

Why does the sex ratio at birth rise? Evidence from Vietnam*

Nghiem Q. Huynh[†] Ngoc Nguyen[‡]

December 30, 2023

Abstract

This study connects Vietnam’s elevated sex ratio at birth (SRB) to the 2001 US-Vietnam Bilateral Trade Agreement. Our model incorporates three major factors that influence SRB: income, relative returns based on the child’s sex, and fertility. The model presents twelve predictions, which are tested using large-scale repeated cross-sectional and panel surveys in a difference-in-difference design. The results indicate that mothers who experience larger tariff reductions tend to have a stronger preference for sons, work more, and desire fewer children. These findings suggest that fertility is the main driver of the elevated SRB. Overall, this paper highlights the interplay between cultural norms, maternal income, childcare, and fertility, revealing the unexpected demographic impact of trade policies.

1 Introduction

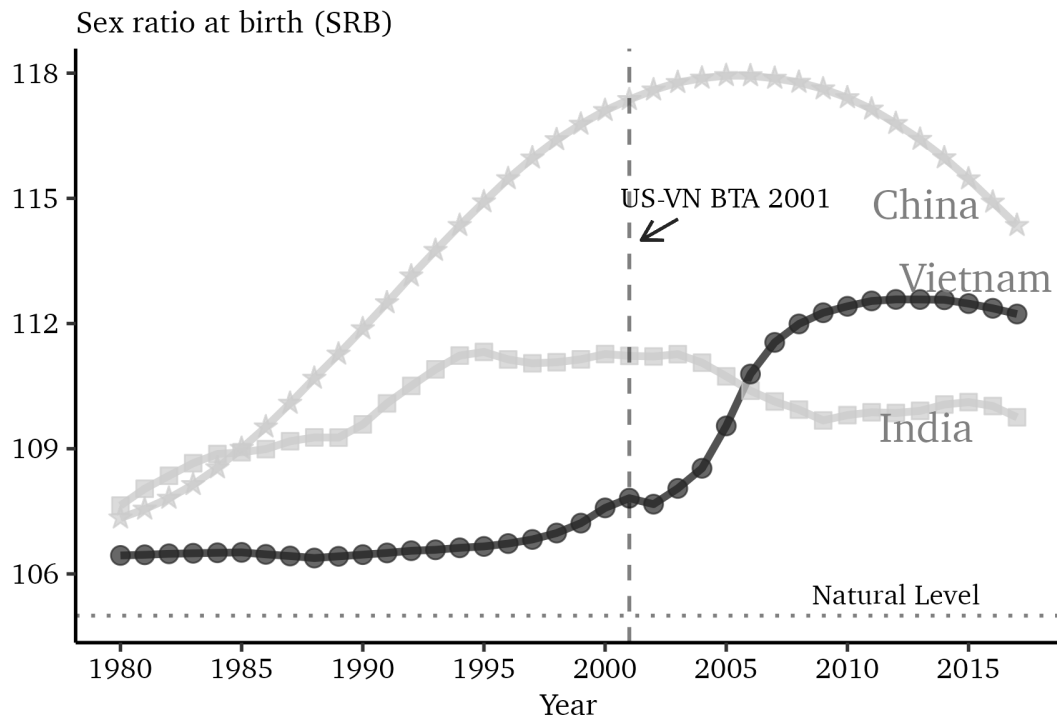
Sex imbalance at birth is a concerning demographic phenomenon in many Asian and Eastern European countries. The sex ratio at birth (SRB), which measures the number of boys per 100 girls at birth, typically stabilizes around 105. However, countries like China,

*We are grateful to S. Anukriti, Tim Guinane, Arthur Lewbel, Kaivan Munshi, Rohini Pande, and Mark Rosenzweig for helpful suggestions. We also thank seminar participants at Yale Development Lunch and Boston College Dissertation Workshop for their comments. We are grateful to Brian McCaig for guidance with VHLSS data.

[†]Yale University, nghiem.huynh@yale.edu

[‡]Lightcast, jade.nguyen@lightcast.io

Figure 1: Sex ratio at birth by country, 1980-2017



Source: Chao et al. (2019). Notes: Sex ratio at birth (SRB) reports the number of boys per 100 girls at birth.

India, and Vietnam, where son preference is prevalent, experience much higher SRBs than this natural level. As shown in Figure 1, the SRB in these countries has exceeded 110, foreshadowing a serious shortage of women in the future. Of particular interest is Vietnam, where the SRB began to rise in the 2000s, much later than in China and India. This period coincides with Vietnam's integration into the world economy, including the signing of a major trade agreement with the US in 2001. Figure 1 illustrates that the SRB started accelerating in Vietnam following the 2001 US-Vietnam Bilateral Trade Agreement (BTA) enactment.

This paper connects these seemingly unrelated events and tests competing economic theories to explain the rise in SRB. In the Vietnamese cultural context, societal norms place significant emphasis on male roles in carrying forward the family lineage, known as "*nổi đời tông đường*". Consequently, couples aspire to have at least one son.

Before the 2001 trade agreement, the average income of Vietnamese individuals was extremely low, approximately 8% of that of Americans in 2000. Subsequently, the trade shock boosted the demand for Vietnamese exports from the US, leading to higher house-

hold incomes. This income surge could enable couples to afford sex-selective technology to fulfill their preference for having a son. This income effect represents the first channel through which the trade shock can influence the SRB, aligning with the findings of Almond et al. (2019) in the context of rising land prices in China. Notably, the household's increase in demand for sons remains unchanged regardless of whether males or females earn income.

The trade shock, however, can also favor industries predominantly dominated by one sex, thereby influencing the relative returns of boys and girls. Consequently, the second channel through which the trade shock can impact the rise in SRB is reducing relative returns to girls, leading parents to favor boys. This pattern has been observed, for instance, in the context of changing tea prices in China, as evidenced by Qian (2008).

Nonetheless, higher returns to female-dominated industries need not increase the likelihood of having daughters. This is particularly true when the mother assumes the primary caregiving role for her children, leading to a trade-off between work and childcare. Reducing industry-specific export tariffs can raise mothers' income and simultaneously increase the opportunity cost of childcare. If the effect of substituting work for childcare outweighs the income effect, mothers who already have a preference for having a son will choose to have fewer children and prioritize the birth of a son over a daughter. Conversely, increasing the father's wage only leads to an income effect that increases the overall demand for children, as his role in childrearing is minimal. We refer to this mechanism as the fertility channel, which influences SRB.

In Section 2, we adapt an off-the-shelf quantity-quality model to our specific context, taking into account the interplay between economic forces and cultural factors. This adaptation enables us to consider the impact of trade shocks on the sex ratio at birth (SRB) through three distinct channels: income, relative returns, and fertility. The model focuses on a couple, with the mother being the primary caregiver, who desires to have at least one son and has the ability to decide both the number and sex of their children. From this model, we derive seven testable predictions pertaining to child sex, labor supply, and the overall number of children.

To test these predictions, we leverage household data, which is described in Section 3. This data includes both repeated cross-sectional surveys and panels that provide direct observations of the parents' industries at the time of their children's birth. By linking the parents' industries to tariff changes, we employ a difference-in-difference design in Section 4 to compare outcomes for parents who experienced different tariff shocks before and after the trade agreement was enacted.

In addition to comparing parents, we construct two shift-share or Bartik measures to

capture the local effects of the trade shocks on non-labor income and the relative returns of female jobs compared to male jobs. These measures account for indirect income to households, in addition to the direct effects of the trade shocks on parents, as well as the relative prices associated with the sex of the child. For instance, even parents working in industries minimally exposed to the trade shock can still exhibit sex-selective behavior if they experience higher income due to general equilibrium effects at the local level or observe changes in the returns associated with having male children resulting from the trade shocks.

The data consistently support our model predictions, and among the three mechanisms analyzed (income, relative returns, and fertility), the strongest evidence favors the fertility channel. Notably, we observe that mothers of infants, on average, experience a ten percentage point tariff reduction, leading to a 4% increase in male infants and a 9% decrease in their yearly likelihood of giving birth. Furthermore, mothers exposed to greater tariff reductions increase their work hours, with a potential 10% rise observed in the garment and textile manufacturing sectors, which are particularly impacted by tariff reductions.

Related literature: Our findings highlighting the quality-quantity tradeoff contribute to the large literature on the factors influencing the skewed sex ratio at birth. While this literature traces back to Becker and Lewis (1973) and the phenomenon of missing women popularized by Sen (1990), the evidence regarding the primary cause of the rise in the sex ratio at birth remains mixed, as discussed in the review by Dasgupta and Sharma (2022).

Previous studies have provided evidence supporting the relevance of the fertility mechanism, which aligns with our findings. Notable examples include Jayachandran (2017) and studies on China's One-Child policy such as Ebenstein (2010) and Li et al. (2011). However, recent evidence challenges this explanation.

The relative returns perspective, as argued by Qian (2008), suggests that improved female labor market returns should enhance the survival of girls. Similarly, Rosenzweig and Schultz (1982) and Carranza (2014) posit that increased returns to women's work, which elevate the opportunity costs of having sons, reduce the preference for male children.¹ Furthermore, Almond et al. (2019) argues that land value plays a more significant role in sex selection than fertility pressure, even within the context of China's One-Child policy. They argue that land reform raises household income, enabling families to afford sex-selective abortions. By encompassing all three channels in our model and empirical design, we are able to compare and contrast their effects.

¹The returns to women's work in these papers correspond to the returns to daughters in our model.

We also contribute to the literature examining the effects of trade policy, aimed at promoting economic development, on gender discrimination. In the context of US-VN BTA, existing studies have shown that this trade shock increased wages, formalization, and reduced poverty (Fukase, 2013; McCaig, 2011; McCaig & Pavcnik, 2018). However, we show that this policy can exacerbate the sex ratio at birth when son preference and time constraints for women are present.

In addition, while Anukriti and Kumler (2019) and Chakraborty (2015) have provided evidence of the impact of trade liberalization on the sex ratio at birth in India using population censuses, which has a substantial time gap between them, our study capitalizes on more frequent and extensive household surveys. This advantage in data allows us to better observe the industries and locations of parents at the time of their children's births.

Our study complements existing literature on the rise in the sex ratio at birth in Vietnam. While Bongaarts and Guilmoto (2015) and Guilmoto et al. (2009) attribute the increase to the adoption of ultrasound technology and sex-selective abortion, Guilmoto (2012) and Becquet and Guilmoto (2018) suggest that declining fertility rates and changes in son preference intensity play a role. In our analysis, we provide direct evidence supporting our model incorporating the fertility channel, and we find it challenging for alternative hypotheses to account for all the empirical findings.

2 A quantity-quality trade-off model

We present a quality-quantity (Q-Q) model that associates the sex of the child with quality. The child's sex is determined by the wages of both parents, the household's non-labor income, and the relative returns of female-to-male jobs. In the Vietnamese context, we make two assumptions. First, households desire to have at least one son to continue the family lineage. Second, mothers bear the sole responsibility for child-rearing while also participating in the labor market.²

We incorporate these assumptions into a standard Q-Q model inspired by Jones et al. (2010), wherein mothers assume the primary caregiver role. Specifically, let's consider a household comprising a woman of reproductive age and her husband. Referring to them as a mom (m) and dad (d), they jointly determine their private consumption c_g , leisure ℓ_g ,

²A time-use study conducted by ActionAid-Vietnam (2016) reveals that women spend approximately 5.24 hours per day on unpaid care work, whereas men dedicate 3.17 hours per day to similar activities. This gender disparity in household chores widens as the number of children increases. For instance, when there are two children, wives spend 7.41 hours on unpaid care work compared to husbands who spend 4.34 hours. Moreover, the female employment rate in Vietnam is considerably high, reaching approximately 77% among prime-aged women (20-64 years old) in VHLSS 2002-2016.

where $g \in m, d$, the number of children n , and the quality of their children q , with $q = 1$ indicating a son and 0 otherwise.³

Alongside interpreting quality as the child's sex, we also establish the household's desire to have at least one son, which we denote as $Q \equiv qn$, representing the effective number of sons. In the context of logarithmic utility assumption, the model implies that a household would prefer having a son ($Q = 1$) over, for instance, having three daughters ($n = 3$) with no sons ($q = 0$). The utility function specific to each parent g is defined as follows:

$$U_g = \alpha_c \log(c_g) + \alpha_\ell \log(\ell_g) + \alpha_n \log(n) + \alpha_q \log(Q),$$

where $g \in m, d$, and α_i represents the weight of preference assigned to the corresponding component, including consumption, leisure, fertility, and the effective number of sons.

While both parents have the same amount of time available to allocate between work and leisure, it is only the women who take care of the children, incurring a time cost of γ per child. Additionally, the household incurs a cost of p_q for the effective quality of the children Q . This cost encompasses the overall expenses associated with sex-selective abortion, including economic, physical, and psychological factors, as well as the benefits it brings to the couple. An example of p_q could be the relative returns of daughters compared to sons, which parents may forecast based on the current relative returns of women's work in the labor market.

In addition to labor incomes, the household also receives non-labor incomes denoted as I . This includes transfers from other members of the household or the value of their land, as discussed in Almond et al. (2019). Furthermore, children and their quality are considered public goods within the household. The optimization problem for the household can be formulated as follows:

$$\begin{aligned} \max_{\{c_m, c_d, n, q\}} \quad & \lambda_d U_d + \lambda_m U_m \\ \text{s.t.} \quad & c_m + c_d + p_q Q \leq I + w_d(1 - \ell_d) + w_m(1 - \ell_m - \gamma n). \end{aligned}$$

Here, λ_m represents the bargaining weight of the mother, and λ_d represents that of the father. We assume that $\lambda_m + \lambda_d = 1$.

By defining the household's total income as $W = I + w_m + w_d$, we can derive the

³We actually let q be a continuous variable in the unit interval $[0, 1]$. All variables n and Q here are continuous to allow simple comparative statics.

following solutions after taking logarithms:

$$\log q = \log(\gamma\alpha_q/\alpha_n) - \log p_q + \log w_m, \quad (1)$$

$$\log \ell_m = \log(\lambda_m\alpha_\ell) + \log W - \log w_m, \quad (2)$$

$$\log n = \log(\alpha_n/\gamma) + \log W - \log w_m. \quad (3)$$

We have derived twelve predictions based on simple comparative statics, which are summarized in Panel A of Table 1. Each column corresponds to the numbered equation mentioned above, focusing on three outcomes: the sex of the children, the mother's labor supply, and the number of children or the probability of giving birth. These outcomes are influenced by four key factors: the mother's wage (w_m), the father's wage (w_d), the household's non-labor income (I), and the relative returns to daughters or the opportunity costs associated with having a son (p_q).

Panel A of Table 1 presents statements in the form "if [row] increases, then [column] [cell], *ceteris paribus*." For instance, the top-left cell states that if the mother's wage (w_m) increases, the probability of the newborn child being a son also increases, *ceteris paribus*. It is important to note that each prediction in this context requires the condition of *ceteris paribus*, meaning that other variables, in this example, the father's wage, household income, and relative returns (as indicated in column (1)), are held constant.

Before testing all of these predictions using the BTA shock and household data, it is essential to highlight the distinctive features of our model compared to existing ones. Specifically, our model incorporates non-labor income (I) and the father's wage as potential drivers of SRB. However, unlike the model in Almond et al. (2019), neither of these factors affects the demand for sons; they only contribute to an increase in the overall number of children. Higher non-labor income or the father's wage motivates households to desire more children in general, including sons, thereby nullifying any specific effect on the demand for sons.

On the other hand, in contrast to Qian (2008), our model predicts that an increase in the mother's income, while keeping the spouse's wage and non-labor income constant, leads to a higher demand for sons rather than daughters. This prediction stems from the trade-off mothers face between child care and work. With a higher income, they exhibit a greater inclination to work more, have fewer children, and prioritize the birth of sons. Our model incorporates both the relative prices of sex (captured by p_q) and the relative wages of parents. In practice, it is challenging to determine whether parents compare their own relative wages or observe the relative returns in the local market to forecast the future relative returns of daughters versus those of sons. Hence, we account for both

Table 1: Model's Predictions and Estimands

Panel A: Model Predictions			
<i>Outcomes:</i>	Prob(son), q	Mom's labor supply, $(1 - \ell_m)$	Fertility, n
	(1)	(2)	(3)
Mom's wage $w_m \uparrow$	+	+	-
Dad's wage $w_d \uparrow$	\sim	-	+
Non-labor income $I \uparrow$	\sim	-	+
Opp. cost of sons $p_q \uparrow$	-	+	\sim
Panel B: Hypotheses for coefficients in (4) and (5)			
β_M	+	+	-
β_D	\sim	-	+
β_I	\sim	-	+
β_F	-	+	\sim

Notes: Each cell in Panel A represents a prediction of the model, "If [row] increases, then [column] [cell]". In the cells, the "+" symbol means increases, "-" decreases, and " \sim " does not change. All predictions assume ceteris paribus. Panel B lists predicted coefficients in regression (4) and (5) that correspond to the model.

aspects—the relative prices at the local level and the household level—in our empirical design.

In summary, our model encompasses three key mechanisms through which the trade shock can influence SRB. Notably, our model generates predictions that diverge from existing models. To further investigate these predictions, we proceed to assemble the necessary data and empirical components that align with our theoretical framework.

3 The trade shock and data

3.1 The 2001 US-Vietnam BTA

The identification strategy in this paper relies on leveraging the variations in tariff exposure resulting from the implementation of the US-VN BTA on December 10, 2001. While

the agreement had minimal impact on Vietnam's import tariffs from the US, it significantly lowered tariff rates for Vietnamese exports to the US market. Specifically, the US reclassified Vietnam's trade status from Column 2, which was reserved for former communist countries, to the Most Favored Nation (MFN) tariff schedule, leading to an immediate reduction in import tariffs for Vietnamese products in all traded industries (Figure A1). While agriculture and mining (codes 1-14) saw slight reductions, manufacturing (codes 15-36) experienced substantial cuts, averaging 30.2 percentage points. Notably, sectors like tobacco, textiles, and garments (codes 16, 17, 18) saw the greatest reductions, ranging from 48.5 to 58.2 percentage points. Vietnam's exports to the US skyrocketed, growing from approximately 0.7 billion US dollars in 2000 to nearly 10 billion US dollars in 2008.

Unlike the US's commitment to reduce tariffs, Vietnam's commitment to the bilateral agreement centered on regulatory and legal reforms. While Vietnam was obligated to reduce tariffs mainly on agricultural and food products, these reductions were quite minimal, ranging from 0.03 to 2.7 percentage points. In addition to tariffs, Vietnam agreed to eliminate various import quotas as per the BTA. Nonetheless, the majority of these quotas had already been eliminated by the end of 2002 (STAR-Vietnam, 2003). These factors underscore that the predominant influence on the Vietnamese economy within the US-VN BTA stems from tariff adjustments on the US side.

Prior research by McCaig (2011) and McCaig and Pavcnik (2018) highlights the significant structural transformation induced by the US-VN BTA in Vietnam's economy. This change has led to reduced poverty rates and increased formalization in the labor market. These studies also present compelling evidence of the exogenous nature of the BTA tariff reductions initiated by the US. Notably, these tariff adjustments exhibit no correlation with preexisting economic conditions in Vietnam, which could potentially be linked to preferences for sons and fertility. Their work provides a compelling basis to view the US-VN BTA as a natural experiment for examining its impact on fertility behaviors among Vietnamese parents.

The swift BTA implementation serves as empirical parallels to the exogenous parameters in our model in section 2 that could impact fertility and sex selection behaviors. These parameters (w_m , w_d , I , p_q) represent maternal wage, paternal wage, non-labor income, and relative returns to daughters. The significant tariff reductions create a positive demand shock for Vietnamese workers. Consequently, wage rates in industries with higher tariff cuts rise in comparison to those with lower tariff cuts. Thus, we employ the variation in industry-specific tariff reduction magnitude as a proxy for the wage rate change experienced by mothers and fathers (w_m and w_d) within their respective industries.

Besides potentially altering wage rates, the BTA can directly impact non-labor income (I) by raising local average income levels. In Vietnam, it's common for extended family members to share a household, allowing adult couples to receive transfers. Consequently, couples in more BTA-exposed provinces might experience a greater non-labor income effect than those in less exposed provinces. Another avenue through which the BTA could affect non-labor income is through increased returns on land holdings, which Almond et al. (2019) linked to skewed sex ratios in China (1978-1986). To capture potential shifts in non-labor income due to the BTA, we create a province-level shift-share measure.

Moreover, as the BTA's largest tariff reductions affect industries predominantly employing women, they could elevate the female relative wage (p_q in the model). This trend could lead parents to perceive an increased relative value of girls, potentially influencing their decisions to have more daughters. We develop a female-specific BTA exposure at the province level using a similar shift-share approach, serving as a proxy for the shift in female relative returns prompted by the trade agreement. The detailed construction of these shift-share exposure variables is discussed in [subsection 3.2](#).

3.2 Data and descriptive statistics

In this subsection, we outline the data and variables used in our empirical analysis, drawing from two primary sources: (i) household data extracted from the Vietnam Household Living Standards Survey (VHLSS), and (ii) tariff exposure data from McCaig and Pavcnik (2018), combined with information from the Vietnamese Census.

The VHLSS data stands as a fitting resource for our study objectives, as it captures parents' information at the moment of their employment, eliminating the need for extrapolation over extended timeframes as seen quite commonly in the literature which often relies on population census data. Furthermore, we discuss how we can represent province-wide exposure to the trade shock, which serves as a proxy for changes in non-labor income and relative returns to daughters, through shift-share (Bartik) variables.

Household data

Our primary data source is the Vietnam Household Living Standards Surveys (VHLSS), conducted biennially by the General Statistics Office (GSO) with technical support from the World Bank. This nationally representative survey employs a 12-month recall period.

From the VHLSS, we extract two main sets of variables: employment and family interrelationship. The employment module furnishes industry codes, allowing us to create dependent variables that quantify each parent's exposure to the trade shock based on

their respective industries. This module also provides information about parents' work hours.

We gather outcome variables related to fertility and sex selection, encompassing whether a woman gave birth recently and the gender of her infant. Ensuring accurate identification of individuals' children is crucial for drawing valid inferences about fertility and sex selection behaviors. The task is intricate in the Vietnamese context due to prevalent extended family cohabitation. The default household relation classifier in VHLSS often struggles to distinguish between an individual's children and their nieces or nephews. To enhance family interrelationship accuracy, we develop an algorithm that identifies spouses, fathers, and mothers for each individual. In Appendix A, we detail the algorithm's construction and its performance in comparison to true parental locators available only for the 2014 and 2016 waves. The algorithm achieves a success rate of 95.27% for mothers, 90.99% for fathers, and 88.95% when both parents are considered. Considering stepmothers and stepfathers, the correction rate for both biological parents is slightly lower. These results affirm the algorithm's efficacy in identifying family interrelationships.

VHLSS proves highly suitable yet under-exploited for studying fertility and sex selection behaviors. Our primary findings use eight repeated cross-sectional datasets spanning 2002 to 2016, capturing effects post the enactment of the US-VN BTA in December 2001. These frequent and representative cross-sections facilitate timely inferences regarding the influence of parents' current jobs on their fertility choices. Unlike prevalent approaches in the literature that rely on infrequent population censuses under the assumption of job and migration stability, VHLSS overcomes this constraint, while offering a sufficiently large sample size for our tasks.

Furthermore, exploiting VHLSS's rotating panel feature, we construct a three-wave women's panel (2002-2004-2006) to complement our cross-sectional findings. This panel empowers us to control for unobservable factors like fertility and sex ratio preferences. Additionally, it enables differentiation between the direct impact of *within* industry exposure and compositional effects arising from workers shifting across industries.

Although the first available data point is from 2002, the 12-month recall characteristic renders it suitable for the pre-BTA timeframe. Specifically, despite the BTA taking effect in December 2001, the employment information from the 2002 wave largely pertains to 2001 due to the 12-month recall period. Furthermore, we focus on a sample of infants aged 0-1 year old. For children selected from the 2002 wave, the BTA likely began during their late gestation phase, meaning that outcomes for these children reflect parental fertility decisions prior to the BTA. Hence, data from the 2002 wave constitute

our pre-BTA observations. Additionally, we define the pre/post period using birth years, not surveyed years, to augment pre-group observations, encompassing children born in 1999, 2000, 2001, and 2002. In total, we acquire a reasonably large sample of around 5,000 infants from each wave.⁴

Exposure measures to the US-VN BTA We use two data sources to create various indicators of BTA exposure. First, we leverage industry-specific tariff information concerning Vietnamese imports to the US, categorized at the 2-digit ISIC2 industry level as aggregated by McCaig and Pavcnik (2018). The authors sourced the primary tariff schedules from the US International Trade Commission. We establish a connection between industry-specific tariff reductions and the industry code of each parent to capture their direct personal exposure to the BTA.

We additionally construct shift-share measures at the province level to serve as proxies for indirect exposure to the trade agreement's impact on (i) non-labor income and (ii) relative returns to daughters. For the former, we use an overall provincial exposure, where the "shifts" refer to industry-specific tariff reductions and the "shares" pertain to pre-BTA industry-specific employment shares at the province level. The "share" component originates from the Population and Housing Census of 1999. For each province, denoted as p , the province-specific exposure to the BTA is given by:

$$\tau_p \equiv \sum_j \frac{L_{pj,1999}}{L_{p,1999}} \Delta\tau_j$$

where $\Delta\tau_j$ is the change in tariff imposed by the US on industry j , $L_{p,1999}$ is the labor force of province p based on Census 1999, and $L_{pj,1999}$ is the number of workers in industry j living in province p . Since McCaig (2011) already computed this variable, we extract it from the author's replication dataset.

To capture the gender-specific relative returns of the child, we construct a gender-specific shift-share variable inspired by Autor et al. (2019). Specifically, the female exposure for province p is defined as:

$$\tau_p^F \equiv \sum_j \frac{W_{j,1999}}{L_{j,1999}} \frac{L_{pj,1999}}{L_{p,1999}} \Delta\tau_j$$

where the only difference between τ_p^F and τ_p is the female share in industry j . Again, using Census 1999, we compute this female share using the industry j 's prime-aged work-

⁴For 2002 only, the sample size is 8,234 infants due to the larger sample in this initial round.

force ($L_{j,1999}$) and the number of female workers in the same industries ($W_{j,1999}$). Unlike Autor et al. (2019), we define this female share, not at the industry-province level, but at the industry level to avoid the selection problem of female-intensive industries located in female-biased provinces.

Descriptive statistics

We start by providing descriptive statistics for sex ratio trends during our study period. [Table A4](#) presents the summary statistics on the repeated cross-sectional sample of infants aged 0 to 1 year at the time of the survey. Slightly under half of the infants are girls. Although this unconditional fraction appears relatively stable over time, we present in [Figure A2](#) that, upon considering survey-year fixed effects, the proportion of girls among newborns decreased during the 2000s. This observation affirms that our data aligns with the SRB trend documented in the literature using alternative data sources (Becquet & Guilmoto, 2018; Chao et al., 2019; Guilmoto et al., 2009). By 2010, the probability that a newborn is a girl had declined by 5 percentage points. This number translates to an SRB increase from 105 to 115, consistent with the pattern shown in [Figure 1](#).

At the initiation of the BTA, mothers and fathers experienced average tariff reductions of 8.8 and 7.8 percentage points through their respective industries. Over time, mothers' tariff cuts intensified by an additional three percentage points, while fathers' cuts deepened by slightly less than 1 percentage point. This indicates that women gained more exposure to export-oriented sectors compared to men over the study duration. Concurrently, mothers' work hours consistently increased, suggesting a positive correlation between their workforce commitment and heightened exposure to significant tariff reductions.

At the province level, children within a specific province encountered an average tariff reduction of 9 percentage points, coupled with a "female-specific" reduction of approximately 4.2 percentage points.⁵ The female-specific provincial tariff cut is slightly less than half of the gender-neutral counterpart, primarily because industries that later received lower tariff cuts, such as sales, hospitality, and farming, employed more women prior to the BTA.

[Table A5](#) reports the summary statistics for the data used in the panel analysis. This panel comprises women aged 20 to 40 who are tracked across three consecutive rounds: 2002, 2004, and 2006. To explore shifts in their sex selection behavior over time, we select

⁵These provincial-level exposures remained constant throughout the study timeframe, as they are inherently time-invariant. Minor year-to-year fluctuations in these values merely reflect changes in the sample size of children aged 0-1 year old.

women who gave birth within one year before the BTA and also experienced at least one birth after its enactment. Consequently, the sample consists of women who had at least one child upon entering the panel, making them slightly older (early 30s) than mothers of newborns in the cross-sectional sample. Over time, as we track these women, their likelihood of giving birth increases as they progress along their fertility trajectory. Other statistics, encompassing women’s personal exposure to the BTA via their industries, their husbands’ exposure, and province-level exposure, yield analogous values as observed in [Table A4](#).

4 Empirical strategy

We exploit the trade shock’s time, industry, and spatial variations to test the model’s predictions in [Table 1](#). Furthermore, we leverage cross-sectional and panel data to alleviate concerns associated with each approach.

Motivated by the comparative statics in [Section 2](#), we specify the following difference-in-difference (DID) regression and estimate it by OLS:

$$y_{imdpt} = \beta_M \cdot \tau_{j(m)} \cdot \text{Post}_t + \beta_D \cdot \tau_{j(d)} \cdot \text{Post}_t + \beta_I \cdot \tau_p \cdot \text{Post}_t + \beta_R \cdot \tau_p^F \cdot \text{Post}_t + \theta_t + \xi_p + \psi_{j(m)} + \phi_{j(d)} + X'_{imdpt} \cdot \lambda + \varepsilon_{imdpt}, \quad (4)$$

where each variable y_{imdpt} is an outcome for child i who is less than two years old at survey time and born to mom m and dad d in province p and in year t . Using cross-sectional data, we observe two outcomes including whether the child is male, denoted by the dummy $\text{Male}_{imdpt} = 1$ if he is, and mom m ’s reported working hours in the main job.

We focus on four coefficients β_M , β_D , β_I , and β_R . Variations in tariff reduction across mom’s industry $j(m)$ and the time of birth relative to the enactment of the trade shock allow us to identify the first coefficient β_M . Consider, for instance, outcome Male_{imdpt} . The interaction between the time-invariant tariff change in moms’ industry $\tau_{j(m)}$ and the indicator Post_t , which is true if child i ’s birth year is after 2002, captures the changes in the sex selection behaviors of mothers working in industries with higher tariff reduction after compared to before the trade policy takes place.

Similarly, identification of β_D follows by replacing $\tau_{j(m)}$ with tariff reductions in dads’ industries $\tau_{j(d)}$. Furthermore, the interaction between the province exposure to the trade shock τ_p and the time dummy identifies β_I , which examines whether households in more exposed provinces adjust their behaviors after the trade shock. Finally, estimates of β_R let us test if the relative returns to daughters, captured at the province level by τ_p^F , impact

sex selection and moms' labor supply.

We include the year fixed effects θ_t (which can be birth year or survey year or both depending on outcomes) to capture any aggregate shocks that equally impact all households in year t , province fixed effects ζ_p to capture any time-invariant differences across provinces, moms' industries fixed effects $\psi_{j(m)}$ any time-invariant differences across moms' industries, and dad's industry fixed effects $\phi_{j(d)}$ any time-invariant differences across dads' industries. The time-varying controls X_{imdp_t} include moms' and dads' birth age, their education, household size, urban dummy, and minority dummy.

In addition to repeated cross-sectional data, we can modify specification 4 to estimate a similar set of coefficients using panel data. Our panel sample consists of women between 20 and 40 years old who have given birth at least once before and after 2002. For each woman i in this sample, we can observe the tariff change in her industry τ_i and her husband's industry τ_h , as well as whether her last-born child is a boy and her labor supply. Thus, the DID specification becomes

$$y_{iht} = \beta_M \cdot \tau_{j(i)} \cdot \text{Post}_t + \beta_D \cdot \tau_{j(h)} \cdot \text{Post}_t + \beta_I \cdot \tau_p \cdot \text{Post}_t + \beta_R \cdot \tau_p^F \cdot \text{Post}_t + \theta_t + \psi_i + X'_{iht} \cdot \lambda + \varepsilon_{iht}, \quad (5)$$

where y_{iht} is an outcome in survey year t of woman i who has children with husband h . The dummy $\text{Post}_t = 1$ if the survey year is larger than 2002. We include individual fixed effects ψ_i and the same set of controls in X_{iht} as X_{imdp_t} but replace a mom and a dad's variables with the individual and her husband's ones. When investigating the effects on whether the last-born child is male, we also include the last-born child's birth year fixed effects.

All regressions are weighted by sampling weights. Standard errors are clustered at the province level. Our hypotheses for outcomes are summarized in Panel B of Table 1.

Fertility outcomes To investigate the effects on fertility, we follow two complementary approaches. First, we construct a retrospective panel from the cross-sectional data in which an observation is at the birth-mom-year level. For each woman, we construct an observation of whether she gives birth to a child at the women-year level, from the year she turns 20 years old to the year she is observed in our data. The outcome of interest is whether there is a birth i to woman w and dad d . The specification is similar to (4), except for the inclusion of mom's fixed effects instead of the province and industries' fixed effects. This approach helps us observe her fertility path (up to the year of the survey) and control for the mother's fertility and sex selection preferences when including

mothers' fixed effects in the specification. However, since we can only observe both parents' employment data at the time of the survey, we must assume that their employment data, which we can only observe once in the survey year, are unchanged across time. Although this assumption may fail to reflect their true employment at the time of birth, it remains informative for families that rarely changed their industries.

Our second approach relies on the *true* panel data as discussed in [section 3](#). We follow women between 20 and 40 years old across three waves 2002, 2004, and 2006, and examine the effects of trade shocks on whether they have a new birth within the year of the survey according to specification [5](#). This approach enables us to test, for instance, whether a higher tariff reduction in a woman's industry reduces the probability that the mother gives birth after versus before the implementation of the policy. Due to the short panel, we can only observe the short-run impact of the policy.

5 Results

5.1 Repeated cross-sections

To begin, we present findings estimated with repeated cross-sectional data regarding the sex ratio at birth among newborn children and the labor supply of their mothers. Subsequently, we delve into the fertility outcomes by examining data from a retrospective panel constructed using cross-sectional data.

Sex ratio at birth

[Table 2](#) presents the outcomes of the regression outlined in [Equation 4](#) across the initial three columns. In Column (1), we include the four exposure measures: mom's, dad's, provincial and female-specific provincial exposures to the BTA, alongside their interaction with the $Post_t$ dummy. Column (2) introduces controls for differences in socioeconomic backgrounds among the children. The rows in the table report the coefficients of interest, namely β_I , β_R , β_M , and β_D , which encapsulate the various channels through which the BTA could affect the gender of newborn children, as suggested by the predictions in Column (1) of [Table 1](#).

As per the predictions presented in Column (1) of [Table 1](#), an increase in a mother's wage due to the BTA tariff reduction in her industry is anticipated to raise the likelihood of a male birth. This prediction is substantiated by the positive and statistically significant coefficient on the interaction between a mother's direct exposure through her

Table 2: Effects on whether a new birth is male and labor supply

	1(Last born is male)		Hrs Worked	1(Recent Birth)		
	(1)	(2)	(3)	(4)	(5)	(6)
Mom Exp x Post	0.0019** (0.0008)	0.0019** (0.0008)	0.2727*** (0.1021)	-0.0001*** (0.0000)	-0.0001** (0.0000)	-0.0001** (0.0000)
Dad Exp x Post	-0.0003 (0.0008)	-0.0003 (0.0009)	0.1547 (0.1197)	0.0002*** (0.0000)	0.0002*** (0.0000)	0.0002*** (0.0000)
Prov Exp x Post	0.0025 (0.0151)	0.0029 (0.0150)	-14.8158*** (2.2548)	0.0018 (0.0015)	0.0013 (0.0014)	0.0042*** (0.0011)
Female Exp x Post	-0.0038 (0.0245)	-0.0044 (0.0244)	28.0347*** (4.9418)	-0.0016 (0.0023)	-0.0002 (0.0021)	-0.0006 (0.0020)
R ²	0.008	0.009	0.333	0.286	0.290	0.290
Observations	30,771	30,770	28,303	2,038,111	2,038,100	2,038,100
Control Group Mean	0.4898	0.4898	134.1040	0.0186	0.0186	0.0186
Controls		✓	✓		✓	✓
Province FE	✓	✓	✓			
Mom Industry FE	✓	✓	✓			
Dad Industry FE	✓	✓	✓			
Survey Year FE	✓	✓	✓	✓	✓	✓
Birth Year FE	✓	✓		✓	✓	✓
Mom FE				✓	✓	✓

Notes: Columns (1)-(2) use a sample of children between 0 and 1 year old; (1) no control, (2) add household controls which include household size, mom's birth age, mom's education, dad's birth age, dad's education, urban and minority dummies. Column (3) is restricted to mothers of the children included in the regressions in columns (1)-(2). The outcome variable is the number of working hours per month of the first or main job of these mothers. The controls include the woman's age, her husband's age, her education, her husband's education, urban, and minority. Columns (4)-(6) use the retrospective panel of women, as explained in the main text. The outcome variable is a dummy for whether the woman has given birth in the past two years. The controls include the mom's birth age, number of older brothers, and number of older sisters. All regressions include survey weights and are estimated by OLS. Standard errors, clustered at the province level, are reported in parentheses.

industry and the $Post_t$ dummy, denoted as β_M . Across the two specifications in Columns (1)-(2), β_M emerges as the only statistically significant effect, and its magnitude remains consistent. Specifically, each one-percentage-point reduction in tariff experienced by the mother through her industry is associated with a 0.19 percentage-point increase in the probability of her newborn child being male.

The model also predicts that an increase in the relative returns to daughters leads to a lower proportion of male newborns. This prediction is tested with the coefficient on the female-specific provincial tariff cuts interacted with the $Post_t$ dummy (β_R). The empirical results offer some limited support for this prediction, particularly when considering the full specifications described by Equation 4 in Column (2). Controlling for the mother's exposure to the tariff cuts, a one percentage point reduction in local female tariff rates is linked to a 0.44 percentage point increase in the likelihood of the infant being female. However, this estimated effect is imprecise.

Lastly, the model posits that changes in the father's wage and non-labor income have no bearing on sex selection. Consistent with this prediction, both the father's exposure to the BTA through his industry (β_D) and the provincial exposure reflecting the impact of changes in non-labor income (β_I) are statistically insignificant once we account for the mother's exposure and a term that captures the relative returns of daughters.

These findings underscore the direct exposure of mothers as the primary driver of the skewed sex ratio at birth induced by the BTA. Although the trade shock may potentially impact sex selection through other avenues, such as relative returns to daughters or income effects, our horse race here lends support to the notion that mothers face a trade-off between quantity and quality when making fertility decisions.

It is useful to contextualize the effect of the mother's exposure estimated above. Over our study period, mothers of these children were exposed to an average tariff reduction of 10.62 percentage points. Our estimate implies that this 10.62 percentage point decrease in mothers' tariff rates during the study period would result in a 2.02 percentage point or 3.8% increase in the proportion of male newborns.⁶ This effect closely aligns with the decline in the likelihood of being female among newborn children documented in Figure A2. An increase in SRB from 105 (the natural level) to 113 (approximately the level reported by GSO towards the end of the study period) implies a 3.76% increase in the share of male infants, a figure remarkably close to the effect size implied by our estima-

⁶Since each percentage point reduction in mother's tariff rates is associated with 0.19 percentage point increase in the male share among newborn children, the effect associated with 10.62 percentage point reduction in mother's tariff rates is: $10.62 \times 0.19 = 2.018$ percentage points, or $2.018\%/53\% = 3.8\%$ increase from the baseline of 53% newborn male share.

tion.⁷ These estimates are in line with findings from a few other studies. For instance, Anukriti (2018) discovered that an Indian program aimed at addressing the fertility-sex ratio trade-off resulted in a 1-2.3% increase in the probability of the first birth being male. Almond et al. (2019) found that land reform in China between 1978 and 1984 led to a 3 percentage point or 5.6% increase in the fraction of males following the birth of a first girl.⁸ It is worth noting that the estimates from these studies are specific to birth parity and the gender composition of the first sibling, whereas our findings are estimated across all birth orders.

Mother's labor supply

The preceding findings regarding sex selection at birth point toward the heightened quality-quantity tradeoff that mothers face as the a result of the BTA. To delve deeper into this work-childcare trade-off, we explore how the BTA affects mothers' work hours. We employ the same regression model as described in Equation 4, but this time with the dependent variable being the mother's monthly work hours in her primary job. The sample here includes the mothers of the infants previously examined in the sex selection analysis, and the results are presented in Column (3) of Table 2.

Firstly, a mother's exposure to tariff reductions, representing changes in her wage (as detailed in Column (2) of Table 1, leads to an increase in her work hours. Specifically, mothers who experienced a one-percentage-point increase in BTA exposure through their primary job tend to work an additional 0.27 hours per month on average compared to less-exposed mothers. While this effect is statistically significant, its magnitude may seem relatively small. Notably, this effect is observed among women who recently gave birth and might be on maternity leave, which could explain the modest increase in hours for this group.⁹ For instance, for a mother employed in the apparel industry, this result implies an increase of 13.09 to 15.80 hours per month, constituting a 9.3% to 9.6% increase from a baseline of 140.74 to 163.45 hours per month among mothers in these industries.¹⁰

Furthermore, the husband's exposure, representing changes in the dad's wage, is associated with an increase in the mom's work hours. While this aligns with the predicted

⁷GSO annual statistics on sex ratio at birth can be found at <https://www.gso.gov.vn/dan-so/>

⁸from a baseline sex ratio of about 112 for this group of children

⁹The sample of mothers here is smaller than the sample of children in the sex selection analysis. This most likely reflects maternity leave, because, since 2010 wave, VHLSS has asked about the number of work hours within the last 30 days, resulting in missing values for work hours. If the mother is on maternity, she will drop out of the sample, although her child appears in the infant sample.

¹⁰Tariff reduction in textile and garments (code 17) and fur processing and fur products (code 18) industries are 48.48% and 58.53%, respectively. The implied effect on monthly work hours for mothers in these industries is an increase of $0.27 \times 48.48 = 13.09$ and $0.27 \times 58.53 = 15.80$ hours.

impact of the dad's exposure mentioned in Column (2) of [Table 1](#), the estimate remains imprecise.

Column (2) [Table 1](#) also predicts that returns to female work, proxied by the female provincial exposure, positively impact the mother's labor supply. However, in comparison to the effect of a mother's industry exposure, the impact stemming from the female provincial exposure is considerably more substantial. This suggests that mothers residing in provinces with more favorable conditions for female work tend to work significantly more hours than their counterparts in less female-exposed provinces. A one-percentage-point increase in female provincial exposure corresponds to an additional 15.5 work hours. Here, the female provincial exposure reflects the demand for female labor in the local economy, which could influence not only traded industries but also spill over into non-traded sectors. Consequently, this measure is likely to capture a much larger effect than a mother's exposure.

Finally, when conditioned a mother's direct exposure and female provincial exposure, the gender-neutral provincial exposure represents non-labor income. The findings reveal that it hurts a mother's work hours. Taken together, these results are rather consistent with the predictions outlined in Column (2) of [Table 1](#).

Fertility

Columns (4), (5), and (6) in [Table 2](#) delve into the fertility aspect of the work-childcare trade-off faced by mothers. These three columns employ the regression in [Equation 4](#) to analyze a retrospective panel derived from cross-sectional data. We report the same coefficients of interest as before. These three estimates allow us to validate the predictions outlined in Column (3) of [Table 1](#).

All else being equal, an increase in the mother's wage leads to a reduced likelihood of having an additional child. Consistent with this prediction, mom's direct exposure through her industry, which proxies for her wage change induced by the BTA, indeed hurts the chance of giving birth. The magnitude of the estimate is stable across specifications: compared to before the BTA, each percentage increase in tariff cuts in the mother's industry is associated with a 0.07-0.08 percentage point decline in her likelihood of giving birth after the policy's enactment. Considering that mothers, on average, experienced a 10.6 percentage point increase in tariff cuts between 2002 and 2016, this translates to a decrease of 0.7-0.8 percentage points, or an 8.6-9.9% reduction, from a baseline of an 8.06% chance that a woman gives birth in a given year.

On the other hand, the change in the father's wage is likely to increase the probability

of having an additional child, as stated in column (3) of [Table 1](#). This reflects the pure income effect of the husband's wage, as he faces a minimal work-childcare trade-off. The (gender-neutral) provincial exposure, which serves as a proxy for non-labor income, is also anticipated to have a positive impact on fertility. However, this estimate is insignificant in most specifications.

Lastly, the relative returns to daughters, proxied by female provincial exposure, do not play a significant role in fertility decisions. This aligns with its absence in the model's solution for optimal fertility choice. Overall, the results on fertility corroborate well with the predictions generated by the model.

5.2 True panel

In this subsection, we discuss the complementary results estimated with the 2002-2004-2006 panel of women. We estimate the same effects in a panel-data specification ([Equation 5](#)) that potentially improves upon the previous identification strategy using repeated cross-sections. Firstly, this approach allows us to account for mothers' unobserved preferences regarding fertility and sex composition of children. Secondly, it enables us to estimate the causal impact of the BTA on mothers and fathers by considering their *initial* industry affiliations before the BTA. This mitigates concerns related to potential compositional bias resulting from parents sorting into industries with greater tariff cuts.

Restricting the sample to 20-40-year-old women who recently gave birth in both the pre and post-periods of the 2002-2004-2006 panel, we have a small sample of 1,443 women and a total of 2,596 women-year observations. By construction, the sample requires the women to have at least one child in 2002, so the effect of the BTA post-2002 will be identified on the second birth and above. We present results on sex selection in Columns (1)-(3), mother's work hours in Column (4), and fertility decisions in Column (5) of [Table 3](#). It's important to note that for the fertility results, we do not require women giving birth in both the pre and post-periods, resulting in larger sample sizes.

Sex ratio at birth. In Column (1) of [Table 3](#), we estimate the equation specified in [Equation 5](#) while controlling for parents' demographic information, household characteristics, and the birth order of the child born in the survey year. As per Column (1) in [Table 1](#), two variables may have an impact on the sex of a newborn child: mom's exposure through her industry (representing her wage) and female provincial exposure (reflecting returns to daughters).

Similar to the cross-sectional findings, only the mom's exposure shows statistically significant impacts on the gender of her newborn. Mothers with greater exposure to the

Table 3: Effect on whether the last birth is male, labor supply, and fertility. Panel data

	1(Last born is male)			Hrs Worked	1(Recent Birth)
	(1)	(2)	(3)	(4)	(5)
Own Exp x Post	0.0074** (0.0031)	0.0035 (0.0031)	0.0091*** (0.0031)	0.6105 (0.5342)	-0.0004 (0.0012)
Husband Exp x Post	0.0067 (0.0042)	0.0062 (0.0037)	0.0068 (0.0043)	-0.6714 (0.7210)	0.0000 (0.0011)
Prov Exp x Post	-0.0617 (0.0756)	-0.0285 (0.0795)	-0.1168 (0.0716)	9.8204 (20.9042)	-0.0219 (0.0257)
Female Exp x Post	0.0619 (0.1272)	0.0424 (0.1235)	0.0829 (0.1110)	3.2277 (35.8795)	0.0142 (0.0348)
Had a son before		-0.5672*** (0.0636)			
R ²	0.82	0.85	0.84	0.76	0.48
Observations	2,596	2,596	2,596	2,514	11,045
Control Group Mean	0.2857	0.2857	0.2857	108.9048	0.0741
Controls	✓	✓	✓	✓	✓
Individual FE	✓	✓	✓	✓	✓
Last Child Birth Year FE	✓	✓	✓		
Survey Year FE	✓	✓	✓	✓	✓

Notes: All columns use a panel of women between 20 and 40 years old from 2002-2006 who have given birth once before and after 2002. Columns (1)-(3) report the effect on whether the woman gave birth to a boy in the last two years. Column (1) controls for household size, the woman's age, her husband's age, her education, her husband's education, urban and minority dummies. In addition to these controls, column (2) adds whether she had a son before her most recent birth. Column (3) includes 1991 agriculture and manufacturing employment shares at the province level interacted with the birth year of the last born child. Column (4) reports the effect on the number of working hours per month in her main job and includes the same controls and interactions between survey years and shares of agriculture and manufacturing. Column (5) reports the effect on whether the woman gave birth in the last two years. The regression includes the same controls as in column (4). All regressions include survey weights and are estimated by OLS. Standard errors, clustered at the province level, are reported in parentheses.

trade shock are more likely to give birth to male offspring following the enactment of the BTA compared to those with less exposure. The effect size here more than triples the analogous estimate in the cross-section regression, possibly because it is conditioned on higher birth order. Notice that the R^2 here is much larger than in cross-sectional regressions (Columns (1)-(4) in [Table 2](#)) because we can now control for mother fixed effects.

Column (2) introduces a control for whether the mother has a son before the current birth. The estimate on the interaction between mom's exposure and $Post_t$ dummy drops by half and becomes insignificant. This result aligns with our model's assumption of a preference for at least one son: once a family has at least one son, the BTA no longer imposes a quality-quantity constraint on women.

Column (3) reverts to the basic specification in column (1) but further controls for the effect of initial economic conditions over time to ensure the robustness of the shift-share variables. This slightly amplifies the effect size of a mother's direct exposure through her industry.

Mother's labor supply. Column (4) regresses the mother's number of work hours on the same independent variables as in Column (3). The sample is also restricted to the same mothers as in Column (3). As in our sample selection for the cross-sectional analysis of mother's labor supply, we focus on mothers who recently gave birth because the work-childcare constraint is most binding for them. While none of the reported estimates are statistically significant, their direction and magnitude closely align with the cross-sectional estimates. Despite their lack of precision, these panel estimates reinforce our earlier findings that support predictions in Column (2) of [Table 1](#).

Fertility. Column (5) examines the effects of the BTA on the probability of mothers giving birth. Unlike the previous four columns, the sample here no longer imposes that the women must give birth once in 2002 and at least once in 2004 or 2006. As a result, the sample size increases substantially compared to the previous four columns. However, it remains relatively small compared to the sample size of the fertility retrospective panel regression in Columns (6)-(8) [Table 2](#). This could explain why the fertility estimates lack statistical power. Nevertheless, the direction and magnitude of these panel estimates are consistent with those in the retrospective panel approach. Overall, the result in Column (5) also corroborates with predictions in Column (3) of [Table 1](#). Specifically, a mother's wage, influenced by her direct industry exposure, reduces the likelihood of her giving birth, whereas a father's wage has the opposite effect. The model also predicts that non-labor income, represented by the gender-neutral provincial exposure, increases fertility. However, this finding is not consistently supported by the data, including the results presented here and in the last three columns of [Table 2](#).

5.3 Robustness

This subsection aims to address potential concerns to identification and explore alternative explanations to our main results in [Table 2](#). [Table A6](#) demonstrates the stability of our results throughout these considerations.

To address concerns regarding the identification of the shift-share variables, we introduce interactions between survey years and the pre-BTA shares of agriculture and manufacturing at the province level, using data from the population census in 1999, in line with the approach proposed by Goldsmith-Pinkham et al. (2020). Furthermore, as argued by Ebenstein (2014), agriculture intensity predicts the strength of patrilocality, which correlates with son preference. Thus, having the initial agriculture shares \times year fixed effects also allows us to account for potential changes associated with the initial levels of patrilocality that might affect the intensity of son preference at the provincial level.

We introduce interactions between time-varying control variables and the $Post_t$ dummy to account for the potential impact of our control variables on outcomes after the trade shock.

To address the potential effects of improvements in infrastructure, such as the road network, on access to sex-selective abortions, as suggested by Almond et al. (2019), we include region-year fixed effects. In Vietnam, inter-provincial development projects are often organized into eight regions: Northeast, Northwest, Red River Delta, North Central Coast, South Central Coast, Central Highlands, Southeast, and Mekong River Delta. Even with this additional inclusion, the effects of maternal exposure on our primary outcomes remain relatively stable.

Lastly, we cluster our standard errors at the mother’s industry level instead of the province level, given that the mother’s industry emerges as the most impactful “treatment” variable. Doing so increases the statistical significance of our result on sex selection to 5%. However, since we incorporate variations in exposure at both the parents’ industry and provincial levels, we maintain our clustering at the province level in our primary results. Overall, our findings demonstrate robustness to these alternative specifications.

6 Conclusion

In this study, we explore the causal link between the work-childcare trade-off faced by mothers and their decisions regarding the gender of their newborns. Mothers grappling with a higher opportunity cost of child-rearing often find themselves compelled to limit the size of their families. However, in societies where sons are preferred, and access to

sex-selection technology is readily available, such as in Vietnam, the choice to have fewer children may be coupled with an intensified desire for male offspring. To explore this hypothesis, we leverage the 2001 US-Vietnam Bilateral Trade Agreement (BTA), a major trade agreement that coincided with the rise in the sex ratio at birth (SRB) in Vietnam.

Our findings reveal that mothers with greater exposure to the BTA through their industries are less inclined to give birth. Moreover, among those who do give birth, these mothers are more likely to have male infants. Additionally, we observe that mothers with higher exposure to the BTA tend to work longer hours, indicating an increased opportunity cost associated with child-rearing for this group. These results remain robust even when considering alternative factors that could potentially influence fertility and sex selection decisions in the context of the BTA.

While we acknowledge the existence of alternative mechanisms in the literature that may explain the rise in SRB, such as the access to ultrasound technology (Chen et al., 2013; Lin et al., 2014), inheritance motive (Bhalotra et al., 2019), divorce laws (Sun & Zhao, 2016), pension programs (Ebenstein & Leung, 2010), marriage markets (Borker et al., 2020; Wang & Chang, 2002), and targeted policy reforms (Anukriti, 2018), we contend that our model offers a more compelling explanation within our specific context. This proposition is supported by the wide range of evidence that aligns well with our model's twelve predictions.

Our findings shed light on a troubling consequence of economic growth in rapidly developing Asian economies with a preexisting preference for sons. While development policies aim to enhance economic opportunities for women, increased female participation in the workforce may inadvertently contribute to a heightened demand for male offspring. Given the pivotal role that women play in economic development, the resulting bias in favor of sons in fertility outcomes is likely to persist in such contexts. Consequently, addressing the gender imbalance in these countries will require concerted efforts, including awareness campaigns and initiatives to promote gender equality norms.

References

- ActionAid-Vietnam. (2016). *Unpaid Care Work: Make a House becomes a Home* (tech. rep.). ActionAid Vietnam. Hanoi. Retrieved October 7, 2022, from <https://vietnam.actionaid.org/en/publications/2016/policy-brief-unpaid-care-work-make-house-becomes-home>
- Almond, D., Li, H., & Zhang, S. (2019). Land reform and sex selection in china. *Journal of Political Economy*, 127(2), 560–585.
- Anukriti, S. (2018). Financial incentives and the fertility-sex ratio trade-off. *American Economic Journal: Applied Economics*, 10(2), 27–57.
- Anukriti, S., & Kumler, T. J. (2019). Women’s worth: Trade, female income, and fertility in India. *Economic Development and Cultural Change*, 67(3), 687–724.
- Autor, D., Dorn, D., & Hanson, G. (2019). When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men. *American Economic Review: Insights*, 1(2), 161–178.
- Becker, G. S., & Lewis, H. G. (1973). On the Interaction between the Quantity and Quality of Children. *Journal of Political Economy*, 81(2, Part 2), S279–S288.
- Becquet, V., & Guilmoto, C. Z. (2018). Sex imbalance at birth in vietnam: Rapid increase followed by stabilization. *Population*, 73(3), 519–544.
- Bhalotra, S., Chakravarty, A., Mookherjee, D., & Pino, F. J. (2019). Property Rights and Gender Bias: Evidence from Land Reform in West Bengal. *American Economic Journal: Applied Economics*, 11(2), 205–237.
- Bongaarts, J., & Guilmoto, C. Z. (2015). How Many More Missing Women? Excess Female Mortality and Prenatal Sex Selection, 1970–2050. *Population and Development Review*, 41(2), 241–269.
- Borker, G., Eeckhout, J., Luke, N., Minz, S., Munshi, K., & Swaminathan, S. (2020). Wealth, Marriage, and Sex Selection.
- Carranza, E. (2014). Soil Endowments, Female Labor Force Participation, and the Demographic Deficit of Women in India. *American Economic Journal: Applied Economics*, 6(4), 197–225.
- Chakraborty, T. (2015). Trade Liberalization in a Traditional Society: Implications for Relative Female Survival. *World Development*, 74, 158–170.
- Chao, F., Gerland, P., Cook, A. R., & Alkema, L. (2019). Systematic assessment of the sex ratio at birth for all countries and estimation of national imbalances and regional reference levels. *Proceedings of the National Academy of Sciences of the United States of America*, 116(19), 9303–9311.

- Chen, Y., Li, H., & Meng, L. (2013). Prenatal sex selection and missing girls in China: Evidence from the diffusion of diagnostic ultrasound. *Journal of Human Resources*, 48(1), 36–70.
- Dasgupta, A., & Sharma, A. (2022). Missing Women: A Review of Underlying Causes and Policy Responses. In *Oxford Research Encyclopedia of Economics and Finance*.
- Ebenstein, A. (2010). The “Missing Girls” of China and the Unintended Consequences of the One Child Policy. *Journal of Human Resources*, 45(1), 87–115.
- Ebenstein, A. (2014). Patrilocality and Missing Women. *SSRN Electronic Journal*.
- Ebenstein, A., & Leung, S. (2010). Son preference and access to social insurance: Evidence from China’s rural pension program. *Population and Development Review*, 36(1), 47–70.
- Fukase, E. (2013). Export Liberalization, Job Creation, and the Skill Premium: Evidence from the US-Vietnam Bilateral Trade Agreement (BTA). *World Development*, 41(1), 317–337.
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8), 2586–2624.
- Guilmoto, C. Z. (2012). Son Preference, Sex Selection, and Kinship in Vietnam. *Population and Development Review*, 38(1), 31–54.
- Guilmoto, C. Z., Hoàng, X., & Van Ngo, T. (2009). Recent increase in sex ratio at birth in Viet Nam. *PLoS ONE*, 4(2).
- Jayachandran, S. (2017). Fertility decline and missing women. *American Economic Journal: Applied Economics*, 9(1), 118–139.
- Jones, L., Schoonbroodt, A., & Tertilt, M. (2010). Fertility Theories: Can They Explain the Negative Fertility-Income Relationship? Retrieved November 30, 2020, from <http://www.nber.org/papers/w14266>
- Li, H., Yi, J., & Zhang, J. (2011). Estimating the Effect of the One-Child Policy on the Sex Ratio Imbalance in China: Identification Based on the Difference-in-Differences. *Demography*, 48(4), 1535–1557.
- Lin, M. J., Liu, J. T., & Qian, N. (2014). More missing women, fewer dying girls: The impact of sex-selective abortion on sex at birth and relative female mortality in Taiwan. *Journal of the European Economic Association*, 12(4), 899–926.
- McCaig, B. (2011). Exporting out of poverty: Provincial poverty in Vietnam and U.S. market access. *Journal of International Economics*, 85(1), 102–113.
- McCaig, B., & Pavcnik, N. (2018). Export markets and labor allocation in a low-income country. *American Economic Review*, 108(7), 1899–1941.

- Qian, N. (2008). Missing women and the price of tea in China: The effect of sex-specific earnings on sex imbalance. *Quarterly Journal of Economics*, 123(3), 1251–1285.
- Rosenzweig, M. R., & Schultz, T. P. (1982). Market opportunities, genetic endowments, and intrafamily resource distribution: Child survival in rural India. *American Economic Review*, 72(4), 803–815.
- Sen, A. (1990). More Than 100 Million Women Are Missing. *The New York Review*. Retrieved September 6, 2022, from <https://www.nybooks.com/articles/1990/12/20/more-than-100-million-women-are-missing/>
- Sobek, M., & Kennedy, S. (2009). The Development of Family Interrelationship Variables for International Census Data. *Minnesota Population Center Working Paper Series*.
- STAR-Vietnam. (2003). *An Assessment of the Economic Impact of the United States - Vietnam Bilateral Trade Agreement* (tech. rep.). The National Political Publishing House. Hanoi. Retrieved October 7, 2022, from https://pdf.usaid.gov/pdf_docs/Pnadc703.pdf
- Sun, A., & Zhao, Y. (2016). Divorce, abortion, and the child sex ratio: The impact of divorce reform in China. *Journal of Development Economics*, 120, 53–69.
- Wang, H.-z., & Chang, S.-m. (2002). The Commodification of International Marriages: Cross-border Marriage Business in Taiwan and Viet Nam. *International Migration*, 40(6), 93–116.

A Appendix: Family interrelationship data

Variables categorizing family interrelationship within the household are crucial in this paper. In particular, we need “pointer” variables identifying each person’s mother, father, and spouse. VHLSS has a variable that classifies each individual’s relationship to household head. However, this variable is ambiguous and erroneous for the following reasons. First, extended family is common in Vietnam: roughly 34.91% of VHLSS sample live in the same household with relatives other than their nuclear family members. Because VHLSS does not identify subfamilies within the extended household, the “relate-to-head” variable will miss out the interrelationship of subfamilies other than the head’s. Second, with the exception of VHLSS 2002, all other waves do not identify “children-in-law.” An immediate consequence is missing spousal links and therefore parental links in many cases. Thus, we write an algorithm to improve the precision of parent-child and spousal pairings in VHLSS.

We adapt a similar algorithm that has been used by the Minnesota Population Center (Sobek & Kennedy, 2009) to generate “pointer” variables for IPUMS–International Census Data. We utilize the four variables, relationship to household head, age, and marital status, in combination with the relative position of household members in the roster listing, to infer relationships. We first establish spousal links (generate variable SPLOC), then we find mother and father for each individual (generate variables MOMLOC and POPLOC). Regarding parental linkage, we look for the mother before looking for the father; once we find the mother for a child, we assign her husband as the father of this child. Only when we cannot a child’s mother, we locate its father (and assign his wife as the mother).

Tables A1 and A2 document the rules we apply to these matching tasks. Our algorithm applies the rules sequentially: if the first rule finds a match (spousal match or parental matches) for a given individual, the second rule no longer applies to this person, and so on. Whenever there are ambiguous multiple potential spouses or multiple potential parents, we drop the entire household from the sample.

Fortunately, VHLSS 2014 and 2016 provide two variables to locate the biological fathers and mothers of children under 16 years old. We use this information to test our algorithm and report the results in Table A3. We find the correct mother for 93.23% of these children and the correct father for 90.99% of them. The algorithm also finds both parents correctly for 88.95% of the children. The correction rate for both biological parents is lower than that for each biological parent because our algorithm also counts stepmothers and stepfathers. With these results, we are very confident that our algorithm does a

good job at identifying family interrelationship.

Table A1: Rules for SPLOC construction

Rule	Individual's relationship to head	Partner's relationship to head	Age difference	Both Married	Require adjacency	Only applicable to 2002	Notes
Strong couple pairing, couple adjacency preferred							
	Head	Spouse	No	No	No		1
	Parent	Parent	No	Yes	Yes		1
	Grandparent	Grandparent	No	Yes	Yes		1
	Child	Child-in-law	No	Yes	Yes	Yes	1
Weak couple pairing, couple adjacent							
	Grandchild	Grandchild	Yes	Yes	Yes		1, 2
	Other relationship	Other relationship	Yes	Yes	Yes		1, 2
	Sibling	Sibling	Yes	Yes	Yes	Yes	1, 2
	Grandchild	Other relationship	Yes	Yes	Yes		1, 2
	Sibling	Other relationship	Yes	Yes	Yes	Yes	1, 2
	Child	Other relationship	Yes	Yes	Yes		1, 2
Weak couple pairing, special type child-child							
	Child	Child	Yes	Yes	Yes		1, 2, 6
Weak couple pairing, couple not adjacent							
	Child	Child-in-law	No	Yes	Closest proximity	Yes	5
	Child	Other relationship	Yes	Yes	No		2, 3
	Child	Child	Yes	Yes	No		1, 2, 4, 6

Notes:

1. Drop the entire household if there is any person that could be assigned to 2 couples by the adjacency rule.
2. A woman can be no more than 20 years older or 35 years younger than a potential male partner.
3. For non-adjacent couple pairing, among the potential spouses who satisfy the age and marital requirements, select the person who is closest in age and impose that the husband is older than the wife.
4. For child-child non-adjacent couple pairing, drop all households where this rule yields multiple potential spouses.
5. These non-adjacent couples are matched based on having closest proximity to each other.
6. Once the couple is identified, the in-law is disinguished. For child-child couples, assume the first listed spouse in the biological child to the household head, hence the second listed spouse is the child-in-law.

Table A2: Rules for MOMLOC and POPLOC construction

Rule	Child's relationship to head (updated after SPLOC is generated)	Parent's relationship to head (updated after SPLOC is generated)	Age difference	Proximity requirement	Only applicable to 2002	Notes
Links involving Head, Spouse, and Grandparent (unambiguous)						
	Child	Head, spouse	No	No		
	Child-in-law	Head, spouse	No	No	Yes	
	Head	Parent	No	No		
	Spouse	Parent	No	No		
	Sibling	Parent	No	No	Yes	
	Parent	Grandparent	No	No		
Links between grandchildren and children						
	Grandchild	Child, child-in-law	15-44	Weak		1
	Grandchild	Other relationship	15-44	Weak		1, 2
Links involving other relatives						
	Other relationship	Grandchild	15-44	Weak		1
	Other relationship	Other relative	15-44	Weak		1
	Other relationship	Sibling	15-44	Weak		1
	Other relationship	Child, child-in-law	15-44	Weak		1, 3

Notes:

1. Weak proximity requires that the child must be listed after its potential mothers (or potential fathers if it has no potential mothers); among them, its mother is the one listed closest to the child.
2. Impose that no person with code "Child" is present in the household. The mother of the grandchildren in these cases tend to be listed as "Others" since there is no category for "Child-in-law."
3. Impose that no person with code "Grandchild" is present in the household. These cases tend to mix up the numerical code for "Grandchild" (code 6) and that for "Other relationship" (code 7).

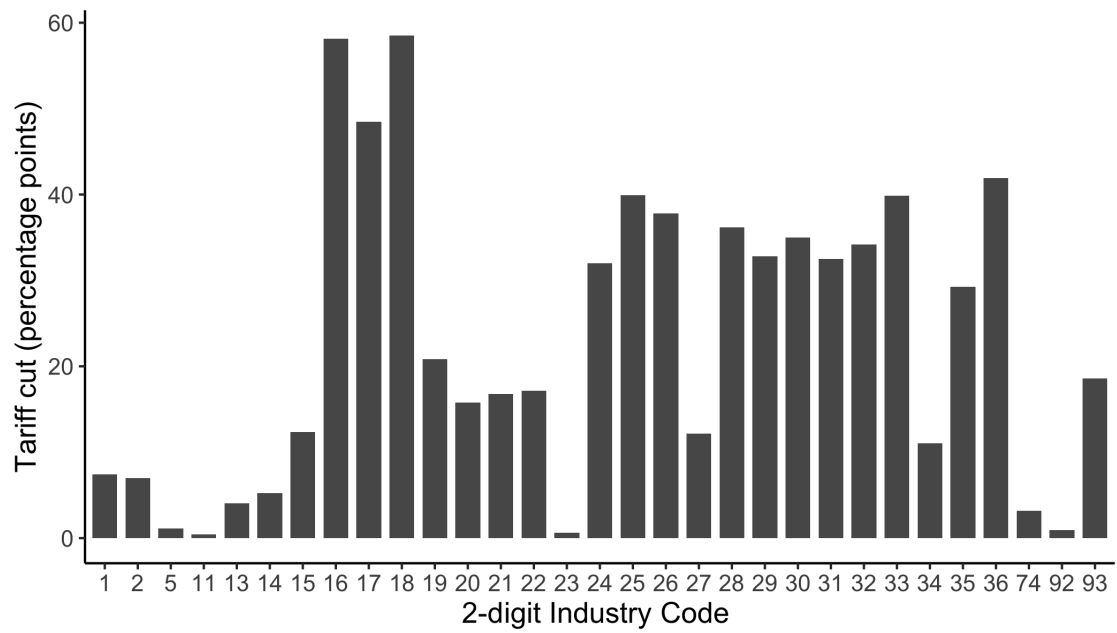
Table A3: Performance of our algorithm in locating parents

	Freq.	Percent	Cum.
Both mom and dad correct	80021	88.95	88.95
Dad correct, mom incorrect	1831	2.04	90.99
Dad incorrect, mom correct	3853	4.28	95.27
Both mom and dad incorrect	4253	4.73	100.00
Total	89958	100.00	

Notes: This table summarizes the comparison between the parent locators generated by our algorithm with the true parents locators provided by VHLSS 2014-2016 for children under 16 years old.


B Appendix: Additional Figures

Figure A1: US tariff cuts by 2-digit ISIC industry codes



Notes: Data come from McCaig and Pavcnik (2018)

Figure A2: Sex ratio trend in VHLSS



Notes: This graph plots the birth year fixed-effect coefficients from a regression of female dummy on birth year fixed effects and survey year fixed effects. The omitted group is children born in 2003. The sample of this regression consists of newborn children between 0 and 1 year old.

C Appendix: Additional Tables

Table A4: Descriptive Statistics Repeated Cross-sectional Data

Characteristic	2002	2004	2006	2008	2010	2012	2014	2016
Male	53%	53%	53%	51%	53%	53%	53%	52%
Mom Exposure	8.82 (11.40)	9.43 (12.85)	10.19 (14.12)	10.03 (14.67)	10.66 (15.36)	10.97 (15.76)	12.27 (17.33)	11.80 (16.91)
Dad Exp.	7.80 (10.52)	7.75 (10.99)	7.95 (11.61)	7.49 (11.78)	8.27 (12.77)	7.68 (12.30)	8.11 (13.00)	8.19 (13.08)
Province Exp.	8.90 (1.75)	9.14 (2.04)	8.97 (1.89)	9.15 (2.03)	9.09 (1.93)	9.04 (1.86)	8.86 (1.49)	8.94 (1.77)
Female Exp.	4.10 (1.04)	4.24 (1.23)	4.15 (1.13)	4.24 (1.21)	4.19 (1.14)	4.15 (1.10)	4.02 (0.85)	4.08 (1.05)
Household size	6.25 (2.29)	6.49 (2.33)	6.31 (2.22)	6.31 (2.29)	5.68 (1.84)	5.81 (1.79)	5.84 (1.73)	5.79 (1.66)
Mom birth age	27.10 (6.43)	27.25 (6.53)	26.85 (6.42)	27.02 (6.25)	26.78 (5.99)	26.80 (5.95)	26.96 (5.98)	27.12 (6.08)
Mom edu	6.57 (3.59)	7.06 (3.70)	7.32 (3.71)	7.94 (3.66)	8.19 (3.65)	8.42 (3.65)	9.07 (3.37)	9.26 (3.41)
Dad birth age	30.38 (6.72)	30.81 (6.87)	30.38 (6.69)	30.79 (6.74)	30.23 (6.45)	30.17 (6.36)	30.71 (6.67)	30.41 (6.64)
Pop edu	7.07 (3.51)	7.48 (3.50)	7.66 (3.56)	8.22 (3.48)	8.35 (3.52)	8.59 (3.44)	9.10 (3.25)	9.20 (3.28)
Mom's Main Job (hr/month)	110.42 (76.34)	122.05 (67.33)	125.26 (67.96)	127.32 (69.88)	150.29 (70.56)	149.68 (70.35)	151.43 (68.88)	153.86 (69.67)
Urban	21%	24%	26%	30%	28%	27%	28%	27%
Minority	20%	20%	21%	19%	20%	20%	21%	24%
Num. Obs.	8,234	4,569	4,482	5,826	5,365	5,089	6,538	4,826

¹ %; Mean (SD); Total N (unweighted)

Table A5: Descriptive Statistics Panel Data

Characteristic	2002	2004	2006
Had a child within the past 2 years	17%	23%	24%
Main job (hr/month)	132.64 (74.26)	140.12 (69.40)	144.04 (68.15)
Own Exposure	8.62 (11.30)	9.43 (12.94)	9.91 (13.87)
Husband Exp.	7.14 (9.33)	7.19 (9.82)	7.16 (10.03)
Province Exp.	9.04 (1.87)	9.13 (1.99)	9.11 (1.98)
Female Exp.	4.16 (1.12)	4.23 (1.20)	4.21 (1.18)
Age	31.49 (5.74)	31.13 (6.21)	30.40 (6.47)
Education	7.07 (3.48)	7.43 (3.54)	7.75 (3.61)
Urban	23%	24%	27%
Minority	15%	15%	16%
Num. Obs.	5,707	5,524	5,486

¹ %; Mean (SD); Total N (unweighted)

Table A6: Robustness: Effects on whether a new birth is male and labor supply

	1(male)	Mom's Hrs Worked	1(male)	Mom's Hrs Worked	1(male)	Mom's Hrs Worked	1(male)	Mom's Hrs Worked
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mom Exp x Post	0.0019** (0.0009)	0.2732*** (0.1022)	0.0019** (0.0009)	0.2792*** (0.1015)	0.0019** (0.0008)	0.2561** (0.1038)	0.0019*** (0.0005)	0.2732* (0.1495)
Dad Exp x Post	-0.0003 (0.0009)	0.1648 (0.1215)	-0.0003 (0.0009)	0.1800 (0.1224)	-0.0002 (0.0009)	0.1968 (0.1210)	-0.0003 (0.0005)	0.1648* (0.0935)
Prov Exp x Post	-0.0251 (0.0191)	-9.4960*** (3.0085)	-0.0295 (0.0211)	-9.4021*** (2.5067)	-0.0166 (0.0258)	-0.4450 (2.3605)	-0.0251 (0.0175)	-9.4960*** (3.1083)
Female Exp x Post	-0.0492* (0.0273)	15.5537** (6.9000)	-0.0379 (0.0283)	12.5827** (6.1951)	-0.0576 (0.0360)	-4.1198 (5.2211)	-0.0492 (0.0306)	15.5537** (7.5672)
Standard-Errors				Province				ISIC2
R ²	0.011	0.337	0.011	0.338	0.013	0.344	0.011	0.337
Observations	30,770	28,303	30,770	28,303	30,770	28,303	30,770	28,303
Control Group Mean	0.4898	134.1040	0.4898	134.1040	0.4898	134.1040	0.4898	134.1040
Controls	✓	✓	✓	✓	✓	✓	✓	✓
ControlsPost			✓	✓				
svyear fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
region8-svyear fixed effects					✓	✓		
Province FE	✓	✓	✓	✓	✓	✓	✓	✓
Mom Industry FE	✓	✓	✓	✓	✓	✓	✓	✓
Dad Industry FE	✓	✓	✓	✓	✓	✓	✓	✓
Birth Year FE	✓		✓				✓	

Notes: All regressions are restricted to a sample of children between 0 and 1 year old. Column (1) reports the effects on whether the infant is a boy. It runs a similar regression to Column (4) in Table 2 but adds 1999 agricultural and manufacturing employment shares at the province level interacted with birth years. Column (2) reports the effects on the number of working hours per month of the main job of these children's mothers. It includes all controls of Column (5) in Table 2 plus the interactions between 1999 agricultural and manufacturing employment shares at the province level and birth years. Columns (3) and (4) add time-varying controls interacting with dummy for Post 2002. Columns (5) and (6) add region-survey year fixed effects to (1) and (2). All regressions in columns (1)-(6) cluster the standard errors at the province level. Columns (7) and (8) run the same regressions as in columns (5) and (6) but cluster the standard errors at the mother's industry levels. All regressions include survey weights and are estimated by OLS. The clustering level of the standard errors are reported above the R² row.