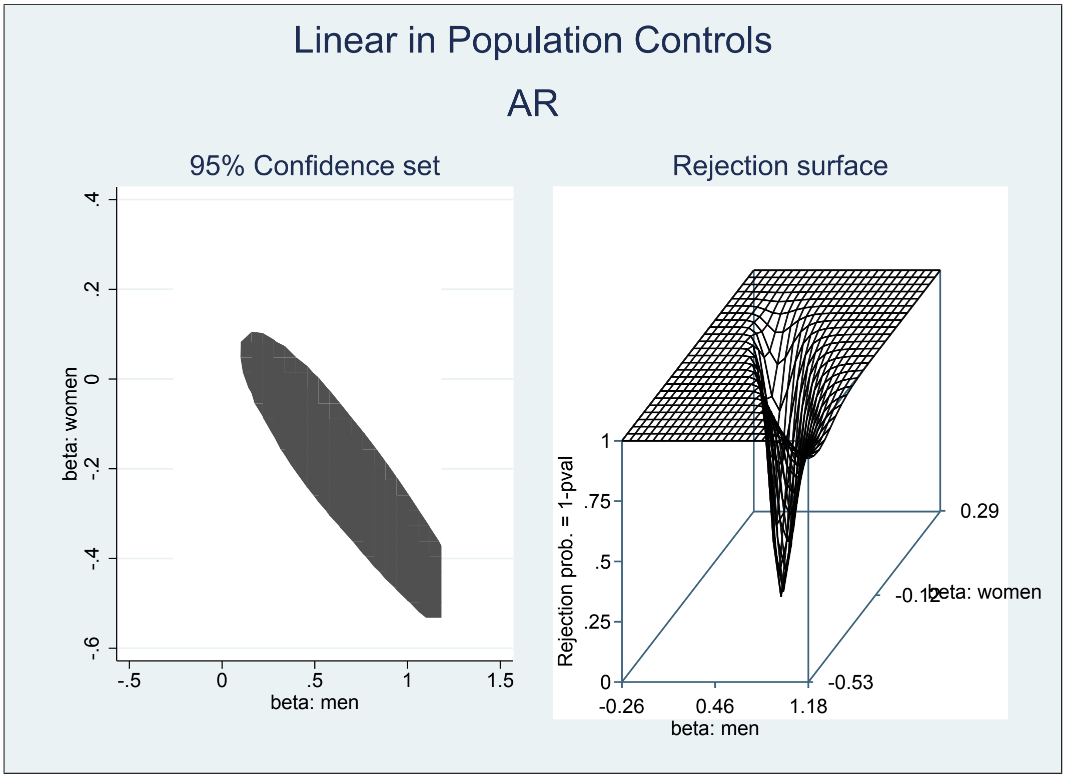

Notes: Confidence set and rejection surface for the Anderson-Rubin test.

Figure 5: Confidence set with quadratic population controls

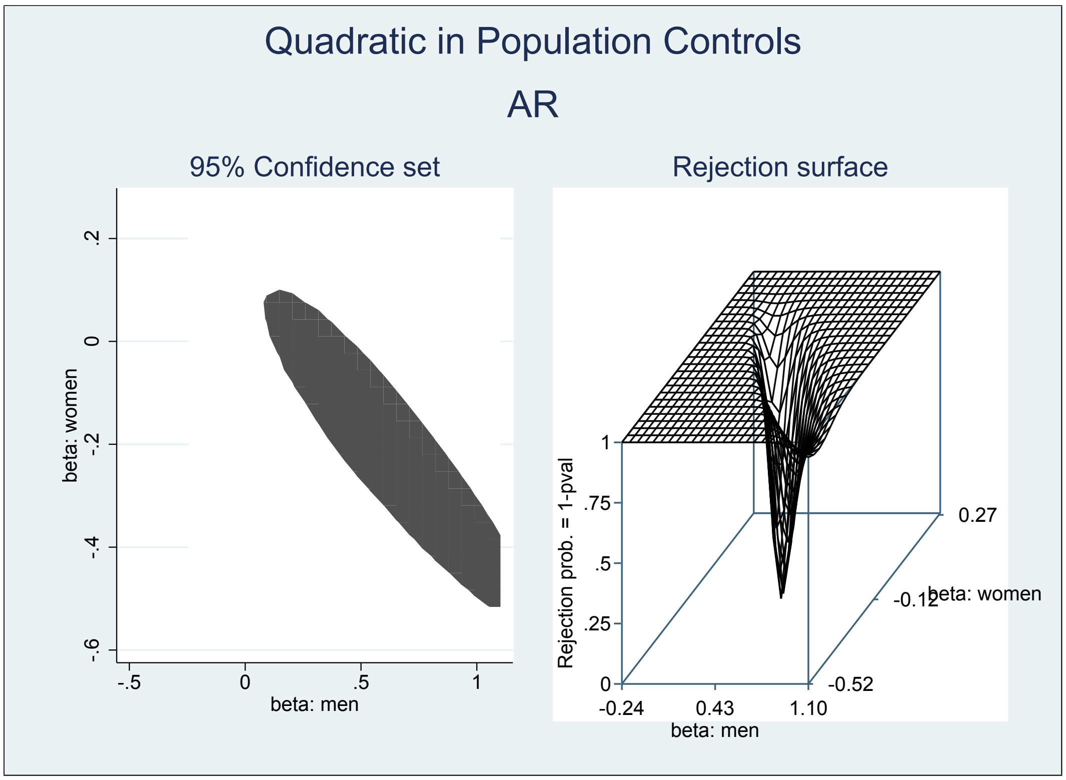

Notes: Confidence set and rejection surface for the Anderson-Rubin test.

Figure 6: Confidence set with cubic population controls

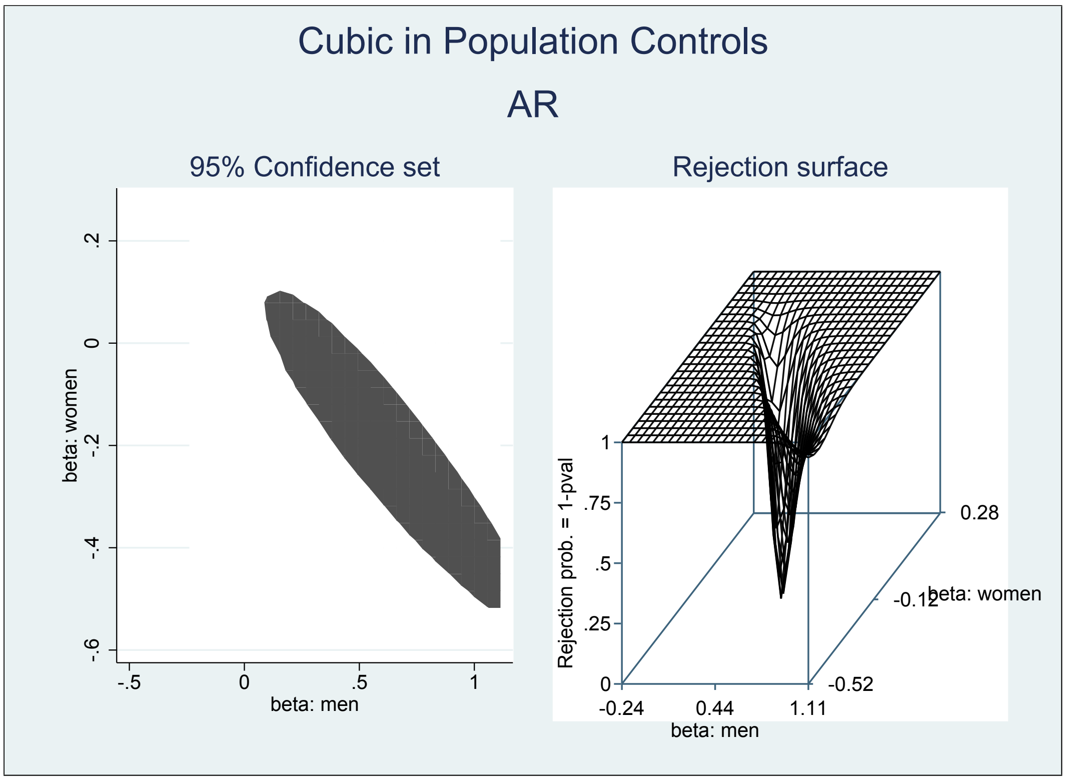

Notes: Confidence set and rejection surface for the Anderson-Rubin test.

Figure 7: Male to female earnings ratio in 2005

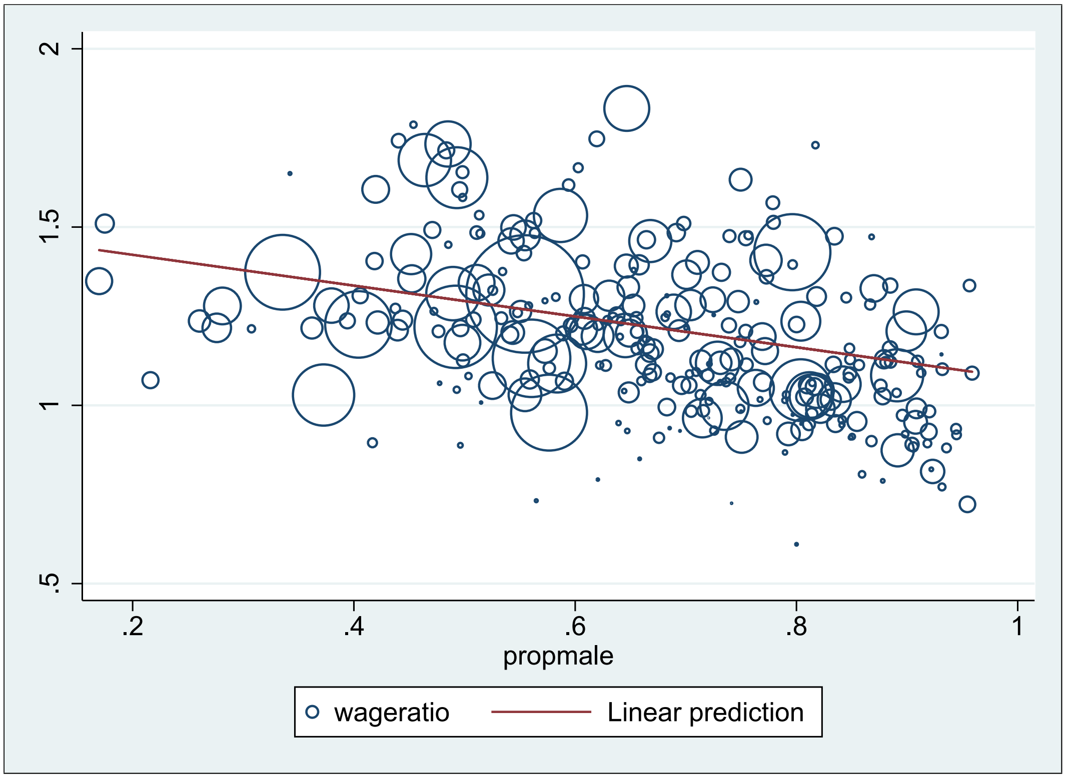

Notes: The size of each circle represents the size of each sector. The line of best fit is shown. Data are calculated from the Mexican Social Security Institute.

Figure 8: Municipalities by maximal share of employment in maquiladoras

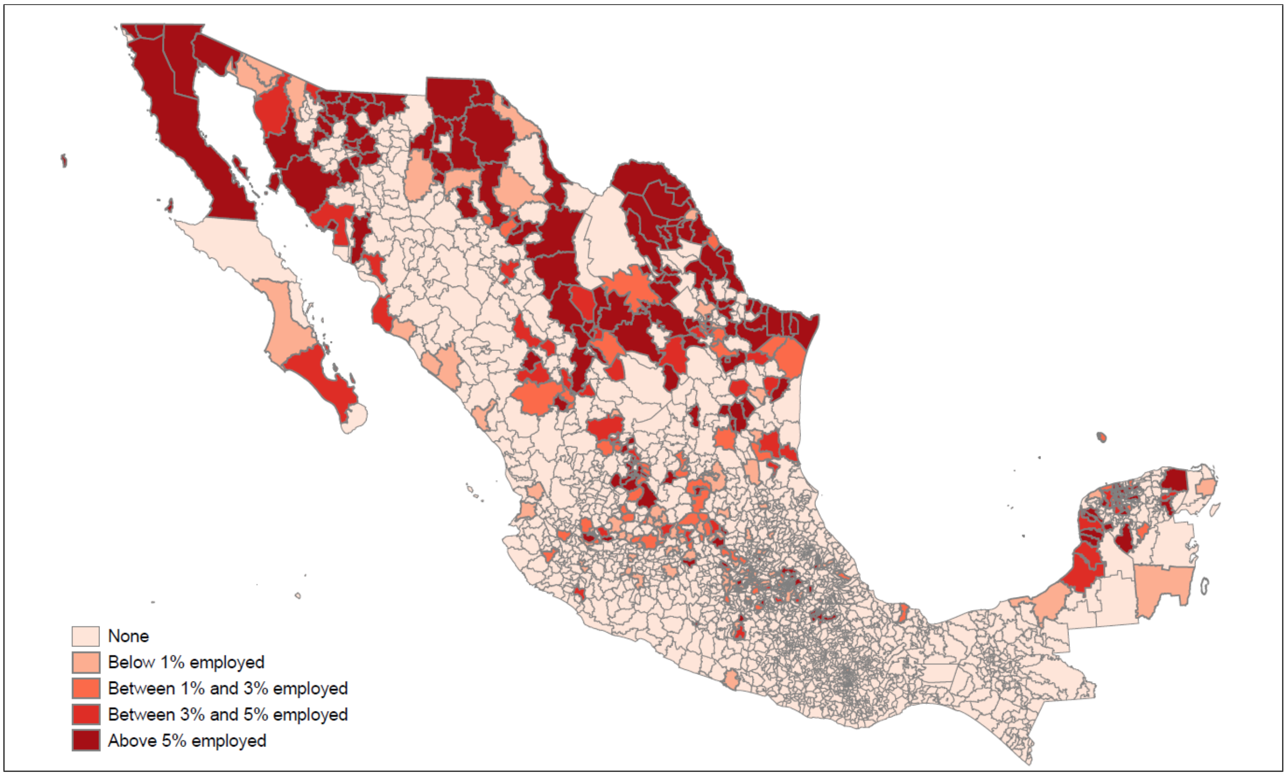

Notes: The share is calculated separately for men and women by taking aggregate employment in maquiladoras and dividing by the estimated population of that group aged 15-44 in each year. Then the maximal share is calculated across all years in the dataset. Data are calculated from the Maquiladora Export Industry Dataset, which is provided by Mexico's National Institute of Statistics and Geography.

Figure 9: Employment for men and women in maquiladoras

Trends in aggregate line employment for men and women at maquiladoras are shown. Data are calculated from the Maquiladora Export Industry Dataset, which is provided by Mexico's National Institute of Statistics and Geography.

Figure 10: Municipalities by female share of maquiladora labor force

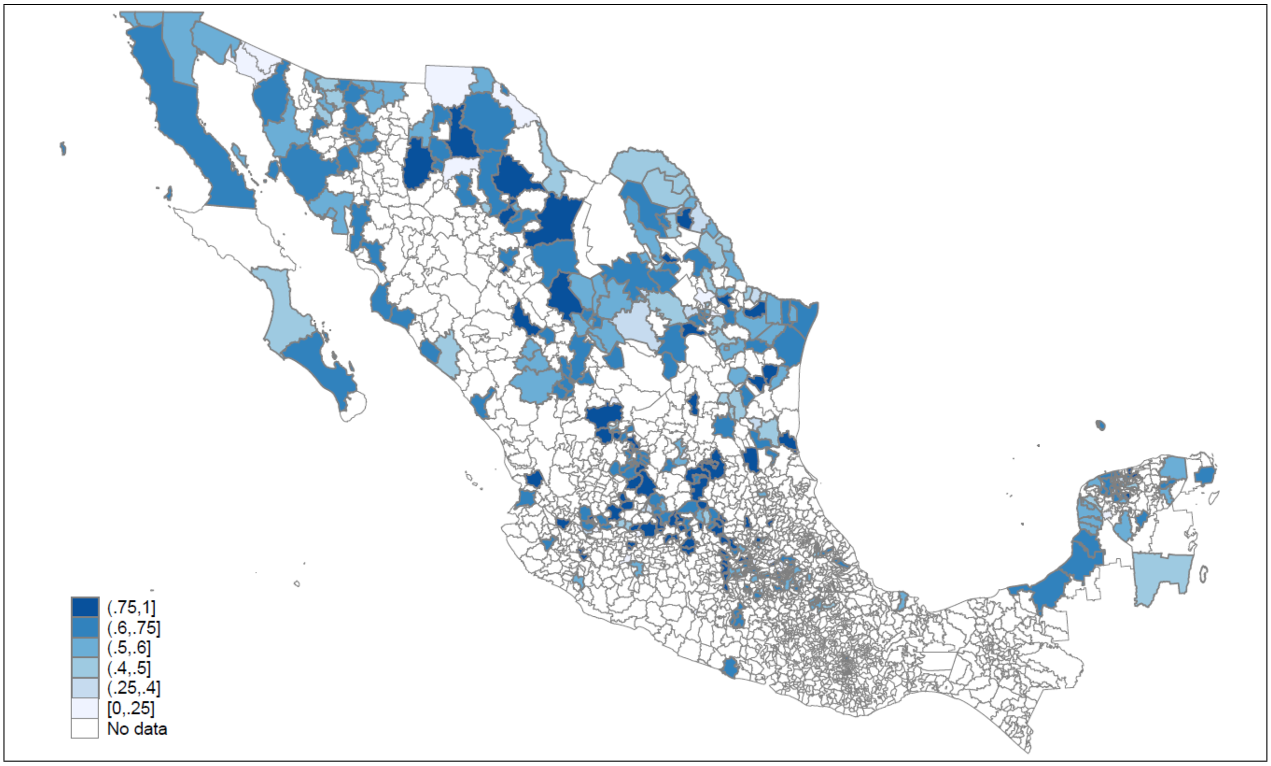

Notes: The share is calculated separately for each municipality. The "No data" category refers to municipalities with no maquiladora exmployment. Data are calculated from the Maquiladora Export Industry Dataset, which is provided by Mexico's National Institute of Statistics and Geography.

Figure 11: Density of employment across maquiladoras

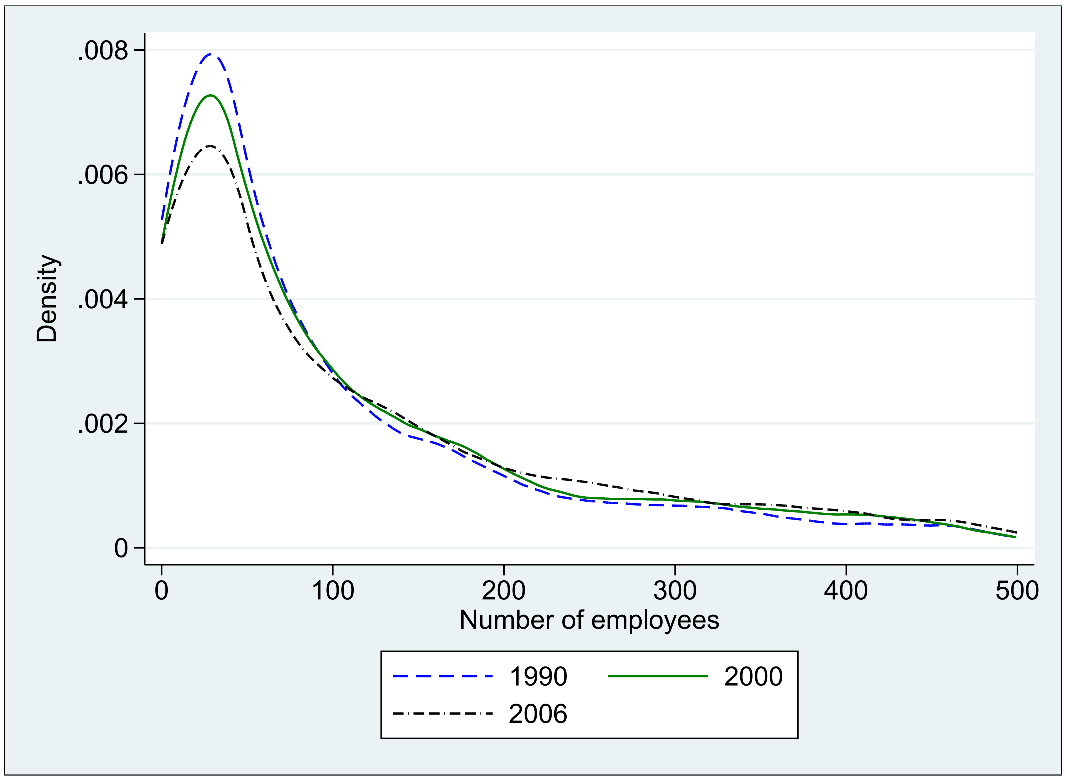

Notes: The density of line employment at maquiladoras is shown for 1990, 2000, and 2006, using the Epanechnikov kernel with the "optimal width" that minimizes mean integrated squared error if the data were Gaussian.

Table 1: General fertility rates across Mexican federal entities

State	Fertility rate in 2005	Fertility Rate in 2013
Aguascalientes	93	88
Baja California	89	77
Baja California Sur	94	71
Campeche	87	86
Coahuila de Zaragoza	92	91
Colima	82	78
Chiapas	137	131
Chihuahua	94	79
Distrito Federal	69	66
Durango	108	101
Guanajuato	97	80
Guerrero	149	114
Hidalgo	108	78
Jalisco	91	83
Mexico	92	79
Michoacn	103	94
Morelos	93	82
Nayarit	103	86
Nuevo Len	83	79
Oaxaca	112	83
Puebla	115	98
Quertaro	96	81
Quintana Roo	91	76
San Luis Potos	101	83
Sinaloa	94	83
Sonora	93	81
Tabasco	102	84
Tamaulipas	97	74
Tlaxcala	103	80
Veracruz	102	83
Yucatn	81	74
Zacatecas	101	90

General fertility rates are calculated as the sum of all births per 1000 women of ages 15-44 in each federal entity. The denominator is linearly interpolated from decennial census data from INEGI.

Table 2: Summary statistics

	Mean	SD
General fertility rate	89.59	18.45
Formal sector employment	59645	84827
Male formal sector employment	23309	34252
Female formal sector employment	36336	50799
Male manufacturing employment	12332	19789
Female manufacturing employment	7139	13634
Female earnings	171.62	50.68
Male earnings	212.32	62.82
Female population (15-44)	115411	118753
Male population (15-44)	108985	113540
GFR - 1st births only	39.57	9.45
GFR - 2nd births only	24.97	4.87
GFR - 3rd births and above	25.03	10.38
Observations	9,108	
Municipalities	1,012	

Summary statistics are shown across years 2005-2013, weighted by population of women aged 15-44 in each municipality. Author's calculations from Mexico's National Institute of Statistics and Geography data.

Table 3: Estimation of the relationship between employment and fertility

	(1) Log (GFR)	(2) Log (GFR)	(3) Log (GFR)	(4) Log (GFR)	(5) Log (GFR)	(6) Log (GFR)
All emp.	0.031***					
	(0.0096)					
Manufacturing emp.		0.012**				
		(0.0058)				
Male emp.			0.033***		0.032***	
			(0.014)		(0.014)	
Female emp.			-0.0096		-0.00048	
			(0.013)		(0.016)	
Male manuf. emp.				0.022**		0.022**
				(0.0088)		(0.0099)
Female manuf. emp.				-0.0055		-0.0037
				(0.0063)		(0.0073)
Male informal emp.					-0.00073	-0.0012
					(0.0047)	(0.0047)
Female informal emp.					-0.0045	-0.0045
					(0.0039)	(0.0039)
Observations	9,108	9,108	9,108	9,108	5,630	5,630
Number of municipalities	1,012	1,012	1,012	1,012	811	811
State trend	Yes	Yes	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors are clustered at the municipality level and reported in parentheses. The dependent variable is the log of the fertility rate [number of children per 1000 women of childbearing age]. Main independent variables are formal sector employment, male and female formal sector employment, manufacturing employment, male and female manufacturing employment, and male and female informal sector employment (all in logs). Regressions are weighted by population of women aged 15-44, and results are for years 2005-2013. Cubic polynomials of the log population of men and women aged 15-44 are included as controls.

Table 4: First stage and OLS, IV, and RF of impact of employment on fertility

	(1) 1st stage	(2) OLS	(3) IV	(4) RF
Log (pred. employment)	0.85***			0.13***
	(0.091)			(0.033)
Log (all employment)		0.031***	0.16***	
		(0.010)	(0.042)	
Observations	9,108	9,108	9,108	9,108
Number of municipalities	1,012	1,012	1,012	1,012
State trend	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Kleibergen-Paap Wald F statistic			87.5	

The first stage for the log of employment is shown in the first column. OLS, instrumented results (using log of employment as IV), and the reduced form are shown in remaining columns, where the dependent variable is the log of the fertility rate. Standard errors are clustered at the municipality level and reported in parentheses. The dependent variable is the log of formal sector employment. Regressions are weighted by population of women aged 15-44, and results are for years 2005-2013. Cubic polynomials of the log population of men and women aged 15-44 are included as controls.

Table 5: First stage regressions for men's and women's employment

	(1)	(2)	(3)	(4)
	First stages			
	Log (male emp)	Log (female emp)	Log (male manuf emp)	Log (female manuf emp)
Log (pred. male employment)	0.45**	-0.13		
	(0.14)	(0.11)		
Log (pred. female employment)	0.30**	1.14***		
	(0.10)	(0.11)		
Log (pred. manuf male employment)			0.54***	-0.11
			(0.15)	(0.18)
Log (pred. manuf female employment)			0.21	1.18***
			(0.13)	(0.16)
Observations	9,108	9,108	9,108	9,108
# of municipalities	1,012	1,012	1,012	1,012
State trend	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

First stage regressions for the log of male and female employment (columns 1 and 2) and the log of male and female manufacturing employment (columns 3 and 4) using the predicted measure of employment described in the text are shown. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 2005-2013. Cubic polynomials of the log population of men and women aged 15-44 are included as controls.

Table 6: Impact of male and female employment on fertility

	(1) OLS	(2) IV	(3) RF	(4) OLS	(5) IV	(6) RF
Log (male employment)	0.0328**	0.32**				
	(0.0139)	(0.13)				
Log (female employment)	-0.00105	-0.087				
	(0.0129)	(0.079)				
Log (predicted male employment)			0.15**			
			(0.061)			
Log (predicted female employment)			-0.0055			
			(0.053)			
Log (male manuf. employment)				0.0189**	0.20**	
				(0.00854)	(0.089)	
Log (female manuf. employment)				-0.00522	-0.060	
				(0.00606)	(0.052)	
Log (pred. manuf. male employment)						0.12**
						(0.047)
Log (pred. manuf. female employment)						-0.027
						(0.043)
Observations	9,108	9,108	9,108	9,108	9,108	9,108
Number of municipalities	1,012	1,012	1,012	1,012	1,012	1,012
State trend	Yes	Yes	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap Wald F statistic		10.0			13.7	

OLS, IV, and reduced form regressions shown. The first three columns include all formal sector employment in logs, while the next set of columns only includes formal sector manufacturing employment in logs. The dependent variable in all cases is the log of the general fertility rate. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 2005-2013. Cubic polynomials of the log population of men and women aged 15-44 are included as controls.

Table 7: Impacts on fertility rates by birth order

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	IV	OLS	IV	OLS	IV
	1st birth	1st birth	2nd birth	2nd birth	3rd+ birth	3rd+ birth
Predicted male employment	0.014	0.059	0.073***	0.62***	0.055*	0.83***
	(0.026)	(0.20)	(0.025)	(0.22)	(0.028)	(0.27)
Predicted female employment	0.027	0.044	-0.021	-0.31**	-0.034	-0.35**
	(0.022)	(0.12)	(0.021)	(0.12)	(0.026)	(0.15)
Observations	9,108	9,108	9,108	9,108	9,105	9,105
Number of municipalities	1,012	1,012	1,012	1,012	1,012	1,012
State trend	Yes	Yes	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
KP Wald F stat		10.0		10.0		10.0

OLS, IV, and reduced form regressions of fertility rates by parity on the log of predicted employment for men and women are shown. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 2005-2013. Cubic polynomials of the log population of men and women aged 15-44 are included as controls.

Table 8: Impact of current and lagged employment on fertility

	(1) GFR	(2) GFR	(3) GFR	(4) GFR
Pred. male employment	0.15**	0.32***	0.32***	0.39***
	(0.061)	(0.098)	(0.10)	(0.11)
Pred. female employment	-0.0055	-0.17	-0.15	-0.20**
	(0.053)	(0.11)	(0.096)	(0.100)
Pred. male employment		-0.14	-0.0017	-0.11
Lag one		(0.093)	(0.16)	(0.17)
Pred. female employment		0.16*	0.21	0.20**
Lag one		(0.094)	(0.14)	(0.10)
Pred. male employment			-0.15	0.074
Lag two			(0.11)	(0.14)
Pred. female employment			-0.0064	-0.055
Lag two			(0.11)	(0.12)
Pred. male employment				-0.11
Lag three				(0.086)
Pred. female employment				0.062
Lag three				(0.098)
Sum of lags for men	0.15**	0.18***	0.16**	0.24**
	(0.061)	(0.067)	(0.072)	(0.098)
Sum of lags for women	-0.0055	-0.0033	0.051	0.09
	(0.053)	(0.057)	(0.061)	(0.068)
Observations	9,108	8,096	7,084	6,072
Number of municipalities	1,012	1,012	1,012	1,012
State trend	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

This table shows results of regressions of the log of the fertility rate on the log of predicted formal sector employment for men and women and their lags. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 2005-2013. Cubic polynomials of the log population of men and women aged 15-44 are included as controls.

Table 9: Decomposition of employment into high and low frequency components

	(1) Log (GFR)
High-frequency men's employment	0.42**
	(0.18)
High-frequency women's employment	-0.29*
	(0.17)
Low-frequency men's employment	0.14***
	(0.058)
Low-frequency women's employment	-0.0033
	(0.057)
Observations	8,096
Number of municipalities	1,012
State trend	Yes
Pop Controls	Yes
Municipality FE	Yes
Year FE	Yes

This table shows results of regressions of the log of the fertility rate on the decomposition of the log of employment for men and women into high and low frequencies. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 2006-2013. Cubic polynomials of the log population of men and women aged 15-44 are included as controls.

Table 10: Relationship between gender intensity of sector and wages

	(1) IMSS Male-to-female wage earnings	(2) IMSS Female earnings	(3) IMSS Male earnings	(4) ENOE Male-to-female wage ratio	(5) ENOE Female wage	(6) ENOE Male wage
Proportion male	-0.42***	0.093	-0.27***	-0.17***	-0.054	-0.20***
	(0.042)	(0.090)	(0.088)	(0.052)	(0.075)	(0.066)
Observations	2,206	2,206	2,206	640	640	640
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Regressions of log wages and log earnings and the wage ratio on proportion of men in each sector, weighted by employment in that sector. Each cell is a sector-year. The ENOE regressions are limited to sectors with a cell size of 10 or greater.

Table 11: Impacts on the wage ratio and wages

	(1) Wage ratio	(2) Wage ratio	(3) Wage ratio	(4) Male wage	(5) Male wage	(6) Male wage	(7) Female wage	(8) Female wage	(9) Female wage
Prop. male	-0.19***	-0.19***	-0.060	0.050	-0.16***	-0.036	0.23***	-0.0018	0.034
	(0.055)	(0.051)	(0.054)	(0.069)	(0.056)	(0.069)	(0.079)	(0.060)	(0.075)
Prop. men	-0.034		0.44***	0.77***		0.44***	0.87***		0.12
in high-skill	(0.088)		(0.13)	(0.070)		(0.097)	(0.100)		(0.13)
Prop. women		-0.26***	-0.57***		0.70***	0.39***		0.98***	0.89***
in high-skill		(0.060)	(0.088)		(0.046)	(0.071)		(0.070)	(0.080)
Observations	631	631	628	631	631	628	631	631	628
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Regressions of either wage inequality, log of male wages, or log of female wages in each sector against the proportion of men in each sector and the proportion of each gender in high-skill occupations. Results are for years 2006-2013.

Table 12: Impacts on earnings

	(1) Male earnings	(2) Female earnings	(3) Male earnings	(4) Female earnings
Male employment	0.15***	0.17***		
	(0.026)	(0.022)		
Female employment	-0.070***	-0.10***		
	(0.017)	(0.017)		
Pred. male employment			-0.015	0.19***
			(0.068)	(0.061)
Pred. female employment			-0.060	-0.21***
			(0.054)	(0.051)
Observations	9,108	9,108	9,108	9,108
Number of municipalities	1,012	1,012	1,012	1,012
State trend	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Type	OLS	OLS	RF	RF

This table shows results of regressions of log earnings on the log of actual and predicted employment (reduced form) for men and women. Results are for years 2005-2013. Standard errors are clustered at the municipality level.

Table 13: First stage regressions for maquiladora employment

	(1) Men	(2) Women	(3) Men	(4) Women
IVmen	0.901***	0.356***	0.855***	0.330***
	(0.0523)	(0.0610)	(0.0536)	(0.0614)
IVwomen	0.212***	0.744***	0.190***	0.709***
	(0.0365)	(0.0596)	(0.0369)	(0.0572)
Observations	38,441	38,441	4,496	4,496
R-squared	0.807	0.764	0.813	0.766
Number of municipalities	2,403	2,403	281	281
Linear state trend	Yes	Yes	Yes	Yes
Sample	Full	Full	Restricted	Restricted
Kleibergen-Paap Wald F statistic	57.06	57.06	51.64	51.64

This table shows results of first stage regressions. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 1991-2006. The dependent variables (men and women) represent the total employment of that gender in the municipality divided by the 1990 population aged 15-44 of that gender. The IV represents replaces total employment with total employment solely due to expansions, contractions, openings, and closings of plants. Regressions contain fixed effects for municipality, year, and linear state trends. Cubic polynomials of the log population of men and women aged 15-44 are included as controls. Restricted sample contains only municipalities that have at least 1% employment of men or women in maquiladoras at least one year.

Table 14: OLS, IV, and reduced form results using maquiladora employment

	(1) OLS Log (GFR)	(2) IV Log (GFR)	(3) RF Log (GFR)	(4) OLS Log (GFR)	(5) IV Log (GFR)	(6) RF Log (GFR)
Men	0.107	0.259*		0.188	0.334**	
	(0.142)	(0.146)		(0.120)	(0.134)	
Women	-0.0319	-0.154		-0.100	-0.215	
	(0.140)	(0.148)		(0.138)	(0.155)	
IV: men			0.178**			0.215***
			(0.0896)			(0.0691)
IV: women			-0.0597			-0.0889
			(0.0871)			(0.0892)
Observations	38,441	38,441	38,441	4,496	4,496	4,496
Number of municipalities	2,403	2,403	2,403	281	281	281
Linear state trend	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap Wald F statistic		57.06			51.64	
Sample	Full	Full	Full	Restricted	Restricted	Restricted

OLS, IV, and RF regressions. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 1991-2006. The independent variables (men and women) represent the total employment of that gender in the municipality divided by the 1990 population aged 15-44 of that gender. Regressions contain fixed effects for municipality, year, and linear state trends. The IV represents replaces total employment with total employment solely due to expansions, contractions, openings, and closings of plants. Cubic polynomials of the log population of men and women aged 15-44 are included as controls. Restricted sample contains only municipalities that have at least 1% employment of men or women in at least one year.

Table 15: Impacts on population

	(1) OLS Men	(2) IV Men	(3) OLS Women	(4) IV Women
Male employment (all)	0.023**	-0.36	0.0073	-0.23
	(0.011)	(0.24)	(0.011)	(0.23)
Female employment (all)	0.034***	-0.065	0.046***	-0.14
	(0.013)	(0.16)	(0.013)	(0.16)
Observations	9,108	9,108	9,108	9,108
Number of municipalities	1,012	1,012	1,012	1,012
State trend	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

OLS and IV regressions for the impact of log male and female employment on log population are shown. Standard errors are clustered at the municipality level.

Table 16: Using a Bartik-style IV to identify the impact of employment on fertility

	(1) OLS	(2) IV	(3) RF	(4) OLS	(5) IV	(6) RF
Male emp.	0.0446***	0.48**				
	(0.0143)	(0.20)				
Female emp.	-0.00963	-0.13				
	(0.0129)	(0.11)				
IV (men)			0.20***			
			(0.069)			
IV (women)			-0.0068			
			(0.056)			
Male manuf. emp.				0.0216**	0.22**	
				(0.00885)	(0.099)	
Female manuf. emp.				-0.00552	-0.074	
				(0.00628)	(0.062)	
IV (men)						0.11**
manuf. only						(0.049)
IV (women)						-0.036
manuf. only						(0.050)
Observations	9,108	6,129	6,129	9,108	9,108	9,108
Number of municipalities	1,012	1,012	1,012	1,012	1,012	1,012
State trend	Yes	Yes	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap Wald F statistic		6.24			10.9	

OLS, IV, and reduced form regressions shown. The first three columns include all formal sector employment, while the next set of columns only includes formal sector manufacturing employment. The dependent variable in all cases is the log of the general fertility rate. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 2005-2013.

Table 17: Impact of male and female employment on fertility

	(1) OLS	(2) IV	(3) RF	(4) OLS	(5) IV	(6) RF
Log (male employment)	0.0323**	0.30**				
	(0.0140)	(0.13)				
Log (female employment)	-0.00470	-0.084				
	(0.0129)	(0.078)				
Log (predicted male employment)			0.14**			
			(0.062)			
Log (predicted female employment)			-0.0100			
			(0.055)			
Log (male manuf. employment)				0.0166**	0.18*	
				(0.00821)	(0.096)	
Log (female manuf. employment)				-0.00667	-0.052	
				(0.00594)	(0.051)	
Log (pred. manuf. male employment)						0.094**
						(0.047)
Log (pred. manuf. female employment)						-0.026
						(0.043)
Observations	9,108	9,108	9,108	9,108	9,108	9,108
Number of municipalities	1,012	1,012	1,012	1,012	1,012	1,012
State-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Pop Controls	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap Wald F statistic		9.52			10.3	

OLS, IV, and reduced form regressions for the impact of male and female employment on general fertility rates are shown. Standard errors are clustered at the municipality level and reported in parentheses. Regressions are weighted by population of women aged 15-44, and results are for years 2005-2013. Regressions contain fixed effects for municipality, year, and year by state interactions. Cubic polynomials of the log population of men and women aged 15-44 are included as controls. Cubic polynomials of the log population of men and women aged 15-44 are included as controls.