Trade, Maternal Time Costs, and Sex Selection: Evidence from Vietnam*

Nghiem Q. Huynh† Ngoc T. T. Nguyen‡

May 2025

Abstract

We study how trade liberalization affects sex selection in a son-preferring society. Using the Vietnam–US 2001 trade agreement as a natural experiment, we exploit industry-level tariff cuts in a difference-in-differences framework. We find that women in exposed industries are more likely to have male children, work more, and have fewer births—effects driven solely by maternal, not paternal, exposure. A quantity-quality model with maternal time costs explains these patterns. Our findings reveal that rising maternal opportunity costs, rather than income or relative returns effects, can increase sex selection during economic transitions.

JEL Classification: F13, J13, J16, J22, O24

*We thank S. Anukriti, Tim Guinanne, Arthur Lewbel, Kaivan Munshi, Rohini Pande, Mark Rosenzweig, and Mu-Jeung Yang for their valuable suggestions. We also appreciate the feedback from seminar participants at Boston College, Yale, and the University of Oklahoma. Special thanks to Brian McCaig for his guidance on the VHLSS data. Huynh acknowledges financial support from the Yale Council on Southeast Asian Studies.

†University of Oklahoma, nqh@ou.edu

‡Lightcast, jade.ng.nguyen@gmail.com

1 Introduction

How does economic development influence sex selection when parents face pressures from work, childcare, and son preference? Since Becker (1960), economists have studied the quantity-quality tradeoff parents face between family size and investment per child. In contexts where parents strongly prefer sons, however, this tradeoff takes on an additional dimension—they must balance their desire for a son against the rising costs of additional children. When mothers bear primary responsibility for childcare while participating in the labor market, improved economic opportunities can create particularly acute pressure on maternal time, affecting both fertility and sex selection decisions.

This focus on maternal time offers a new perspective on how economic development shapes sex selection. Previous studies have proposed three main explanations for biased sex ratios in developing countries. First, Almond et al. (2019) argue that rising household incomes increase sex selection by making it more affordable as a normal good. Second, Qian (2008) emphasizes the role of relative economic returns to daughters versus sons. Third, Jayachandran (2017) provides indirect evidence for a fertility-selection mechanism by showing how fertility decline intensifies pressure for sex selection as parents face a stricter constraint on achieving a son. However, evidence remains mixed on these channels, partly because they are difficult to disentangle empirically. Moreover, these explanations largely overlook the asymmetric burden of childcare between mothers and fathers, a global regularity (Doepke et al., 2023).

Vietnam's experience provides an ideal setting to test these competing mechanisms. The country's sex ratio at birth rose from 107 boys per 100 girls in 1995 to 112 in 2015 amid rapid economic integration, mirroring a phenomenon observed across many developing countries (Figure 1). Crucially, the timing of this acceleration coincides with the 2001 Bilateral Trade Agreement (BTA) with the United States, which generated large, industry-specific changes in labor market opportunities. We posit that this trade agreement could affect sex selection through four distinct channels: maternal employment, paternal employment, household non-labor income, and relative return to daughters.

To discipline and disentangle these forces, we formalize our intuition through a standard quality-quantity (Q-Q) model that emphasizes families' desire for at least one son and mothers' work-childcare tradeoff (Jones et al., 2011). Our model incorporates two critical features: asymmetric childcare responsibilities for mothers and son preference with diminishing returns to additional sons beyond the first. It shows how rising maternal market outcomes can increase sex selection even while improving women's economic status, contrasting with standard relative return explanations. Furthermore, our model predicts that

Figure 1: Sex Ratio at Birth Trends Across Countries

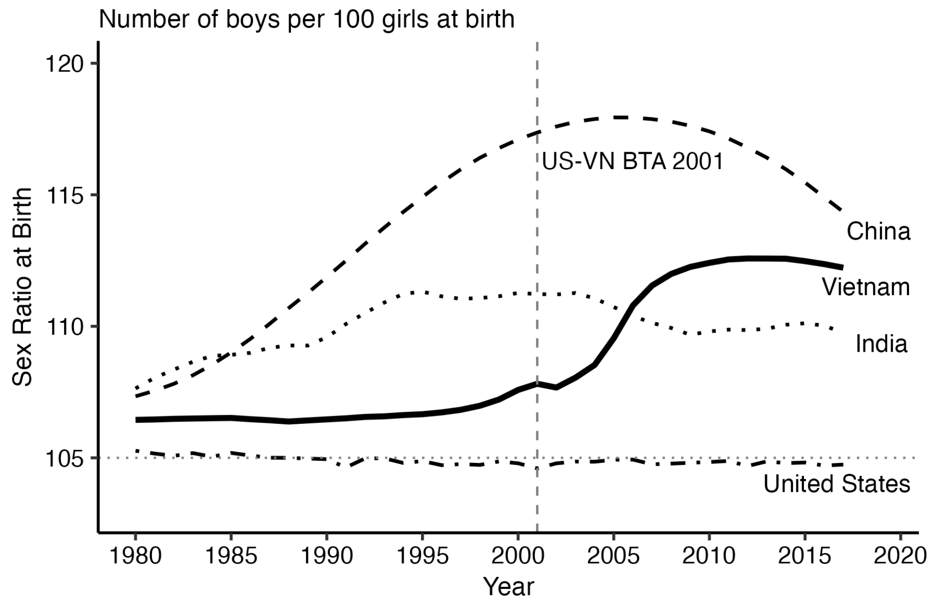

Data source: Chao et al. (2019)
Notes: The figure shows the sex ratio at birth (boys per 100 girls) for several countries. The vertical dashed line indicates the year when the bilateral trade agreement was signed. The horizontal dotted line is the biological sex ratio.

only maternal wages affect sex selection, while household non-labor income and paternal wages influence fertility but not sex selection. These predictions diverge from income-based explanations like Almond et al. (2019), in which all income sources raise both fertility and sex selection. They also differ from bargaining models, which posit that paternal wages can shift bargaining power, thereby affecting sex selection.

We overcome data limitations in previous research on sex selection by using large household surveys, the Vietnam Living Standard Surveys (VHLSS), from 2002 to 2016, rather than relying on decennial census data. Thus, we observe parents' industry and location at the time fertility decisions are made. Furthermore, we exploit variation in BTA exposure across industries and provinces to capture four channels: mother's industry-specific tariff cuts, father's industry-specific tariff cuts, local province-level economic effects, and local female labor demand. This approach allows us to disentangle how maternal opportunity costs, household income, and returns to daughters separately influence sex selection and fertility.

The results strongly support our model's emphasis on maternal time costs as the key mechanism linking trade liberalization to demographic outcomes. Women whose industries faced tariff cuts were 9 percentage points more likely to have male children, which translates to approximately a 0.72 percentage point increase with the median tariff cut of 8 percentage points (increasing sex ratio from 107 to 110.2 boys per 100 girls). These women also worked 23% more hours per month (approximately 29 additional hours) and had a 0.6 percentage point lower annual birth probability (a 30% reduction from the baseline). Consistent with our model, exposure through fathers' industries or local labor markets shows no impact on sex selection, though paternal exposure increases fertility by about 25% as predicted by income effects.

A potential limitation of our empirical setting is that the VHLSS begins in 2002, leaving us with only one survey wave before the December 2001 BTA[1]. Although pre-BTA data are limited, the biological nature of fertility enhances our identification strategy. The nine-month gestation lag ensures that all births recorded before 2002 resulted from decisions made prior to the labor market's reaction to the tariff reductions. Moreover, this gradual nature of fertility adjustments allows us to consider births in 2002 and early 2003 as pre-treatment cohorts. Consistent with this timing, our event study indicates that there are no detectable treatment effects until 2004.

Furthermore, we conduct mother-panel analyses with individual fixed effects, tracking the same women over time. This approach controls for unobserved time-invariant preferences regarding fertility and son preference. Among women who gave birth both before and after the policy change, the sex-selection effect was approximately twice as large as observed in cross-sectional analysis. This panel approach also allowed us to test a key model prediction: families become less responsive to economic incentives once they have achieved their goal of having a son. When controlling for prior son status, the maternal exposure effect decreased by about half and lost statistical significance, while the coefficient for already having a son remained large and highly significant.

Our findings contribute to several literatures. First, we advance research on sex selection determinants by testing our Q-Q model against competing explanations. By highlighting maternal time costs, we provide a novel perspective on why economic development can worsen gender imbalances despite improving women's economic status. This mechanism has important policy implications. Supporting maternal childcare responsibilities could mitigate rising sex ratios, as trade-induced female employment opportunities can paradoxically reduce female births when son preference and childcare burdens coexist.

Second, we extend research on quantity-quality tradeoffs in fertility decisions. The classic literature examines how parents balance family size against investments in each child's human capital, such as education (Rosenzweig & Wolpin, 1980; Rosenzweig & Zhang, 2009). Recent work shows how son preference can skew these post-birth investments toward boys

[1] VHLSS has a 12-month recall period, so the 2002 wave substantially captures data from 2001.

(Anukriti et al., 2022; Jayachandran & Kuziemko, 2011). We demonstrate that parents face an additional pre-birth tradeoff between family size and gender composition. This insight is particularly relevant to contexts where ex-ante birth planning decisions can be more important than subsequent human capital investment, such as Vietnam (Do & Phung, 2010).

Third, we contribute to the literature on trade's impact on gender outcomes in developing countries. While previous studies have documented mixed effects of trade on female labor market participation (Gaddis & Pieters, 2017; Heath & Mushfiq Mobarak, 2015), fewer have examined the consequences for sex selection and fertility decisions. Previous studies examining trade's effects on sex selection in India (Anukriti & Kumler, 2019; Chakraborty, 2015) rely on decennial census data, requiring strong assumptions about the stability of parents' employment and location over long periods. Our analysis exploits detailed household surveys conducted every two years, allowing us to observe parents' industry and location when fertility decisions are made. This data advantage, combined with our theoretical framework that delivers sharp predictions across multiple outcomes, enables more precise identification of how trade liberalization affects fertility choices.

Finally, we build on studies of Vietnam's BTA with the US (McCaig, 2011; McCaig & Pavcnik, 2018; McCaig et al., 2022) by examining its demographic consequences. While others have documented Vietnam's rising sex ratio, our study is the first to propose and test a specific economic mechanism using exogenous variation from trade liberalization. Our findings align with demographic patterns documented by Becquet and Guilmoto (2018), who observed similarly skewed sex ratios at first births and higher parities (3+), but we provide an economic explanation for these patterns rooted in maternal time constraints.

The remainder of the paper is organized as follows. Section 2 outlines the simple Q-Q model with testable predictions. Section 3 provides the institutional background on the 2001 US-Vietnam BTA and the data employed in our empirical analysis. Section 4 lays out the empirical strategy for our tests of the model predictions. Section 5 discusses the results, first for the repeated cross-section sample, then for the panel sample. We go through many robustness checks in Section 6. Lastly, Section 7 concludes.

2 Theoretical Framework

We adapt a standard quantity-quality (Q-Q) model from Jones et al. (2011), integrating son preference and asymmetric maternal childcare responsibilities, to examine how trade liberalization affects sex selection in Vietnam. Specifically, our model captures two distinctive features relevant to the Vietnamese context: (1) households strongly prefer having *at least one son*, and (2) mothers bear the primary responsibility for childcare, which imposes significant constraints on their labor supply.

Intuitively, the model highlights that when maternal wages rise, the opportunity cost of childcare also rises, strengthening parents' incentives to ensure a son is born early and reducing overall fertility. In contrast, increases in household income or paternal wages primarily affect total fertility but have little effect on the child's gender.

Formally, consider a household that comprises a mother (m) and a father (d). They jointly decide consumption (c_m, c_d), leisure time (ℓ_m, ℓ_d), the number of children (n), and sex selection (q). Each household derives greater utility from having at least one son. We model this preference by defining the *effective* number of sons as $Q = qn$, where $q \in [0, 1]$ is the probability of having a male child. This captures households' preference for having one son ($Q = 1$) rather than multiple daughters without any sons ($Q = 0$).

Formally, each parent's utility U_g (where $g \in \{m, d\}$) is:

$$U_g = \alpha_c \log(c_g) + \alpha_\ell \log(\ell_g) + \alpha_n \log(n) + \alpha_q \log(Q), \tag{1}$$

where the preference weights $(\alpha_c, \alpha_\ell, \alpha_n, \alpha_q)$ sum to one.

Each parent has a time endowment normalized to one, allocated between leisure and labor. Crucially, mothers incur a childcare time cost γn per child, directly reducing their available labor time. Fathers face no similar constraint[2]. Thus, the parents' time constraints are:

$$\ell_m + h_m + \gamma n = 1, \tag{2}$$

$$\ell_d + h_d = 1, \tag{3}$$

where h_m and h_d are labor hours for mothers and fathers, respectively. Given these constraints, higher maternal wages directly increase the opportunity cost of childcare, which influences fertility and sex selection decisions uniquely through mothers.

In addition to labor income (w_m, w_d), the household receives non-labor income I, such as transfers or returns from land holdings (Almond et al., 2019). The household's budget constraint is:

$$c_m + c_d + p_q Q \leq I + w_m h_m + w_d h_d, \tag{4}$$

where p_q represents the marginal cost of sex selection (including economic, physical, psychological costs, and forgone benefits of daughters).

[2]In Vietnam, the female employment rate is high, approximately 77% among prime-aged women (20-64 years old) in our household data. Furthermore, women spend over five hours per day on childcare compared to men's three hours, and this gap widens with larger family sizes (ActionAid-Vietnam, 2016).

The household maximizes a weighted sum of parental utilities, reflecting bargaining power (λ_m, λ_d):

$$\max_{c_m, c_d, \ell_m, \ell_d, n, q} \lambda_m U_m + \lambda_d U_d, \tag{5}$$

subject to the above budget and time constraints, with bargaining weights satisfying $\lambda_m + \lambda_d = 1$.

By defining the household's total income as $W = I + w_m + w_d$, the optimal choices in log form are:

$$\log q = \log(\gamma \alpha_q / \alpha_n) - \log p_q + \log w_m, \tag{6}$$

$$\log \ell_m = \log(\lambda_m \alpha_\ell) + \log W - \log w_m, \tag{7}$$

$$\log n = \log(\alpha_n / \gamma) + \log W - \log w_m. \tag{8}$$

Our model generates twelve predictions summarized in Table 1. Notably, maternal wages uniquely affect sex selection by altering the mother's opportunity cost of childcare, whereas paternal wages and non-labor income influence fertility solely through income effects, not affecting sex selection.

In contrast, Equation (6) shows that paternal wages (w_d) and non-labor income (I) do not influence sex selection, as these do not alter the mother's childcare time constraints. Furthermore, an increase in the cost of selecting sons (p_q), which can be proximate by higher local returns to female labor, reduces sex selection because higher returns to daughters decrease the relative preference for sons.

Regarding maternal labor supply in Equation (7), higher maternal wages increase maternal labor supply because the substitution effect dominates the income effect. Conversely, higher paternal wages and non-labor income reduce maternal labor supply through pure income effects, as leisure becomes more affordable.

For fertility, higher maternal wages reduce fertility ($\frac{\partial \log n}{\partial \log w_m} < 0$) due to the increased opportunity cost of childcare as seen in Equation (8). In contrast, higher paternal wages and non-labor income increase fertility via income effects, with no influence from the cost of sex selection (p_q), which affects only the quality dimension (q).

Our model's predictions contrast with alternative theories. For instance, unlike Almond et al. (2019), where higher non-labor income or paternal wages increase both fertility and sex selection through pure income effects, our model predicts these factors influence only fertility, not sex selection.

Similarly, Qian (2008) suggests that higher returns to female labor increase preference for daughters. Our model clarifies that this effect depends crucially on whether returns affect

Table 1: Comparative Statics of the Model

	Pr(Male Birth)	**Mom's Labor Supply**	**Fertility**
$\uparrow w_m$ (Mother's Wage)	$+$	$+$	$-$
$\uparrow w_d$ (Father's Wage)	0	$-$	$+$
$\uparrow I$ (Non-labor Income)	0	$-$	$+$
$\uparrow p_q$ (Cost of Son Selection)	$-$	$+$	0

Notes: $+$ indicates a positive effect; $-$ indicates a negative effect; 0 indicates no effect. Maternal wages (w_m) uniquely affect sex selection and maternal labor supply through changes in childcare opportunity costs. In contrast, paternal wages and non-labor income affect fertility purely via income effects.

maternal wages directly (which increase son preference due to childcare opportunity costs) or reflect broader local labor market conditions (captured by p_q), which could indeed favor daughters.

Our model yields testable predictions that distinguish the maternal time-cost mechanism from competing explanations for sex selection. To empirically evaluate these predictions, we need data on three key aspects: (1) exogenous variation in economic incentives that differentially affect mothers, fathers, and their nonlabor income, (2) information on both parents' industry affiliations at the time of fertility decisions, and (3) measures of sex selection, labor supply, and fertility outcomes. Vietnam's 2001 Bilateral Trade Agreement with the United States, coupled with detailed household survey data, provides an ideal setting to test these predictions and identify the causal impact of maternal opportunity costs on demographic outcomes.

3 Data and Measures

3.1 Trade Liberalization and Measures of Exposure

The 2001 US-Vietnam Bilateral Trade Agreement provides an ideal setting to study how trade liberalization affects fertility and sex selection decisions. Before the BTA went into effect on December 10, 2001, Vietnamese exports to the US faced high tariffs under the Column 2 schedule applied to non-market economies. The BTA granted Vietnam Most Favored Nation (MFN) status, dramatically reducing US tariffs on Vietnamese exports while having minimal impact on Vietnam's import tariffs.[3] Following the agreement, Vietnam's exports to the US surged from 0.7 billion dollars in 2000 to nearly 10 billion dollars in 2008.

The size of tariff reductions varied substantially across industries (Figure A1). Manufacturing sectors experienced cuts averaging 30 percentage points, with textiles and garments seeing reductions of more than 50 percentage points. We construct three types of measures to capture how these differential tariff cuts affect fertility and sex selection through four labor market channels asserted by the model.

1. **Mothers' and Fathers' Industry-Level Exposure.** First, we measure direct exposure through parents' industries following McCaig et al. (2022):

$$\tau_j = \log(1 + \text{Column } 2_j) - \log(1 + \text{MFN}_j) \tag{9}$$

   where j denotes a parent's industry. A higher value of τ_j indicates a larger effective tariff cut for that industry. As we will show later, mothers and fathers typically work in different industries, allowing these separate measures for each parent to capture relative changes in labor market returns between mothers and fathers. Thus, these industry-level exposures proxy for mother's and father's wage changes due to tariff cuts.

2. **Province-Level Exposure.** Second, we construct a province-level shift-share measure to capture broader local economic effects:

$$\tau_p = \sum_j s_{jp}\tau_j$$

   where s_{jp} is industry j's employment share in province p from the 1999 Census. This measure reflects changes in local labor market conditions as well as non-labor income—such as rising returns to land holdings—stemming from overall economic expansion.

3. **Female-Specific Exposure.** Third, we create a female-specific exposure measure to capture changes in market returns to women's work:

$$\tau_p^w = \sum_j (f_j \times s_{jp})\tau_j$$

[3] 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. Most of these quotas had, however, been eliminated by the end of 2002 (STAR-Vietnam, 2003).

where f_j is industry j's female employment share from the 1999 Census, inspired by Autor et al. (2019). We use national (but industry-specific) female shares to avoid non-random sorting of female-intensive industries across provinces. This measure captures how broader changes in female labor demand at the provincial level, particularly in sectors such as textiles and garments that experienced significant tariff cuts, might affect perceived returns to daughters.

Binary Treatment Indicators. To avoid issues with continuous treatment variables in difference-in-differences settings (Callaway et al., 2024), our main analysis uses binary versions of these exposure measures. For individual-level exposure, we create indicators where $\mathbf{1}_m$ equals one if the mother's industry tariff experienced any tariff cut ($\tau_{j(m)} > 0$) and zero otherwise. Similarly, $\mathbf{1}_d$ equals one if the father's industry experienced a tariff cut and zero otherwise. For location-based measures, we define $\mathbf{1}_p$ as an indicator for provinces with above-median tariff cuts (τ_p) and $\mathbf{1}_p^w$ as an indicator for provinces with above-median female-specific tariff cuts (τ_p^w). In our robustness checks, we also use the continuous measures directly.

3.2 Household Data

The literature studying sex selection and fertility typically relies on decennial census data, which requires strong assumptions about the stability of parents' employment and location over long periods. We overcome this limitation using the Vietnam Household Living Standards Surveys (VHLSS) from 2002-2016, a biennial nationally representative survey conducted by Vietnam's General Statistics Office. The VHLSS offers large enough samples and two distinct advantages: it captures parents' industry and employment at the moment fertility decisions are made, rather than years before or after; and its frequent waves enable us to track responses to the trade shock with greater precision than decennial observations would allow.

Sample Construction. Our main sample uses repeated cross-sections of approximately 4,000 infants aged 0-1 years from each survey wave. For these infants, we observe their gender and their parents' contemporaneous industry codes and work hours. Although our first observation is from 2002, the survey's 12-month recall period means the 2002 data largely reflects pre-BTA conditions since the agreement took effect in December 2001. Additionally, since fertility decisions and sex selection are slow-moving outcomes, children born in the immediate post-BTA years (2002-2003) likely reflect conception decisions made before any substantial labor market changes could influence parental behavior.

A key measurement challenge in developing countries is identifying spouses and parent-child relationships within extended family households, where multiple families and generations often live together. Standard household surveys like the VHLSS typically classify household relationships relative to the household head, making it difficult to distinguish between children and their cousins or between parents and aunts/uncles. To overcome this challenge, we develop an algorithm that identifies both couple and parent-child relationships within extended families. We validate our algorithm using known parent-child relationships available in the 2014 and 2016 VHLSS waves, achieving 95% accuracy in identifying mothers and 91% accuracy for fathers. This methodological advance, detailed in Appendix A, allows us to link parents' industry exposure to their fertility outcomes.

For fertility analysis, we create a retrospective panel that tracks annual birth outcomes. Specifically, for each woman in a given survey year, we construct observations spanning from when she turned 20 until her survey year, with our key outcome being whether she gave birth in a given year. This approach allows us to control for mother-specific fertility preferences through mother-fixed effects.

Panel Sub-sample. To assess the robustness of our main results from the repeated cross-sections, we exploit the VHLSS's rotating panel feature to follow a subset of women aged 20-40 across three consecutive waves (2002-2004-2006). This small panel helps address potential bias from workers selectively switching industries in response to trade shocks. By fixing parents' industry affiliations before the policy change, we can isolate the effect of trade exposure from compositional changes driven by worker sorting. The panel allows us to control for unobservable factors like fertility and sex ratio preferences through individual fixed effects. However, identifying the impact on sex selection requires mothers to have multiple births across the panel period, with at least one birth before the policy change. Given this requirement and the short panel length, we rely on the cross-sectional variation for our main analysis.

One potential concern is multicollinearity among our treatment variables. In Figure 2, we report the correlation matrix across all binary treatment indicators. Reassuringly, we find a low correlation between mother and father treatments, with correlation coefficients ranging from 0.3 to 0.4 across survey years. Similarly, the correlation between parent-level treatments and location-level treatments remains low throughout the sample period. There is a moderate correlation of approximately 0.7 between the province-level shift-share ($\mathbf{1}_p$) and female-specific shift-share ($\mathbf{1}_p^w$), which warrants some caution but is not severe enough to prevent identification. In our analysis, we report the individual effects of each treatment

Figure 2: Correlations between Different Treatment Variables

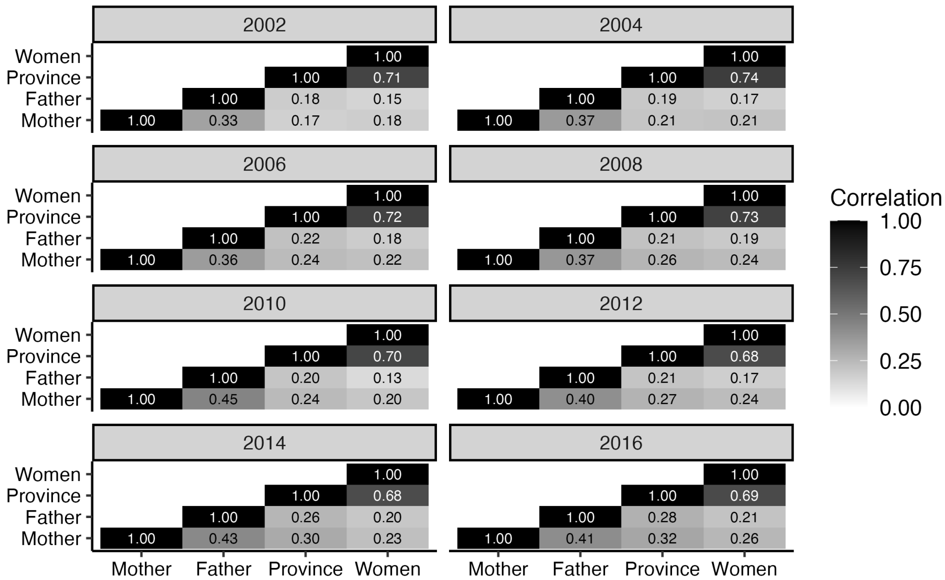

Data Source: Vietnam Household Living Standards Survey, 2002-2016 waves.
Notes: The correlation matrices display relationships between different measures of exposure to the BTA detailed in subsection 3.1. Mother stands for binary treatment $\mathbf{1}_m$, Father for $\mathbf{1}_d$, Province for $\mathbf{1}_p$ and Women for $\mathbf{1}_p^w$.

and gradually add different treatment combinations to assess sensitivity. Variance inflation factor (VIF) tests across all specifications remain below 3 (generally below 2), indicating that multicollinearity is not biasing our estimates.

4 Empirical Strategy

The model's predictions in Table 1 motivate our difference-in-differences specification:

$$\begin{aligned} y_{imdpt} = {} & \beta_1 \mathbf{1}_m \times \text{Post}_t + \beta_2 \mathbf{1}_d \times \text{Post}_t + \beta_3 \mathbf{1}_p \times \text{Post}_t + \beta_4 \mathbf{1}_p^w \times \text{Post}_t \\ & + \alpha_t + \delta_p + \gamma_{j(m)} + \gamma_{j(d)} + X'_{imdpt}\lambda + S'_p \times \alpha_t + \varepsilon_{imdpt} \end{aligned} \tag{10}$$

where the dependent variable y_{imdpt} represents outcomes for child i born to mother m and father d in province p in year t. For child gender analysis, y_{imdpt} is an indicator for male births. For the mother's labor supply analysis, y_{imdpt} measures the log of monthly work hours in the mother's main job. The indicator Post_t equals 1 for years after 2002. The specification includes province fixed effects δ_p, year fixed effects α_t, and industry fixed effects for both parents $\gamma_{j(m)}$ and $\gamma_{j(d)}$.

Our control variables X_{imdpt} include the mother's age, education, urban location, and minority status. The term S_p represents the pre-BTA employment shares in agriculture and manufacturing sectors in each province p. These employment shares are interacted with year dummy variables to account for different economic trends across provinces in which preexisting conditions might persist.

For fertility outcomes, we construct a retrospective panel that tracks mothers' birth histories over time. We modify the baseline specification:

$$\begin{aligned} \text{AnyBirth}_{mt} = \beta_1 \mathbf{1}_m \times \text{Post}_t + \beta_2 \mathbf{1}_d \times \text{Post}_t + \beta_3 \mathbf{1}_p \times \text{Post}_t + \beta_4 \mathbf{1}_p^w \times \text{Post}_t \\ + \alpha_t + \mu_m + X'_{mt}\lambda + S'_p \times \alpha_t + \varepsilon_{mt} \end{aligned} \tag{11}$$

where AnyBirth_{mt} is an indicator for whether mother m had any birth in the past year, and we substitute mother fixed effects μ_m for the province and industry fixed effects. This allows us to control for time-invariant fertility preferences.

We cluster standard errors at the province level across all specifications. We confirm that our main findings are robust to alternative clustering methods: clustering at the industry level or using two-way clustering (province and industry) yields similar significance levels. These robustness checks are discussed in Section 6.

Our use of repeated cross-sectional data introduces potential selection bias through endogenous sorting. If women with stronger son preferences systematically moved into industries with higher tariff cuts after the policy implementation, our cross-sectional estimates would capture both the causal impact of trade policy and this non-random selection, potentially inflating the estimated treatment effects.

To address this concern, we implement a panel approach that fixes each woman's industry assignment to her *pre-BTA* (2001) industry affiliation. By fixing treatment status before the policy shock, the design mirrors an intent-to-treat experiment and prevents endogenous post-2001 sorting from biasing the coefficients:

$$\begin{aligned} y_{iht} = \beta_1 \mathbf{1}_i \times \text{Post}_t + \beta_2 \mathbf{1}_h \times \text{Post}_t + \beta_3 \mathbf{1}_p \times \text{Post}_t + \beta_4 \mathbf{1}_p^w \times \text{Post}_t \\ + \alpha_t + \mu_i + X'_{iht}\lambda + S'_p \times \alpha_t + \varepsilon_{iht}, \end{aligned} \tag{12}$$

where $\mathbf{1}_i$ and $\mathbf{1}_h$ are industry-specific tariff treatments for woman i and her husband based on their 2001 (pre-BTA) industries, following McCaig and Pavcnik (2018).

For sex-selection and maternal-labor-supply analyses, the sample consists of women aged 20–40 who gave birth within the two years preceding the BTA (observed in the 2002 VHLSS) and at least once more in 2004 or 2006. For fertility, the panel is widened to all women aged

20–40, and the outcome is an indicator that the woman gave birth in the previous two years. Mother fixed effects μ_i absorb all time-invariant heterogeneity—such as underlying son preference or baseline fertility ideals—so identification relies solely on within-woman changes in outcomes following an exogenous change in her potential wage.

Since treatment status is locked to 2001 industries, post-BTA mobility cannot alter exposure. This panel specification thus resembles an intent-to-treat design that isolates the causal effect of trade-induced wage changes from endogenous sorting. If women with strong son preference were disproportionately drawn to high-cut sectors after 2001, a cross-sectional estimator would confound selection with treatment. The panel estimator eliminates that channel. If sorting were driving our results, we would expect smaller treatment effects in the panel approach compared to the cross-sectional estimates.

The predicted signs of the coefficients of interest β_1, β_2, β_3, and β_4 correspond to the comparative statics summarized in Table 1. Our model predicts $\beta_1 > 0$ for male births and labor supply and $\beta_1 < 0$ for fertility, as rising maternal wages intensify the childcare-work tradeoff. In contrast, increases in the father's wage or non-labor income should raise fertility ($\beta_2 > 0$ and $\beta_3 > 0$), with no impact on sex selection. Finally, higher relative returns to female work should reduce sex selection ($\beta_4 < 0$) and increase mother labor supply ($\beta_4 > 0$) but have no effect on fertility.

Our identification strategy leverages two sources of variation: cross-sectional differences in exposure to tariff cuts across industries and provinces, and temporal variation before versus after BTA implementation. The key identifying assumption for causal inference is that outcomes would have evolved similarly across industries and provinces with different exposure levels absent the BTA. For the panel analysis, we additionally assume that women's pre-BTA industry choices were not made in anticipation of the policy's effects.

A natural concern with our identification strategy is limited pre-policy data for assessing parallel trends. However, a key feature of our research context provides additional validation.

Fertility outcomes inherently respond with biological and decision-making lags. Births recorded immediately after BTA implementation (2002-2003) primarily reflect conception decisions made before labor market changes could substantially influence parental behavior. This biological delay effectively extends our pre-treatment period, allowing us to consider 2003 onwards as the true post-treatment period for sex selection and fertility outcomes.

Our subsequent event study analysis confirms this pattern—treatment effects emerge gradually following BTA implementation and only increase meaningfully after approximately a one-year lag, consistent with the expected delayed response in fertility outcomes. The lack of significant differences between treatment and control groups in the immediate post-policy years further validates our identification assumption.

Table 2: Effects of BTA Exposures on Sex Selection (Male Birth Probability)

	MaleBirth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$1_m\times$ Post	0.09***				0.09***	0.09***	0.09***	0.09***
	(0.02)				(0.02)	(0.02)	(0.02)	(0.03)
$1_d\times$ Post		0.04			0.007	0.010	0.010	0.007
		(0.02)			(0.02)	(0.03)	(0.03)	(0.03)
$1_p\times$ Post			-0.0004			-0.01	0.003	0.001
			(0.01)			(0.01)	(0.01)	(0.02)
$1_p^w\times$ Post				-0.007			-0.02	-0.02
				(0.01)			(0.01)	(0.01)
R^2	0.009	0.009	0.009	0.009	0.009	0.009	0.009	0.01
Observations	28,419	28,419	28,419	28,419	28,419	28,419	28,419	28,419
Dep. Var. Mean	0.52	0.52	0.52	0.52	0.52	0.52	0.52	0.52
Controls								✓

Notes: The dependent variable is an indicator equal to 1 if the child is male. Exposure variables $\mathbf{1}_m$, $\mathbf{1}_d$, $\mathbf{1}_p$, and $\mathbf{1}_p^w$ are defined in Section 3.1. All specifications include birthyear, province, and parents' 2-digit ISIC fixed effects. Controls in column (8) include mother's age, education, urban location, minority status, and pre-BTA provincial employment shares in agriculture and manufacturing sectors interacted with year dummies. Standard errors clustered at the province level are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5 Main Results

5.1 Effects on Sex Selection

The results presented in the Table 2 strongly support most predictions from our model as in column 1 of Table 1.

In columns (1)-(4), each key independent variable is tested individually. The most striking finding is the strongly significant positive coefficient (0.09) for the mother's exposure effect ($\mathbf{1}_m \times \text{Post}_t$) on male births, which remains consistently significant at the 1% level across all model specifications. This robust result aligns perfectly with the theoretical prediction ($\beta_1 > 0$), suggesting that as maternal wages increase, families are more likely to select male children. This supports the hypothesis that higher opportunity costs for mothers intensify the childcare-work tradeoff.

To interpret the magnitude of the maternal exposure effect (coefficient of 0.09), we need to account for the size of tariff reductions experienced by mothers in exposed industries. With an average tariff cut of approximately 8 percentage points conditional on exposure, the scaled effect would be $0.09 \times 0.08 = 0.0072$ (or a 0.72 percentage points increase in male birth probability). Starting from a baseline sex ratio of 107 males per 100 females (51.7% male births), this effect would increase the male birth percentage to 52.42%, corresponding to a sex ratio of approximately 110.2 males per 100 females. This magnitude aligns with Vietnam's demographic trends during the study period documented in Figure 1 and suggests that maternal exposure to trade liberalization made a meaningful contribution to the observed increase in male-biased births.

The effects of fathers' industry exposure ($\mathbf{1}_d \times \text{Post}_t$) show a different pattern, generally consistent with our model's predictions. As shown in column (2), it has a positive coefficient (0.04) but lacks statistical significance. When included in combined models (columns (5)-(8)), this coefficient diminishes substantially and remains statistically indistinguishable from zero, suggesting limited support for the hypothesis that paternal income influences sex selection decisions and consistent with our model prediction that $\beta_2 = 0$.

Regarding household non-labor income ($\mathbf{1}_p \times \text{Post}$), column (3) shows a near-zero and statistically insignificant effect (-0.0004) on male births. This null result persists across specifications, indicating that general household income does not play a significant role in sex selection, contrary to what might be expected if sex selection were purely a normal good, as argued by Almond et al. (2019).

The coefficient for relative returns to women work ($\mathbf{1}_p^w \times \text{Post}_t$) is negative (-0.007 in column (4), increasing to -0.02 in columns (7)-(8)) as predicted ($\beta_4 < 0$), though it fails to reach statistical significance. This directionally aligns with the hypothesis that improved economic opportunities for women could reduce son preference, but the evidence is not strong enough to make definitive claims. Furthermore, when considered alongside the positive coefficient on maternal exposure ($\mathbf{1}_m \times \text{Post}_t$), these results offer limited support for the relative returns mechanism proposed by (Qian, 2008; Rosenzweig & Schultz, 1982).

Although the female-specific exposure measure aligns directionally with our theoretical predictions, the effects are smaller and statistically imprecise. This result likely reflects that province-level female labor demand is a noisy proxy for the opportunity cost of sex selection, p_q in our model. By contrast, maternal industry exposure provides a cleaner measure of changes in individual opportunity costs and relative returns.

Our findings are comparable in magnitude to results from other studies when properly scaled. For instance, Anukriti (2018) found that an Indian program addressing the fertility-sex ratio trade-off resulted in a 1-2.3% increase in the probability of first births being male.

Figure 3: Event Study Estimates of BTA Exposure Effects on Male Births

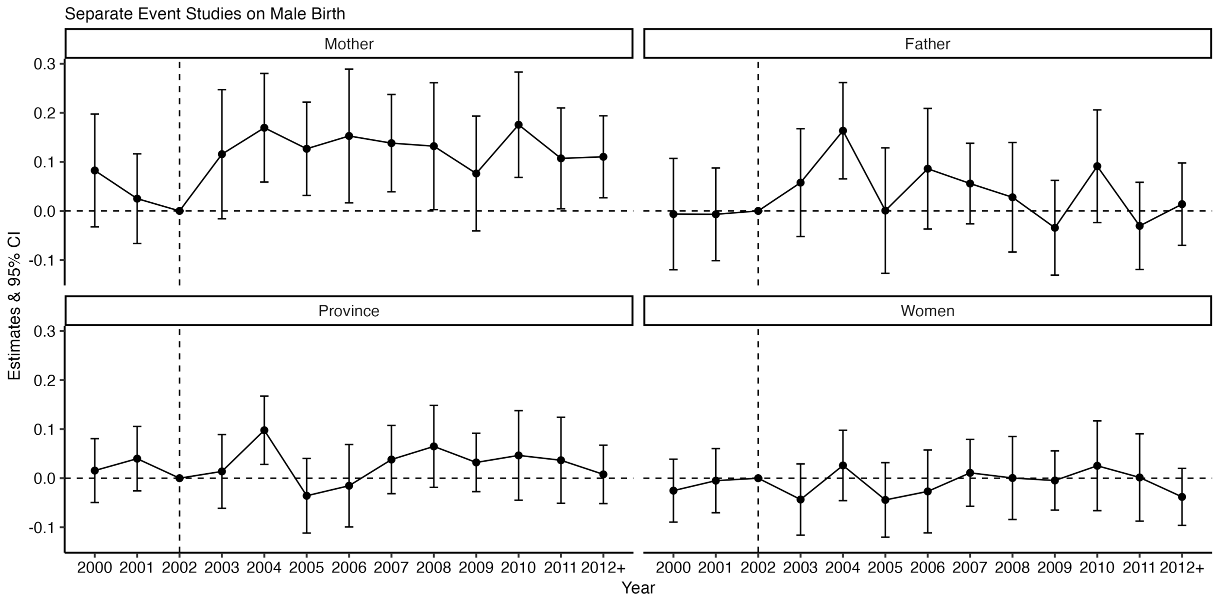

Notes: This figure displays event study estimates examining the effects of the BTA on male births across four dimensions of trade exposure. Each panel represents a separate regression where the dependent variable is an indicator for male birth, and the independent variables are year dummies interacted with the respective exposure measure, controlling for province, year, and parental industry fixed effects. The panels show coefficients for mothers in exposed industries (top left), fathers in exposed industries (top right), provinces with above-median exposure (bottom left), and provinces with above-median female-specific exposure (bottom right). The vertical dashed line indicates the implementation of the BTA in late 2001. Error bars represent 95% confidence intervals with standard errors clustered at the province level.

Similarly, Almond et al. (2019) discovered that land reform in China between 1978 and 1984 led to a 3 percentage point (5.6%) increase in the fraction of males following the birth of a first girl, from a baseline sex ratio of about 112. While these studies focused on specific birth parities and sibling compositions, our estimate of a 0.72 percentage point increase in male birth probability (translating to a sex ratio increase from 107 to 110.2) represents an effect across all birth orders, suggesting a substantial population-wide impact of maternal economic opportunities on sex selection behavior.

Next, we explore the event-study counterpart of regression specification (10) to examine both the pre-trends and dynamic effects of the BTA exposures. Figure 3 presents event study estimates on male births across four dimensions of trade exposure. Each panel represents a specific channel through which trade liberalization might affect sex selection. Thus, we examine in this figure the individual effect of each exposure without conditioning on the others.

The top-left panel displays estimates for mothers employed in industries experiencing tariff cuts post-BTA. Mothers in affected industries show higher male birth probabilities starting in 2004, approximately two years after the BTA implementation. The effect stabilizes at around 10 percentage points in subsequent years, with most post-treatment coefficients statistically significant as indicated by confidence intervals that do not cross zero. This delayed pattern aligns with the biological timeframe required for conception and birth following economic policy changes.

The top-right panel focusing on fathers' industry exposure reveals mixed effects on male births, with mostly null effects post-BTA. This pattern for fathers shows much more volatility than for mothers, which explains the null result for paternal exposure found in Table 2.

The lower panels examine geographic dimensions of trade exposure. Neither province-level nor female-specific provincial exposure seem to have consistent significant impacts individually on sex selection, with coefficients fluctuating around zero and confidence intervals frequently crossing the zero line.

Importantly, all four panels demonstrate minimal pre-trend differences both before the BTA implementation (marked by the vertical dashed line) and a couple of years immediately following it. These results support the assumption of parallel trends and take advantage of slow-moving adjustment due to the biological nature of fertility. These findings suggest that household-specific exposure, particularly maternal employment in industries benefiting from trade liberalization, plays a more direct role in sex selection decisions than broader geographical exposures.

Figure 4 presents a more comprehensive analysis of the event study where all the main regressors of Equation (10) are included simultaneously, along with the full set of controls. The figure provides a more nuanced understanding of how different dimensions of trade exposure affect male births when accounting for potential simultaneous changes in other factors due to the trade cost.

The most striking pattern appears in the maternal exposure panel (top), which shows consistently positive and statistically significant effects on male births across almost all post-BTA years. After a brief adjustment period in 2002, the coefficient stabilizes around 0.10, indicating approximately a 10 percentage point increase in the probability of male births for mothers employed in industries benefiting from tariff reductions. This robust pattern persists and remains stable throughout the study period.

In contrast, the remaining exposure measures show minimal effects. Father's industry exposure exhibits brief deviations from zero in 2004 but no consistent pattern. Both province-level measures remain statistically insignificant.

Across all panels, pre-trend differences before the 2001 BTA implementation remain

Figure 4: Event Study Estimates of BTA Exposures on Male Births - Full Model

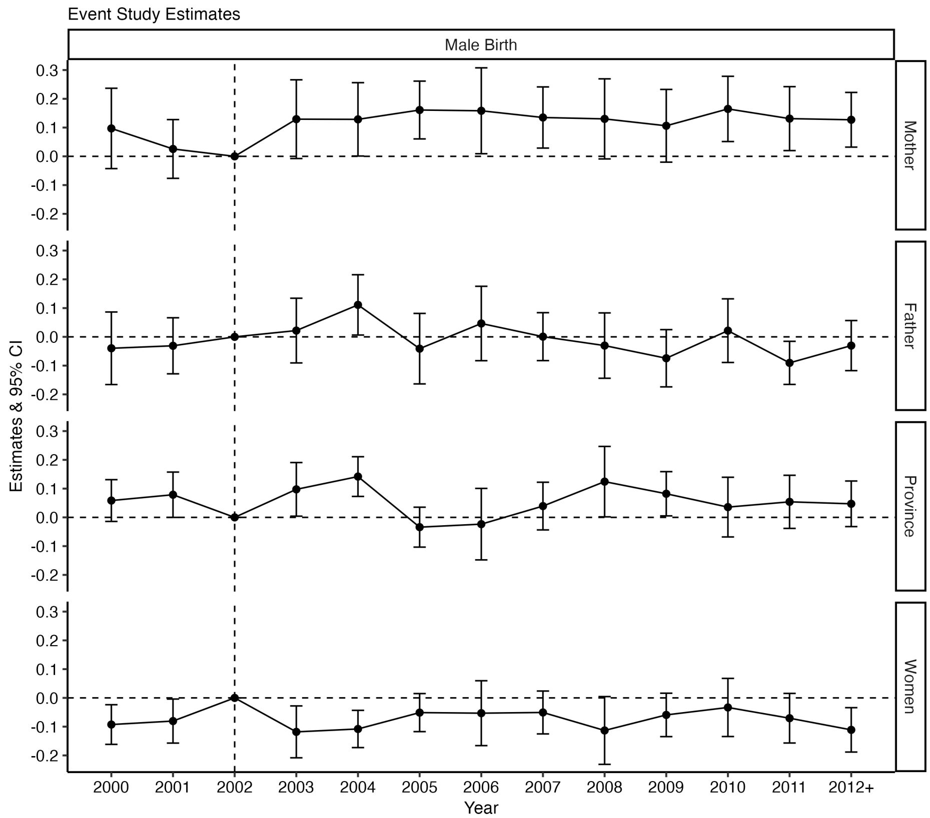

Notes: This figure displays event study estimates from the full model examining the effects of BTA exposures on male births. Each panel shows coefficients for a specific dimension of trade exposure: mothers in exposed industries (top), fathers in exposed industries (second row), provinces with above-median exposure (third row), and provinces with above-median female-specific exposure (bottom row). All dimensions are included simultaneously in a single regression with the full set of controls. The vertical dashed line indicates the implementation of the BTA in late 2001. Error bars represent 95% confidence intervals with standard errors clustered at the province level.

minimal, supporting the validity of our difference-in-differences approach. The contrast between the consistent maternal effects and the more muted patterns for other dimensions suggests that household decisions about sex selection respond most strongly to improvements in mothers' labor market opportunities, consistent with our model's predictions.

We examine how the effect of maternal industry exposure on sex selection varies across subgroups. For each subgroup, we estimate our baseline specification (10) separately. Figure 5 reveals several important patterns of heterogeneity.

First, maternal industry tariff cuts are associated with increased male birth probability in both urban and rural areas, with point estimates of approximately 0.08 and 0.10, respectively. While the rural effect is slightly larger, the difference is not statistically significant. This

Figure 5: Heterogeneous Effects of Maternal Industry Exposure on Male Birth Probability

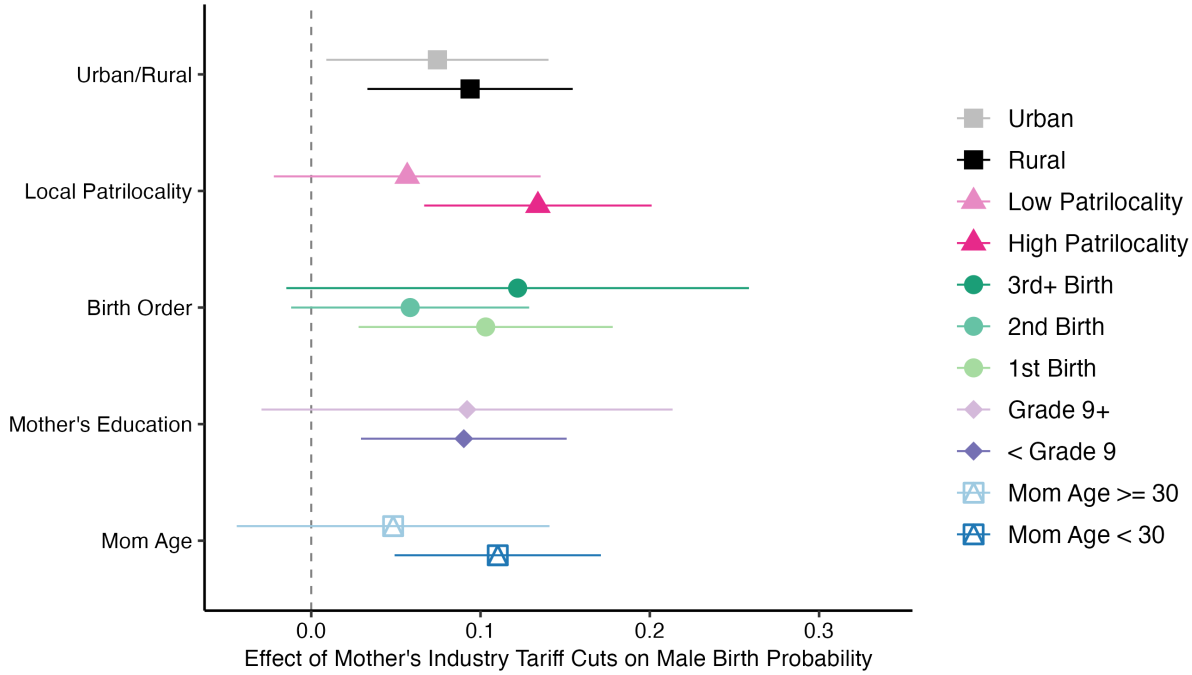

Notes: This figure shows the heterogeneous effects of maternal industry tariff cuts across various subgroups, including residence, patrilocality, birth order, education, and age. Each point estimate and its corresponding 95% confidence interval represent the coefficient on maternal industry exposure (β_1) from separate regressions of equation (10) estimated within the indicated subgroup. The subgroups include urban/rural residence status, high/low patrilocality areas, birth order (1st, 2nd, and 3rd+ births), maternal education level (less than grade 9 or grade 9+), and maternal age (below 30 or 30 and above). The dashed vertical line at zero represents no effect. All regressions include the full set of controls and fixed effects as specified in equation (10).

suggests that trade liberalization may have a more pronounced maternal impact in less economically developed regions, where son preference tends to be stronger.

Second, the effects are roughly twice as large in areas with high patrilocality compared to those with low patrilocality. We measure patrilocality using the 1999 Census, calculating the share of married couples living with the husband's parents in each province, then classifying provinces above the median as high patrilocality. This residence pattern traditionally reinforces son preference by ensuring property inheritance through male lines (den Boer and Hudson, 2017; Ebenstein, 2014). The stronger effect in patrilocal areas likely reflects the interaction between new economic incentives and deep-rooted cultural norms that favor sons.

Third, birth order matters. We detect a large, statistically significant increase in male-birth probability at first parity, no discernible effect at second parity, and a renewed, though smaller, increase at third and higher parities.[4] Our estimates align with evidence from the 2014 Vietnamese Intercensal Survey reported by Becquet and Guilmoto (2018), who show that the sex ratio at birth is most skewed at first and third-or-higher parities, while remaining near the biological norm at parity 2.

Two mechanisms consistent with our model likely drive this pattern. First, mothers face a sharp work-child-care trade-off when deciding on their first child and may use sex selection to satisfy the cultural expectation of producing a son before returning to the labor market. Second, if the first two children are daughters, son preference intensifies at higher parities, reviving selection at births 3+.

Fourth, maternal education does not materially alter the trade-induced shift in sex selection. The point estimate is roughly 0.10 for both mothers who completed grade 9 or higher and those with less schooling, although only the latter estimate is statistically precise. Thus, the economic incentives generated by trade liberalization appear to influence fertility decisions across the education distribution, with the clearest evidence among less-educated women.

Fifth, age matters. Mothers under 30 experience roughly double the effect (≈ 0.12) observed for mothers 30 and older (≈ 0.06). Younger women may feel greater pressure to secure a son early—both to satisfy cultural norms and to capitalize on a longer remaining work horizon—mirroring the strong selection we found at first parity.

Although the point estimates line up with where we expect son preference to be most intense, formal tests cannot reject equality of the effects across most subgroups. In other words, increases in maternal trade exposure appear to raise sex selection in much the same way for all groups, just as our simple model predicts when the wage mechanism operates independently of the underlying strength of inherent son preference.

5.2 Effects on Mom's Labor Supply

Each column in Table 3 reports a difference-in-differences coefficient on an interaction between the post-BTA indicator, *Post*, and one of the four binary exposure measures defined in Section 3.1. The dependent variable is log(MomWorkHrs) and all specifications absorb survey-year, province, and parents' 2-digit ISIC fixed effects. Column (8) additionally includes the same set of controls as Column (8) of Table 2.[5]

[4] Births of order 3 or higher are uncommon in our data; Vietnam's total fertility rate had already fallen to 2.01 by 2000.

[5] The sample of mothers here is smaller than that of children in the sex selection analysis because of maternity leave. Since the 2010 wave, VHLSS has asked about the number of work hours within the last 30 days. If a mother is on maternity, she will drop out of the sample, although her child appears in the infant sample.

Table 3: Effects of BTA Exposures on Mom's Labor Supply

	log(MomWorkHrs)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$1_m\times$ Post	0.23***				0.23***	0.22***	0.21***	0.19***
	(0.04)				(0.05)	(0.04)	(0.04)	(0.04)
$1_d\times$ Post		0.07**			-0.009	-0.02	-0.02	-0.01
		(0.03)			(0.03)	(0.04)	(0.04)	(0.03)
$1_p\times$ Post			0.06			0.04	-0.07	-0.09
			(0.05)			(0.05)	(0.07)	(0.07)
$1_p^w\times$ Post				0.13**			0.16**	0.14**
				(0.05)			(0.07)	(0.07)
R^2	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.30
Observations	26,147	26,147	26,147	26,147	26,147	26,147	26,147	26,147
Dep. Var. Mean	4.7	4.7	4.7	4.7	4.7	4.7	4.7	4.7
Controls								✓

Notes: The dependent variable is the log of the mother's monthly hours worked in her main job. Controls in column (8) match those in Table 2. All specifications include survey-year, province, and parents' 2-digit ISIC fixed effects. Standard errors, clustered at the province level, are reported in parentheses.

Column (1) shows that mothers employed in tariff-cut industries ($\mathbf{1}_m = 1$) work 0.23 log-points more after the BTA, significant at the 1% level. Exponentiating gives a 26% increase in monthly hours. With the sample mean of 110 hours, this translates into roughly +29 hours per month—about one additional full-time work-week.

The estimate remains virtually unchanged when other exposure dummies are added (Columns 5–7) and only attenuates slightly—to 0.19 log-points (21%)—once the full control set is included (Column 8).

The father's industry exposure ($\mathbf{1}_d$) is positive and significant when entered alone (Column 2, +7%) but becomes small and insignificant once maternal exposure or controls are added. Provincial exposure capturing general economic expansion ($\mathbf{1}_p$) is never significant, while the female-specific provincial measure ($\mathbf{1}_p^w$) is positive and significant. Entered alone (column 4), the women-intensive province measure raises work hours by 0.13 log points (14%). Importantly, it remains significant in columns (7) and (8) at 0.16–0.14 log points (15–17%, or roughly +16–18 hours per month). Thus, even after we net out the mother's own industry shock, mothers living in provinces where tariff cuts disproportionately boosted

Table 4: Effects of BTA Exposures on Fertility (Probability of Giving Birth)

	AnyBirth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$1_m\times$ Post	-0.004***				-0.006***	-0.006***	-0.006***	-0.006***
	(0.001)				(0.001)	(0.001)	(0.001)	(0.001)
$1_d\times$ Post		0.003**			0.006***	0.005***	0.005***	0.005***
		(0.002)			(0.001)	(0.001)	(0.001)	(0.001)
$1_p\times$ Post			0.001			0.002	0.002	0.002
			(0.002)			(0.002)	(0.002)	(0.002)
$1_p^w\times$ Post				0.0003			-0.001	-0.002
				(0.002)			(0.002)	(0.002)
R^2	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.29
Observations	2,011,978	2,011,992	2,012,007	2,012,007	2,011,963	2,011,963	2,011,963	2,011,963
Dep. Var. Mean	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
Controls								✓

Notes: The dependent variable is an indicator that the mother had any birth in the past 12 months. Controls in column (8) match those in Table 2. All specifications include birth year and mother fixed effects. Standard errors, clustered at the province level, are reported in parentheses.

female employment still expand their hours.

Across eight increasingly demanding specifications, the mother-industry interaction remains large and precisely estimated, while alternative channels fade once maternal exposure is taken into account. Therefore, the evidence strongly supports the paper's central mechanism: tariff-induced improvements in women's labor-market opportunities substantially increase their work hours, tightening the work - childcare trade-off that later manifests itself in fertility and sex selection behavior.

5.3 Effects on Fertility

The repeated cross-section estimates in Table 4 show a clear divergence between maternal and paternal exposure to tariff cuts in their effects on annual fertility.

Across all eight specifications, exposure through the mother's industry lowers the probability that she gives birth in a given year by about 0.6 percentage points (columns 5–8). Because the dependent variable mean is only 2 percent, this translates into a sizeable 30 percent reduction in the likelihood of childbearing in any given year. The coefficient is remarkably stable—dropping only slightly from -0.004 when entered alone to -0.006 once the full set of controls is included—and is precisely estimated with a standard error of 0.001.

In sharp contrast, fathers working in liberalized industries become *more* likely to have a child. The fully controlled estimate of 0.005 implies a 25 percent increase in annual birth probability relative to the mean. As soon as the maternal dummy is included, the paternal effect actually *strengthens*, suggesting that the two channels operate independently rather than offsetting one another. Taken together, the coefficients imply that in couples where both parents are exposed, the net effect on fertility is close to zero, but when only the mother is exposed fertility falls sharply—consistent with the model's emphasis on the mother's time constraint.

Neither the general province-level exposure indicator ($\mathbf{1}_p \times Post$) nor the female-specific provincial measure ($\mathbf{1}_p^w \times Post$) is statistically significant in any specification. This absence of local-market effects echoes the labor-supply results and reinforces the interpretation that it is the individual mother's opportunity cost—rather than aggregate demand shocks—that drives the fertility response.

Taken together, the empirical evidence lines up closely with the twelve comparative-static tests summarised in Table 1. We confirm three of the four predicted signs for the sex ratio, two of the four for maternal labour supply, and three of the four for fertility. The few departures involve coefficients that are qualitatively consistent yet statistically insignificant, so the overall pattern strongly supports the work–child-care mechanism at the heart of our model.

5.4 Panel Evidence

We estimate the panel specification in (12), which fixes each woman to her *pre-BTA* (2001) industry and absorbs individual fixed effects. If women with stronger son preference systematically moved into high-tariff-cut sectors after BTA, the cross-sectional estimates would blend causal effects with selection and should, therefore, exceed the panel estimates. The data tells a different story.

Columns (1)–(2) of Table 5 show that mothers whose initial industries were treated are 0.32–0.34 log-points more likely to bear a son after the BTA—roughly a 60 percent increase over the sample mean of 0.53. These effects are larger than, not smaller than, their cross-sectional counterparts, implying that endogenous sorting is unlikely to inflate our main findings.

Column (3) tests a key prediction of our model: once a family achieves their desired goal of having a son, their sex selection behavior should become less responsive to economic incentives. To examine this mechanism, we add a control for whether the mother already had

Table 5: Panel Estimates with Pre-BTA Industry Assignment

	MaleBirth			log(WorkHrs)		AnyBirth	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$1_m \times$ Post	0.32***	0.34***	0.18	-0.05	-0.11	0.05	0.05
	(0.12)	(0.12)	(0.11)	(0.15)	(0.17)	(0.03)	(0.04)
$1_d \times$ Post	0.03	0.01	0.05	-0.15	-0.18	0.006	0.004
	(0.16)	(0.17)	(0.15)	(0.17)	(0.16)	(0.03)	(0.03)
$1_p \times$ Post	0.15*	0.14*	0.12	0.12	0.10	-0.02	-0.02
	(0.08)	(0.08)	(0.07)	(0.16)	(0.17)	(0.02)	(0.02)
$1_p^w \times$ Post	-0.13*	-0.18**	-0.09	-0.05	-0.18	0.03	0.02
	(0.08)	(0.07)	(0.07)	(0.16)	(0.18)	(0.02)	(0.02)
Had a son before			-0.56***				
			(0.06)				
R^2	0.82	0.84	0.86	0.73	0.74	0.43	0.44
Observations	2,615	2,615	2,615	2,533	2,533	11,164	11,164
Dep. Var. Mean	0.53	0.53	0.53	4.7	4.7	0.10	0.10
Individual FE	✓	✓	✓	✓	✓	✓	✓
Birth Year FE	✓	✓	✓				
Survey Year FE	✓	✓	✓	✓	✓	✓	✓
Controls		✓	✓		✓		✓

Notes: The table reports coefficients from Eq. (12) estimated on the 2002–2006 VHLSS panel. Standard errors, clustered at the province level, are in parentheses.

a son before the current birth. The coefficient on maternal exposure falls substantially to 0.18 and loses statistical significance, while the "Had a son before" dummy is large, negative, and highly significant at -0.56. This pattern suggests that prior son status is a powerful predictor of subsequent sex selection, consistent with our model's premise that families primarily desire at least one son rather than maximizing the number of sons.

The attenuation of the maternal exposure effect when controlling for prior son status indicates that the impact of the BTA on sex selection operates primarily through families still seeking to have their first son, consistent with the heterogeneous result reported in Figure 5. This conditioning effect provides additional evidence for our proposed mechanism and rules out alternative explanations that would predict uniform effects regardless of prior family composition.

For maternal labor supply (Columns 4-5), the panel estimates show small, negative, and statistically insignificant effects of maternal industry exposure (-5% to -11%). This result differs from our cross-sectional findings of increased work hours. However, this difference is consistent with life-cycle labor supply patterns for our panel sample of women who have given birth twice in a relatively short period. These mothers likely experience temporary workforce detachment due to childcare responsibilities, especially during the early years after childbirth. The absence of immediate labor supply increases does not contradict our mechanism but rather reflects the short-term adjustment costs mothers face when balancing work and childcare. The longer-term labor supply effects captured in our cross-sectional analysis likely emerge as children age and mothers reintegrate into the workforce.

For fertility (Columns 6-7), we find no statistically significant effect of maternal exposure on birth probability within our short 2002-2006 panel window. Given that our cross-sectional analysis spanning fifteen years shows a clear negative fertility response, this null result likely reflects the limited time horizon rather than contradicting our proposed mechanism. Fertility decisions typically involve long-term planning and adjustment, which may not be fully captured in our short panel. Additionally, the panel sample consists of women who had at least one birth during the study period, potentially selecting for those with higher fertility preferences.

Taken together, the panel evidence reinforces our main conclusion: improved labor market opportunities for women in industries benefiting from trade liberalization substantially increase male-biased births, particularly among families still seeking their first son. The results for labor supply and fertility, while less precisely estimated in the panel, remain consistent with the slower adjustment process we would expect for these outcomes given maternal childcare responsibilities in the short run.

6 Robustness Checks and Discussion

In this section, we systematically assess the robustness of our main results to various estimation choices, including alternative clustering strategies, additional controls for region-specific factors, survey weighting, nonlinear specifications, and continuous treatment exposures. Overall, our core findings prove robust across these exercises. Maternal industry exposure consistently increases sex selection, raises maternal labor supply, and reduces fertility, regardless of estimation method or clustering choice. While alternative clustering approaches and unweighted regressions slightly attenuate significance levels, the qualitative and quantitative interpretations remain unchanged. Introducing continuous treatment measures further enhances clarity and confirms our baseline results. These robustness checks thus reinforce our conclusion that the maternal opportunity-cost channel, rather than alternative mechanisms such as income effect, relative returns, or bargaining effects, is the primary driver behind observed demographic adjustments following the BTA.

6.1 Alternative Clustering Approaches

A natural concern with inference in our setting is potential correlation in shocks and behavioral responses within provinces or industries. Our baseline estimates cluster standard errors by province (61 clusters), capturing geographic correlation arising from local policies or economic conditions. However, with mothers distributed across 59 distinct industries, industry-level clustering might also be relevant if shocks are industry-specific.

Table A5 reports standard errors clustered by industries. Clustering by industry slightly increases standard errors across all outcomes relative to province-level clustering, modestly reducing statistical significance. For instance, in the maternal labor supply regression, standard errors rise from approximately 0.04 (province) to 0.07–0.08 (industry), shifting significance from the 1 to the 5 percent level. Similarly, for fertility, maternal exposure remains significant but moves from the 1 percent to the 5 percent level as standard errors triple from province to industry clustering. Nonetheless, these modest changes do not alter our main conclusions: maternal industry exposure consistently predicts higher sex selection, increased maternal labor supply, and reduced fertility.

Given that provinces and industries have similar numbers of clusters, neither approach is clearly more conservative based on cluster count alone. To ensure robust inference that accommodates potential correlations simultaneously within provinces and industries, we also employ two-way clustering (Table A6). As theory predicts, two-way standard errors fall between the single-dimension province-only and industry-only values, and maternal exposure effects retain their signs, magnitudes, and significance across outcomes. This reinforces that our main findings are robust to alternative clustering strategies. We maintain province-level clustering as our baseline, as it best aligns with our geographically-driven identification strategy and captures key local correlations.

6.2 Region-Year Fixed Effects

A potential concern is that regional infrastructure improvements could independently affect access to sex-selective technologies or childcare services, confounding our estimates of trade exposure effects. Vietnam's infrastructure development typically follows its eight administrative regions (Northeast, Northwest, Red River Delta, North Central Coast, South Central

Coast, Central Highlands, Southeast, and Mekong River Delta). To address this possibility, we incorporate region-by-year fixed effects into our baseline specifications to absorb any time-varying regional confounders.

Table A7 shows that including these region-by-year fixed effects leaves our main results essentially unchanged. For sex selection, the maternal exposure coefficient remains stable at 0.09, highly significant at the 1 percent level, virtually identical to our baseline estimate. The other exposure measures remain small or statistically insignificant, aligning with our prior conclusions.

Regarding maternal labor supply, the estimated impact of maternal industry exposure remains consistent at approximately 0.18 log-points, translating into about a 20 percent increase in monthly hours, again statistically significant at the 1 percent level. Province-level exposure variables lose significance and become smaller in magnitude, reinforcing that maternal own-industry exposure—not regional shocks—drives labor market responses.

Finally, for fertility, maternal exposure continues to decrease the annual probability of giving birth by around 0.6 percentage points (from a mean of 2 percent), which is significant at the 1 percent level. The paternal coefficient remains positive and highly significant, closely matching earlier estimates, and the provincial measures remain negligible and insignificant.

In sum, accounting for regional time-varying infrastructure improvements leaves the core maternal exposure findings intact across sex selection, maternal work hours, and fertility. These robustness checks further strengthen our confidence in the main results, demonstrating that they are not driven by omitted region-specific shocks.

6.3 Survey Weighting

Because VHLSS oversamples certain subpopulations, our baseline specifications use survey weights for representativeness. Table A8 repeats analyses without weights, treating observations equally. Unweighted estimates closely mirror weighted results: maternal exposure remains statistically significant across sex ratio (0.06–0.07), maternal hours (around 20 percent increase), and fertility (-0.6 to -0.7 percentage points), confirming robustness.

This robustness is reassuring and unsurprising. Although unweighted estimates can sometimes outperform weighted estimates (Solon et al., 2015), the effects of trade exposure on sex selection appear relatively stable across subgroups (Figure 5). Thus, we maintain the use of sampling weights in our baseline specifications since they ensure our estimates are representative of the population-level impacts of trade liberalization.

6.4 Nonlinear (Logit) Estimation

Given our binary outcome (sex selection), we check robustness to nonlinear specification using a fixed-effects logit model, which is more computationally feasible than probit given numerous fixed effects. Table A9 reports the logit estimates. The results clearly reinforce our baseline findings. The maternal exposure coefficient remains strongly positive and highly significant at 0.4, with virtually no change upon the addition of controls. Given the mean male birth probability of approximately 50 percent, this translates into an odds ratio of about $e^{0.4} \approx 1.49$, indicating roughly a 49 percent increase in the odds of a male birth relative to the unexposed group.

The coefficients on paternal and province-level exposures remain small and statistically insignificant across both specifications, confirming again the centrality of the mother's own industry exposure in driving sex-selection responses. The consistency between logit and linear probability model results further underscores that our main estimates are not driven by model specification choices.

6.5 Continuous Treatment Exposures

Our main analysis used binary indicators to measure exposure to tariff cuts, categorizing industries simply as either "exposed" or "unexposed." While intuitive, this approach ignores potentially meaningful variation in exposure intensity within the treated group. Continuous measures of tariff reductions (τ variables in Section 3.1) can leverage this richer variation, enhancing statistical power and interpretability. In particular, continuous variables allow us to quantify the incremental effects of larger tariff cuts on behavioral responses, thus providing more nuanced insights into the underlying mechanisms.

Table A10 presents estimates based on continuous tariff-cut measures. The results strongly confirm and refine our main findings. For the sex ratio at birth, the maternal industry tariff cut ($\tau_m \times$ Post) coefficient is positive (0.26) and significant at the 5 percent level. This indicates that a one-unit increase in maternal industry tariff exposure (approximately a 100-percentage-point tariff reduction) would raise the likelihood of a male birth by roughly 26 percentage points. Given that the median tariff cut is around 7 percentage points, this translates into about a 1.8 percentage point increase in male births (around 3.5 percent relative to the mean). Other channels—paternal and province-level tariff cuts—remain small and insignificant, consistent with our earlier results.

Regarding maternal labor supply, maternal tariff exposure shows a clear and significant positive effect (0.29 log-points, significant at the 5 percent level). A median tariff cut of 7 percentage points would therefore predict about a 2 percent ($0.29 \times 0.07 \approx 0.02$) increase in monthly hours worked, reinforcing that higher tariff exposure intensifies the maternal work–childcare trade-off. Female-specific provincial tariff exposure also positively and significantly affects labor supply, while general province exposure negatively influences work hours, aligning closely with the binary results and our model's predictions.

For fertility, maternal tariff exposure significantly decreases the annual birth probability by -0.01, which implies a roughly 0.07 percentage-point reduction at median tariff cuts (about 3.5 percent of the mean fertility rate). Paternal tariff cuts positively influence fertility, consistent with the binary analysis, while the general provincial tariff exposure now also significantly increases fertility. This detailed analysis underscores the distinct roles that maternal and paternal channels play in shaping fertility outcomes.

In sum, shifting to continuous tariff exposure measures maintains—and in some cases clarifies—the main conclusions of our baseline analysis. The continuous specification leverages richer variation and further supports the robustness of the core maternal industry exposure effects on sex selection, maternal labor supply, and fertility.

6.6 Alternative Mechanisms

Our analysis strongly supports the maternal opportunity-cost mechanism derived from our quantity-quality model, and our robustness exercises further reinforce this conclusion. To briefly restate the logic, the model predicts that trade liberalization improves women's labor market returns, raising the maternal opportunity cost of childcare and intensifying the trade-off between work and child-rearing. Consequently, we expect maternal industry exposure to increase sex selection, raise maternal labor supply, and reduce fertility. By contrast, paternal exposure or general household income shocks should primarily influence fertility decisions rather than sex selection.

Our empirical results consistently align with these theoretical predictions across various robustness checks. These persistent patterns differentiate our maternal opportunity-cost mechanism from competing explanations that emphasize household income effects or changes in relative returns to daughters.

Moreover, our findings also present challenges for bargaining-based explanations. A bargaining model might replicate our results if mothers have a stronger preference for sons and desire fewer children compared to fathers. However, such a model also predicts that higher paternal wages should reduce sex selection by shifting bargaining power away from mothers, a prediction inconsistent with our data.

To illustrate this formally, we extend our baseline model in Appendix B by allowing bargaining weights to depend on relative wages ($\lambda(w_m/w_d)$) and permitting gender-specific preferences for both fertility ($\alpha_{nm} \neq \alpha_{nd}$) and sons ($\alpha_{qm} \neq \alpha_{qd}$). When preferences are identical across spouses, the predictions mirror those of our baseline model, driven primarily by the asymmetric childcare burden. Differences arise only when mothers have a stronger relative preference for sons ($\alpha_{qm}/\alpha_{nm} > \alpha_{qd}/\alpha_{nd}$) and desire fewer children ($\alpha_{nm} < \alpha_{nd}$) than fathers. As summarized in Panel C of Table A4, this bargaining extension can match most of our key empirical results—maternal wage increases leading to more sex selection and reduced fertility, and paternal wage increases raising fertility. However, it simultaneously predicts that higher paternal wages should lower sex selection by shifting bargaining power away from mothers—a prediction clearly rejected by our empirical evidence.

7 Conclusion

Using Vietnam's 2001 Bilateral Trade Agreement as a natural experiment, we find that mothers who are more exposed to tariff reductions through their industries have fewer children overall but are more likely to have sons when they do give birth. These effects appear to operate through increased maternal work hours and thus higher opportunity costs of childcare. Our theoretical framework reconciles competing explanations for skewed sex ratios in the literature, while our use of household data allows us to track parental employment at the time of birth, addressing limitations of previous work relying on less frequent censuses.

The results strongly support our model's emphasis on maternal time costs as the key mechanism linking trade liberalization to demographic outcomes. Women whose industries faced tariff cuts were 9 percentage points more likely to have male children, worked 23% more hours per month, and had a 0.6 percentage point lower annual birth probability. These effects stem from maternal exposure rather than paternal exposure or local economic changes, contradicting income-based or relative-returns explanations. Furthermore, panel analysis with individual fixed effects demonstrates that once a family achieves their desired goal of having a son, their sex selection behavior becomes substantially less responsive to economic incentives.

Our findings highlight how trade liberalization can create economic opportunities for women while exacerbating gender imbalances in son-preferring societies. This result suggests that policies aimed at promoting female labor force participation in developing economies should be coupled with efforts to address underlying cultural preferences for sons and asymmetric childcare responsibilities.

Although our analysis focuses on Vietnam's 2001 BTA, this mechanism likely applies to other developing economies experiencing trade integration alongside son preference and gendered childcare. Countries like Bangladesh, Indonesia, and parts of sub-Saharan Africa combine these elements in varying degrees. Our findings suggest that without complementary interventions addressing maternal time constraints—such as childcare subsidies—trade liberalization could exacerbate demographic imbalances. Future research should examine how institutional and cultural contexts modify the relationship between maternal opportunity costs and sex selection.

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., Bhalotra, S., & Tam, E. H. F. (2022). On the Quantity and Quality of Girls: Fertility, Parental Investments and Mortality. *The Economic Journal*, *132*(641), 1–36.

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. (1960). An Economic Analysis of Fertility. In *Demographic and Economic Change in Developed Countries* (pp. 209–240). Columbia University Press. Retrieved April 16, 2025, from https://www.nber.org/books-and-chapters/demographic-and-economic-change-developed-countries/economic-analysis-fertility

Becquet, V., & Guilmoto, C. Z. (2018). Sex imbalance at birth in vietnam: Rapid increase followed by stabilization. *Population*, *73*(3), 519–544.

Callaway, B., Goodman-Bacon, A., & Sant'Anna, P. H. C. (2024). Difference-in-differences with a Continuous Treatment. *National Bureau of Economic Research*.

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.

den Boer, A., & Hudson, V. (2017). Patrilineality, Son Preference, and Sex Selection in South Korea and Vietnam. *Population and Development Review*, *43*(1), 119–147.

Do, Q. T., & Phung, T. D. (2010). The importance of being wanted. *American Economic Journal: Applied Economics*, *2*(4), 236–253.

Doepke, M., Hannusch, A., Kindermann, F., & Tertilt, M. (2023). The economics of fertility: A new era. In *Handbook of the Economics of the Family* (pp. 151–254, Vol. 1). Elsevier.

Ebenstein, A. (2014). Patrilocality and Missing Women. *SSRN Electronic Journal*.

Gaddis, I., & Pieters, J. (2017). The Gendered Labor Market Impacts of Trade Liberalization: Evidence from Brazil [Publisher: [University of Wisconsin Press, Board of Regents of the University of Wisconsin System]]. *The Journal of Human Resources*, *52*(2), 457–490. Retrieved April 28, 2025, from https://www.jstor.org/stable/26450452

Heath, R., & Mushfiq Mobarak, A. (2015). Manufacturing growth and the lives of Bangladeshi women. *Journal of Development Economics*, *115*, 1–15.

Jayachandran, S. (2017). Fertility decline and missing women. *American Economic Journal: Applied Economics*, *9*(1), 118–139.

Jayachandran, S., & Kuziemko, I. (2011). Why Do Mothers Breastfeed Girls Less than Boys? Evidence and Implications for Child Health in India. *The Quarterly Journal of Economics*, *126*(3), 1485–1538.

Jones, L. E., Schoonbroodt, A., & Tertilt, M. (2011, January). Fertility Theories: Can They Explain the Negative Fertility-Income Relationship? In J. B. Shoven (Ed.), *Demography and the Economy* (pp. 43–100). University of Chicago Press.

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.

McCaig, B., Pavcnik, N., & Wong, W. F. (2022, December). Foreign and Domestic Firms: Long Run Employment Effects of Export Opportunities.

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.

Rosenzweig, M. R., & Wolpin, K. I. (1980). Testing the Quantity-Quality Fertility Model: The Use of Twins as a Natural Experiment. *Econometrica*, *48*(1), 227.

Rosenzweig, M. R., & Zhang, J. (2009). Do population control policies induce more human capital investment? Twins, birth weight and China's "one-child" policy. *Review of Economic Studies*, *76*(3), 1149–1174.

Sobek, M., & Kennedy, S. (2009). The Development of Family Interrelationship Variables for International Census Data. *Minnesota Population Center Working Paper Series*.

Solon, G., Haider, S. J., & Wooldridge, J. M. (2015). What are we weighting for? *Journal of Human Resources*, *50*(2), 301–316.

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

A Family Interrelationship Algorithm

Variables categorizing family interrelationships 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 the household head. However, this variable is ambiguous and erroneous for the following reasons. First, extended family is common in Vietnam: roughly 34.91% of the VHLSS sample lives 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 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 the 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 find a child's mother do 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 95.27% 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 interrelationships.

B Model Extension: Endogenous Bargaining

We now extend our model to incorporate household bargaining. Each parent's utility function includes gender-specific preferences for both fertility and sons:

$$U_g = \alpha_c \log(c_g) + \alpha_\ell \log(\ell_g) + \alpha_{ng} \log(n) + \alpha_{qg} \log(Q)$$

The household maximizes a weighted sum of utilities where bargaining weights depend on relative wages:

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

where $\lambda'(\cdot) > 0$, so higher relative wages increase bargaining power. This model yields three key equations for sex selection, mother's leisure, and fertility:

$$\begin{aligned} \log q &= \log\left(\frac{\lambda\alpha_{qm} + (1-\lambda)\alpha_{qd}}{\lambda\alpha_{nm} + (1-\lambda)\alpha_{nd}}\right) + \log\gamma - \log p_q + \log w_m \\ \log \ell_m &= \log\alpha_\ell + \log\lambda - \log w_m + \log W \\ \log n &= \log\left(\frac{\lambda\alpha_{nm} + (1-\lambda)\alpha_{nd}}{\gamma}\right) + \log W - \log w_m - \log\left(1 + \frac{\lambda\alpha_{qm} + (1-\lambda)\alpha_{qd}}{\lambda\alpha_{nm} + (1-\lambda)\alpha_{nd}}\right) \end{aligned}$$

The comparative statics reveal that this model can match most but not all of our empirical findings. Specifically:

1. Mother's wages increase sex selection ($\partial q/\partial w_m > 0$) if mothers have stronger relative preference for sons: $\alpha_{qm}/\alpha_{nm} > \alpha_{qd}/\alpha_{nd}$

2. Mother's wages reduce fertility ($\partial n/\partial w_m < 0$) if mothers prefer fewer children than fathers: $\alpha_{nm} < \alpha_{nd}$

3. Father's wages increase fertility ($\partial n/\partial w_d > 0$) under the same conditions

However, under these conditions, the model predicts father's wages would reduce sex selection ($\partial q/\partial w_d < 0$), contrary to our finding of no effect. This result suggests that while bargaining may play a role, it cannot fully explain the asymmetric effects of parental industry exposures that we observe empirically.

C Additional Figures and Tables

Figure A1: Changes in U.S. Tariffs on Vietnamese Imports Following the 2001 BTA

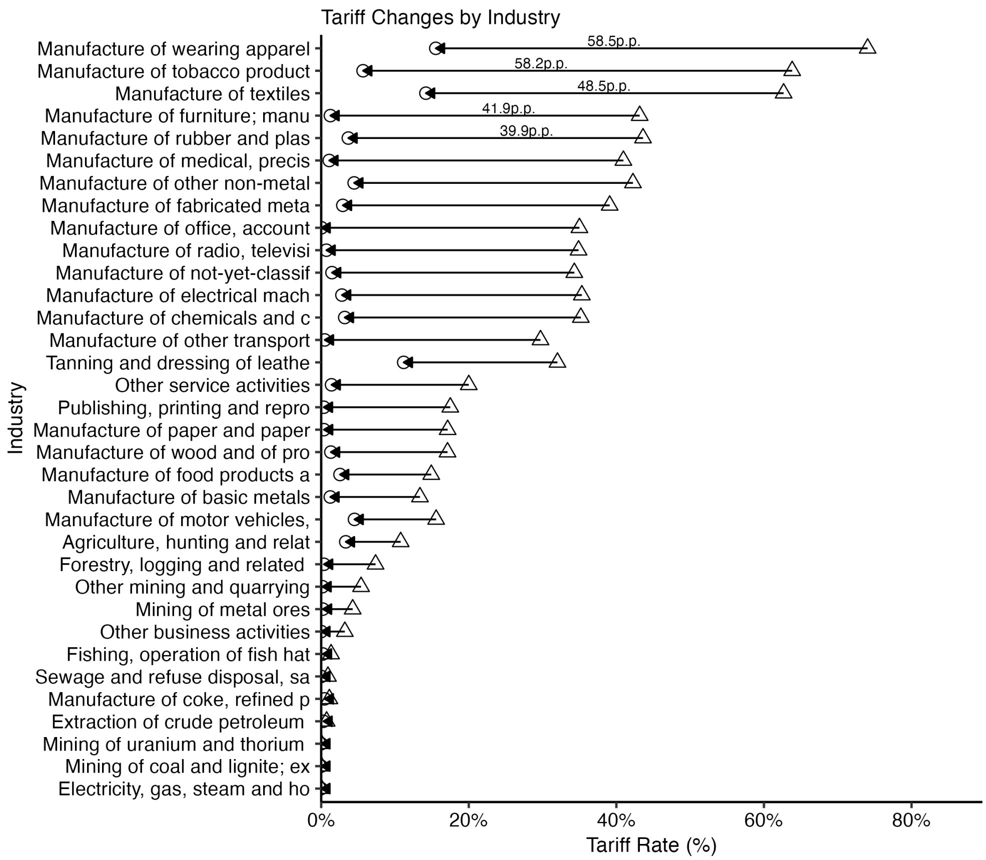

Notes: This figure shows the changes in U.S. tariff rates on Vietnamese imports by industry following the 2001 U.S.-Vietnam Bilateral Trade Agreement. The start of each arrow represents the Column 2 tariff rate, while the arrowhead shows the new Most Favored Nation (MFN) tariff rate. Percentage points indicate the size of the tariff reduction for the five most affected industries. Data source: McCaig and Pavcnik (2018).

Figure A2: Event Study of Mother's Industry Exposure to BTA on Male Birth

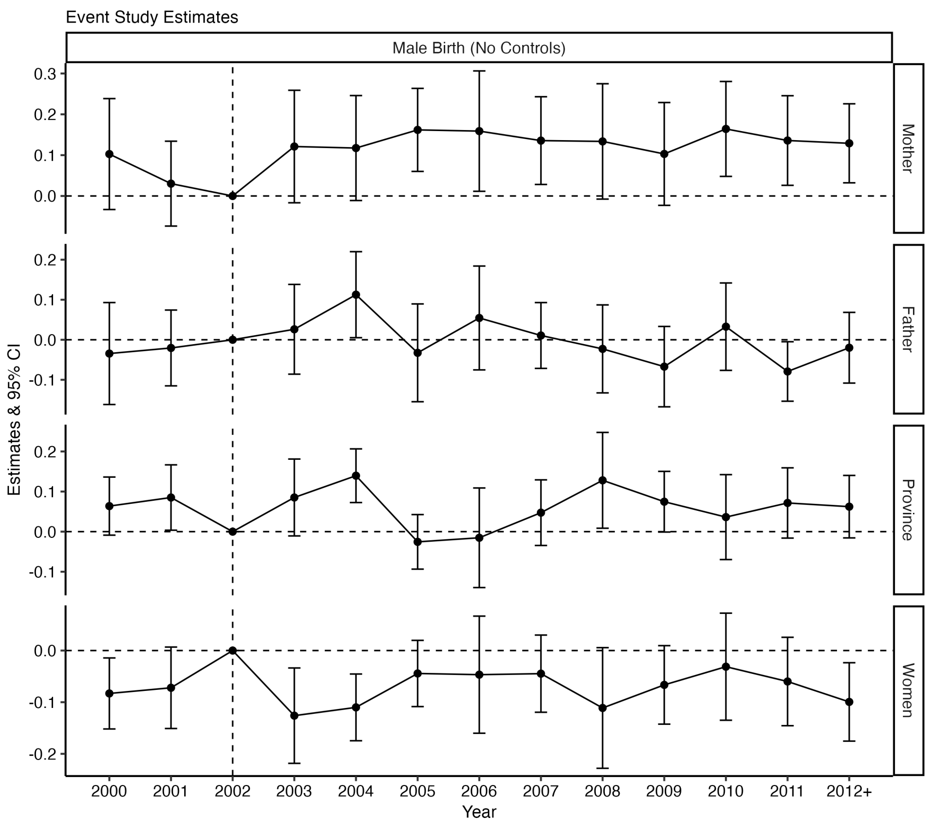

Notes: Panel (a) plots the coefficients from event study regressions showing how the effect of maternal industry exposure to trade liberalization evolves over time, with 2001 as the reference year. Panel (b) shows how the effect of maternal industry exposure varies across different groups. For each subgroup, we estimate the coefficient on maternal industry exposure ($\tau_m \times Post$) from our baseline specification (10). Both panels report 95% confidence intervals using standard errors clustered at the province level. Sample includes children aged 0-1 from VHLSS 2002-2016.

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 distinguished. For child-child couples, assume the first listed spouse is the biological child to the household head, hence the second listed spouse is the child-in-law.

Table A2: Rules for Child-Parent Relationship Construction (after SPLOC is Generated)

Rule	**Child's relationship to head**	**Parent's relationship to head**	**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.

Table A4: Key Predictions of Our Model and Alternative Theories

	Sex Ratio	**Mother's Labor Supply**	**Fertility**
A. Baseline Model (Work-Childcare Trade-off)			
Increase in w_m (Mother's Wage)	+	+	−
Increase in w_d (Father's Wage)	0	−	+
Increase in I (Non-labor Income)	0	−	+
Increase in p_q (Cost of Sex Selection)	−	+	0
B. Alternative Theories			
Fertility Decline[1]	+	n/a	−
Income Effects[2]	+	0	+
Relative Returns to Daughters[3]	−	+	0
C. Extended Model with Endogenous Bargaining[4]			
Increase in w_m (Mother's Wage)	+	+	−
Increase in w_d (Father's Wage)	−	−	+
Increase in I (Non-labor Income)	0	−	+
Increase in p_q (Cost of Sex Selection)	−	+	0

Notes: + indicates a positive effect; − indicates a negative effect; 0 indicates no effect; n/a indicates not applicable or no specific prediction.
[1] Jayachandran (2017) shows that the desired sex ratio becomes more male-skewed as fertility declines in India.
[2] See Almond et al. (2019) for the impact of nonlabor income effects on sex selection in China.
[3] Based on Qian (2008), higher returns to female labor increase the preference for daughters.
[4] These predictions from our extended model with endogenous bargaining assume mothers have stronger son preference ($\alpha_{qm} > \alpha_{qd}$) and prefer fewer children ($\alpha_{nm} < \alpha_{nd}$) than fathers. Fertility predictions become unambiguous only when incorporating asymmetric childcare costs.

Table A5: Robustness: Standard errors clustered by industry

	MaleBirth		log(MomWorkHrs)		AnyBirth	
	(1)	(2)	(3)	(4)	(5)	(6)
$1_m \times$ Post	0.09***	0.09***	0.21**	0.19***	-0.006**	-0.006**
	(0.03)	(0.03)	(0.08)	(0.07)	(0.003)	(0.003)
$1_d \times$ Post	0.010	0.007	-0.02	-0.01	0.005***	0.005***
	(0.02)	(0.02)	(0.06)	(0.06)	(0.001)	(0.001)
$1_p \times$ Post	0.003	0.001	-0.07*	-0.09*	0.002**	0.002
	(0.01)	(0.01)	(0.04)	(0.05)	(0.0009)	(0.001)
$1_p^w \times$ Post	-0.02*	-0.02	0.16***	0.14***	-0.001	-0.002**
	(0.01)	(0.02)	(0.02)	(0.02)	(0.001)	(0.0009)
R^2	0.009	0.01	0.29	0.30	0.29	0.29
Observations	28,419	28,419	26,147	26,147	2,011,963	2,011,963
Dep. Var. Mean	0.52	0.52	4.7	4.7	0.02	0.02
Mom FE					✓	✓
Birth Year FE	✓	✓			✓	✓
Survey Year FE			✓	✓		
Dad ISIC FE	✓	✓	✓	✓		
Mom ISIC FE	✓	✓	✓	✓		
Province FE	✓	✓	✓	✓		
Controls		✓		✓		✓

Notes: Identical specifications to the baseline tables, but standard errors are computed with two-way clustering by province and by 2-digit ISIC industry. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A6: Robustness: Two-way clustered standard errors (province & industry)

	MaleBirth		log(MomWorkHrs)		AnyBirth	
	(1)	(2)	(3)	(4)	(5)	(6)
$1_m \times$ Post	0.09***	0.09***	0.21**	0.19***	-0.006**	-0.006**
	(0.02)	(0.03)	(0.08)	(0.07)	(0.002)	(0.003)
$1_d \times$ Post	0.010	0.007	-0.02	-0.01	0.005***	0.005***
	(0.02)	(0.02)	(0.06)	(0.05)	(0.001)	(0.001)
$1_p \times$ Post	0.003	0.001	-0.07	-0.09	0.002*	0.002
	(0.01)	(0.01)	(0.05)	(0.05)	(0.001)	(0.002)
$1_p^w \times$ Post	-0.02***	-0.02	0.16***	0.14***	-0.001	-0.002
	(0.007)	(0.01)	(0.04)	(0.03)	(0.002)	(0.001)
R^2	0.009	0.01	0.29	0.30	0.29	0.29
Observations	28,419	28,419	26,147	26,147	2,011,963	2,011,963
Dependent variable mean	0.52	0.52	4.7	4.7	0.02	0.02
Mom FE					✓	✓
Birth Year FE	✓	✓			✓	✓
Survey Year FE			✓	✓		
Dad ISIC FE	✓	✓	✓	✓		
Mom ISIC FE	✓	✓	✓	✓		
Province FE	✓	✓	✓	✓		
Controls		✓		✓		✓

Notes: Identical specifications to the baseline tables, but standard errors are clustered by 2-digit ISIC industry. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A7: Robustness: Inclusion of Region-by-Year Fixed Effects

	MaleBirth		log(MomWorkHrs)		AnyBirth	
	(1)	(2)	(3)	(4)	(5)	(6)
$1_m\times$ Post	0.09***	0.09***	0.18***	0.18***	-0.006***	-0.006***
	(0.02)	(0.02)	(0.04)	(0.04)	(0.001)	(0.001)
$1_d\times$ Post	0.009	0.009	0.01	0.02	0.005***	0.005***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.001)	(0.001)
$1_p\times$ Post	0.008	0.009	0.02	0.02	0.002	0.002
	(0.02)	(0.02)	(0.07)	(0.07)	(0.002)	(0.002)
$1_p^w\times$ Post	-0.03*	-0.03	0.02	0.03	-0.001	-0.002
	(0.02)	(0.02)	(0.08)	(0.08)	(0.002)	(0.002)
R^2	0.01	0.01	0.30	0.30	0.29	0.29
Observations	28,419	28,419	26,147	26,147	2,011,963	2,011,963
Dep. Var. Mean	0.52	0.52	4.7	4.7	0.02	0.02
Region-Year FE	✓	✓	✓	✓	✓	✓
Mom FE					✓	✓
Birth Year FE	✓	✓			✓	✓
Dad ISIC FE	✓	✓	✓	✓		
Mom ISIC FE	✓	✓	✓	✓		
Province FE	✓	✓	✓	✓		
Controls		✓		✓		✓

Notes: Specifications replicate baseline results but now include region-by-year fixed effects to absorb any region-specific, time-varying factors such as infrastructure improvements. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A8: Robustness: Unweighted Estimates (No Survey Weights)

	MaleBirth		log(MomWorkHrs)		AnyBirth	
	(1)	(2)	(3)	(4)	(5)	(6)
$1_m \times$ Post	0.07**	0.06**	0.20***	0.18***	-0.006***	-0.007***
	(0.03)	(0.03)	(0.04)	(0.04)	(0.001)	(0.001)
$1_d \times$ Post	0.01	0.01	0.01	0.01	0.005***	0.005***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.002)	(0.002)
$1_p \times$ Post	-0.008	-0.01	-0.09	-0.11	0.002	0.002
	(0.01)	(0.01)	(0.07)	(0.07)	(0.002)	(0.002)
$1_p^w \times$ Post	-0.007	-0.010	0.16**	0.13*	-0.003	-0.003
	(0.01)	(0.01)	(0.07)	(0.07)	(0.002)	(0.002)
R^2	0.007	0.008	0.26	0.26	0.29	0.29
Observations	28,419	28,419	26,147	26,147	2,012,206	2,012,206
Dep. Var. Mean	0.52	0.52	4.7	4.7	0.02	0.02
Mom FE					✓	✓
Birth Year FE	✓	✓			✓	✓
Survey Year FE			✓	✓		
Dad ISIC FE	✓	✓	✓	✓		
Mom ISIC FE	✓	✓	✓	✓		
Province FE	✓	✓	✓	✓		
Controls		✓		✓		✓

Notes: Specifications match those in the main text but exclude the VHLSS sampling weights. All regressions include the same set of fixed effects and controls as indicated. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A9: Robustness: Logit Estimation for Male Birth Outcome

	MaleBirth	
	(1)	(2)
$1_m \times$ Post	0.4***	0.4***
	(0.1)	(0.1)
$1_d \times$ Post	0.04	0.03
	(0.1)	(0.1)
$1_p \times$ Post	0.01	-0.005
	(0.06)	(0.06)
$1_p^w \times$ Post	-0.09	-0.09
	(0.06)	(0.06)
Observations	28,417	28,417
Dep. Var. Mean	0.5	0.5
Birth Year FE	✓	✓
Province, Industry FEs	✓	✓
Controls		✓

Notes: The table presents logistic regression (logit) estimates of the effect of parental and province-level exposures on the probability of a male birth. Both specifications include birth-year, province, and parents' industry fixed effects. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A10: Continuous Tariff Exposure Measures

	MaleBirth		log(MomWorkHrs)		AnyBirth	
	(1)	(2)	(3)	(4)	(5)	(6)
$\tau_m\times$ Post	0.26**	0.26**	0.29**	0.29**	-0.01***	-0.01***
	(0.10)	(0.11)	(0.14)	(0.14)	(0.003)	(0.003)
$\tau_d\times$ Post	-0.05	-0.05	0.07	0.09	0.02***	0.02***
	(0.11)	(0.11)	(0.14)	(0.14)	(0.006)	(0.006)
$\tau_p\times$ Post	0.03	0.009	-0.22**	-0.29***	0.003	0.010***
	(0.03)	(0.04)	(0.10)	(0.10)	(0.004)	(0.003)
$\tau_p^w\times$ Post	-0.03	-0.02	0.25***	0.22**	-0.001	-0.003
	(0.03)	(0.03)	(0.09)	(0.08)	(0.004)	(0.002)
R^2	0.009	0.01	0.29	0.30	0.29	0.29
Observations	28,419	28,419	26,147	26,147	2,011,963	2,011,963
Dep. Var. Mean	0.52	0.52	4.7	4.7	0.02	0.02
Mom FE					✓	✓
Birth Year FE	✓	✓			✓	✓
Survey Year FE			✓	✓		
Dad ISIC FE	✓	✓	✓	✓		
Mom ISIC FE	✓	✓	✓	✓		
Province FE	✓	✓	✓	✓		
Controls		✓		✓		✓

Notes: Specifications replicate baseline results using continuous tariff-cut measures (τ variables) rather than binary exposure indicators. Columns (1)–(2) use *MaleBirth*, columns (3)–(4) use log(MomWorkHrs), and columns (5)–(6) use *AnyBirth* as dependent variables. All regressions include the same fixed effects and controls as baseline models. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.