Better Jobs for Mothers, More Sons at Birth*

Nghiêm Q. Huỳnh† Ngoc T. T. Nguyen‡

March 26, 2026

Abstract

Women's economic gains are expected to favor daughters, yet Vietnam's sex ratio at birth rose from 107 boys per 100 girls to 112 between 1999 and 2015 as trade liberalization expanded female employment. The missing element is asymmetric childcare costs. Because mothers bear primary childcare responsibility, a maternal wage gain raises the opportunity cost of children in a way a paternal wage gain does not. Using industry-level tariff cuts from the 2001 US-Vietnam trade agreement, we assign exposure separately to each parent. Mothers in high-exposure industries are more likely to have a male birth and have lower fertility. Fathers show higher fertility but no change in sex selection. The contrast identifies maternal time costs, not income or relative returns, as the channel linking female employment growth to rising sex ratios.

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

Keywords: Son preference, Sex ratio at birth, Fertility, Female labor supply, Missing women

*We thank S. Anukriti, Klaus Desmet, 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 and conference participants at Boston College, Yale, the University of Oklahoma, Fulbright University Vietnam, International School at Vietnam National University (Hanoi), the Chinese Economists Society North America Annual Conference (University of Michigan), and NEUDC (Tufts University). Special thanks to Brian McCaig for his guidance on the VHLSS data. Huỳnh acknowledges financial support from the Yale Council on Southeast Asian Studies. All errors are our own.

†University of Oklahoma, nqh@ou.edu

‡Disney Streaming, jade.ng.nguyen@gmail.com

1 Introduction

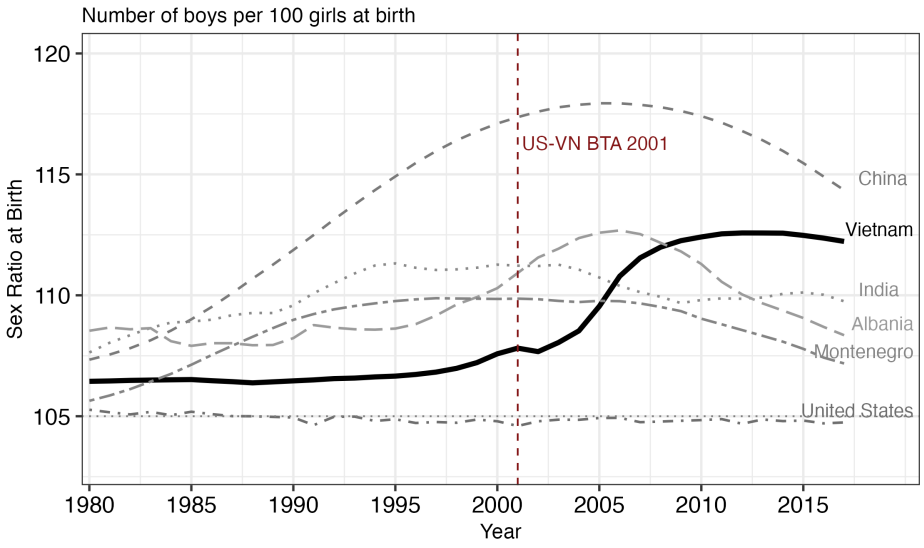
Vietnam's sex ratio at birth rose from 107 boys per 100 girls in 1999 to 112 by 2015, a period in which trade liberalization drew millions of women into formal employment. The pattern is not unique to Vietnam. Figure 1 shows that sex ratios at birth have worsened across a range of countries as trade integration deepened. These trends sit uncomfortably alongside a well-established finding in development economics. Female economic empowerment improves daughter outcomes, whether through grandmother pensions (Duflo, 2003) or higher female income in tea-growing regions (Qian, 2008). If women's economic gains help daughters, why do sex ratios worsen precisely when women enter formal labor markets?

The missing element is a constraint that binds mothers and fathers asymmetrically. Mothers bear primary responsibility for childcare in virtually all societies, a regularity that Doepke et al. (2023) identify as central to modern fertility economics. This asymmetry means that a maternal wage gain and a paternal wage gain are not equivalent. When maternal wages rise, the opportunity cost of time devoted to children rises, creating pressure to reduce fertility and to concentrate remaining births on preferred children. A paternal wage gain raises household resources without tightening any parent's childcare time budget, so it increases fertility without shifting sex composition. Income-based and relative-returns accounts, the two leading frameworks in the sex-selection literature, do not make this distinction. Models that treat parental income as fungible predict that any income gain raises sex selection regardless of its source (Almond et al., 2019). Relative-returns accounts predict that higher female economic returns in local labor markets reduce sex selection by raising daughters' value (Rosenzweig and Schultz, 1982; Qian, 2008). The maternal time-cost mechanism predicts neither. It implies that maternal wages raise sex selection while paternal wages raise fertility but leave sex composition unchanged.

We formalize the maternal time-cost mechanism through a household fertility model in which mothers bear disproportionate childcare costs and households exhibit son preference, and exploit the 2001 US-Vietnam Bilateral Trade Agreement (BTA) to test its predictions. The agreement generated industry-specific tariff cuts that reached mothers and fathers at different intensities depending on where each parent worked. Separately identifying maternal and paternal channels requires observing each parent's industry around the time of conception. Prior studies link trade shocks or female income gains to sex ratios using geographic variation, assigning a common district-, county-, or province-level shock to all households in the same location (Anukriti and Kumler, 2019; Chakraborty, 2015). Geographic exposure cannot separately identify maternal wages, paternal wages, and local income effects, because all three vary together at the location level. These designs also assume that families did not relocate after the trade shock, since they assign exposure based on current residence rather than residence at the time of conception.[1]

[1] Anukriti and Kumler (2019) use DLHS-2 (2002–2004) retrospective birth histories linked to district-level tariff exposure, assuming births occurred in the mother's district of survey residence. Chakraborty (2015) use NFHS-2 (1998–1999) birth histories linked to district tariff decline and restrict robustness checks to women resident for more than nine

Figure 1: Sex Ratio at Birth Trends Across Countries



Notes: Sex ratio at birth (boys per 100 girls) for Vietnam and selected countries with documented imbalances, from Chao et al. (2019). The vertical dashed line marks 2001, the year the US-Vietnam Bilateral Trade Agreement took effect. The horizontal dotted line indicates the natural sex ratio at birth (approximately 105 boys per 100 girls).

We address both problems by combining the 1999 Vietnam Population Census with the Vietnam Household Living Standards Surveys (VHLSS), a biennial nationally representative survey spanning 2002 to 2010. The Census provides a pre-treatment birth baseline. The VHLSS records industry, location, and employment for both parents at each survey wave, and we develop an algorithm to link children to both parents within extended households and assign each parent the tariff cut applied to their industry. Including both parents' industry exposures simultaneously within the same province separates maternal from paternal channels and absorbs location-level confounders, so the identification does not rely on assumptions about the absence of geographic sorting. Assigning exposure separately to mothers, fathers, and provinces allows us to identify maternal opportunity costs, paternal income, non-labor income, and local returns to female labor as distinct channels.

The evidence is consistent with the maternal time-cost mechanism. Women in high-exposure industries are approximately 6 percentage points more likely to have a male child, 16 percentage points more likely to hold a wage job (from a base of 25 percent), and work approximately 14 percent more hours per month. Annual birth probability falls by 0.6 percentage points, a 30 percent decline relative to the mean. Exposure through fathers' industries or local labor markets has no detectable effect on sex selection, though paternal exposure increases fertility by about 25

years. Qian (2008) uses county-level crop shares to proxy relative female and male income, inferring county of birth from county of residence for long-stayers.

percent, consistent with the model's prediction that paternal wages operate primarily through income rather than time costs.

The mother-panel analysis strengthens the causal interpretation. Using the rotating panel structure of the VHLSS, we track the same women across the 2002, 2004, and 2006 waves and assign each mother's treatment based on her pre-BTA industry, fixing exposure at the 2002 baseline, shortly after the BTA took effect, rather than allowing it to update with post-BTA occupation changes. The within-mother fixed-effects estimate of the sex-selection effect is larger than the cross-sectional estimate, a pattern consistent with attenuation in the cross-section rather than amplification by positive sorting on pre-existing son preference. Adding an indicator for whether the household already has a son substantially attenuates the maternal coefficient, a pattern descriptively consistent with sex selection responding to prior child sex composition.

The contrast between maternal and paternal effects is difficult to reconcile with the leading alternatives. Income-based explanations predict that any income gain raises sex selection regardless of source, yet paternal exposure raises fertility without affecting sex composition. A bargaining extension in which relative wages shift decision power predicts that paternal wage gains should raise sex selection by shifting power toward fathers, yet paternal exposure leaves sex composition unchanged (Section 7). Relative-returns explanations predict that stronger female labor demand should favor daughters, yet in provinces where female-intensive industries expanded most, maternal wage employment rises while sex-selection rates do not fall. The broader literature on female income gains and daughter outcomes reaches mixed conclusions (Jayachandran and Voena, 2026), a pattern consistent with institutions mediating how shocks translate into outcomes.

Our model provides the economic mechanism behind Jayachandran (2017)'s finding that fertility decline accounts for one-third to one-half of India's rising sex ratio. In that account, smaller families mechanically raise the sex ratio because each birth carries more weight. Our framework identifies the economic force driving both outcomes. Rising maternal wages reduce fertility by raising the opportunity cost of children and simultaneously intensify sex selection by concentrating remaining births on preferred sons.

A growing literature studies how trade and manufacturing growth affect women's labor market and demographic outcomes (Gaddis and Pieters, 2017; Heath and Mobarak, 2015), but does not separate maternal from paternal exposure channels. By assigning exposure separately to mothers, fathers, and local labor markets, we can test which channel drives fertility and sex selection.[2] We also complement the Vietnam BTA literature (McCaig, 2011; McCaig and Pavcnik, 2018; McCaig et al., 2025) by documenting that the same trade shock can improve women's labor-market outcomes while worsening prenatal gender imbalance.

[2] A related literature studies sex selection through marriage-market and family-formation channels. Drozdoff et al. (2025) show that foreign manufacturing jobs in Vietnam delay marriage and childbirth for young women, especially through migration. Their paper and ours speak to different margins affected by female-oriented labor demand. They study premarital migration and delayed family formation, whereas we study parental exposure at conception and show that, among births that occur, higher maternal opportunity costs reallocate births toward sons. Borker et al. (2024) show that relative wealth within caste in India predicts sex-selection intensity through dowry and matching equilibria.

Because the operative channel runs through maternal time costs rather than income, income transfers or wage subsidies alone will not break the link between female employment and sex selection.[3] Policies that directly reduce the maternal time constraint, such as expanding access to institutional childcare, address the channel through which female employment growth worsens sex ratios.

The paper proceeds as follows. Section 2 presents the model and testable predictions. Section 3 describes the BTA and the VHLSS data. Sections 4–5 present the empirical strategy and results for the cross-section and panel samples. Section 6 reports robustness checks, and Section 7 tests alternative mechanisms. Section 8 concludes.

2 Model

Understanding how economic development affects fertility and sex composition requires a framework that distinguishes between different sources of household income. Standard models that treat household income as fungible cannot explain why trade liberalization might affect sex ratios differently depending on which family member experiences the wage gain. A mother's wage increase, a father's wage increase, and a windfall transfer all raise household resources, yet they need not have equivalent effects on fertility decisions. The reason lies in how each income source interacts with the constraints that bind household members differentially.

We build on the framework of Jones et al. (2011) that embeds gender-asymmetric childcare costs into a unitary household model. In their framework, mothers bear a disproportionate share of the time costs of raising children, a fact that holds across most societies (Doepke et al., 2023).[4] When this asymmetry exists, the mother's wage carries information beyond its contribution to household income. It also represents the opportunity cost of the time she devotes to childcare. A higher maternal wage tightens her time constraint in a way that a higher paternal wage or non-labor income does not. We extend their model by introducing son preference, which allows households to choose not only total fertility but also the sex composition of births. This extension generates distinct channels through which different income sources affect both the number of children and the probability that a birth is male.

The framework separates four channels of influence on fertility and sex selection. Maternal wages operate through the opportunity cost of childcare time. Paternal wages contribute to household resources without altering the binding time constraint. Non-labor income provides a pure income effect. Local labor market conditions for women affect the economic value of daughters

[3] Anukriti (2018) finds that a cash-transfer program in Haryana lowered birth rates but worsened the sex ratio at birth, as households ensured at least one son before stopping. Our design goes further by separating maternal and paternal income, showing not only that income gains that do not tighten the maternal time constraint leave sex selection unchanged, but why.

[4] Vietnam is no exception. In our data, 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). Formal childcare services remain limited outside major urban centers, and cultural norms strongly favor maternal care for young children.

relative to sons. By allowing each channel to operate independently, the model clarifies when rising female wages will increase sex selection and when improved female labor market conditions will reduce it. These two forces can coexist and even operate simultaneously within the same household.

We now develop the formal model. Consider a household with a mother (m) and father (d) who jointly decide consumption (c_m, c_d), leisure time (ℓ_m, ℓ_d), total fertility (n), and the sex composition of their children. We model this through $q \in [0, 1]$, the probability that a birth is male, which the household can influence through sex-selective abortion or other practices.[5] The expected number of sons is $Q = qn$. Each parent's utility 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.

Son preference in Vietnam predates our study period and reflects traditional patrilineal kinship structures where sons provide old-age support and continue the family line (Belanger, 2002). We take this preference as given rather than modeling its origins, focusing instead on how economic shocks interact with pre-existing preferences to affect fertility choices.[6]

We now formalize the time constraints that generate these differential effects. Each parent has a time endowment normalized to one. Only mothers incur childcare time cost γ per child, reducing their available time for labor and leisure. Fathers face no such constraint ($\gamma_d = 0$). The mother's time constraint is:

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

where h_m denotes maternal labor hours and ℓ_m denotes leisure. The father's time constraint is:

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

where h_d and ℓ_d are paternal labor hours and leisure. 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 from sources such as transfers from extended family members or returns from land and other assets. Non-labor income provides a pure income effect that relaxes the budget constraint without altering the

[5] The variable q translates directly to the sex ratio at birth: Sex Ratio $= 100 \times q/(1-q)$. For example, $q = 0.51$ corresponds to 104 males per 100 females, while $q = 0.52$ corresponds to 108. The log formulation permits positive but declining marginal utility from additional sons, appropriate for Vietnam where son preference remains culturally important but allows for diminishing returns to male children.

[6] The model focuses on immediate time costs rather than the long-run old-age support value of sons. That value affects the baseline level of son preference (α_q) but does not alter how maternal wages shift sex selection at the margin. An alternative reading is that the old-age support motive operates through relative labor market returns, with higher female earnings reducing son preference as daughters become more economically valuable. The female-weighted provincial exposure P_p^w tests this directly, and its effect is discussed in Section 7.

opportunity cost of either parent's time. This income source is conceptually distinct from wages because it does not tighten any time constraint. 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. This cost encompasses both direct expenses (medical fees for sex-determination tests and procedures) and indirect costs (physical and psychological burden on mothers, social stigma, and the forgone economic and social benefits that daughters would provide through marriage networks and elder care (Rosenzweig and Schultz, 1982)).

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 $\lambda_m + \lambda_d = 1$.[7]

By defining the household's total income as $W = I + w_m + w_d$, the optimal choices in log form are given by three equations. The probability of male birth is:

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

This expression is an interior solution that assumes $q \in (0, 1)$.[8] Maternal leisure is:

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

Total fertility is:

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

To understand these equations, consider how different income sources affect household decisions. When the mother's wage w_m rises, Equation (6) shows the sex ratio q increases. Higher maternal wages raise the opportunity cost of time spent on childcare, and since each child requires maternal time γ, the household responds on two margins. Total fertility falls (Equation 8) while the probability of male births rises. Mothers also work more because their time becomes more valuable (Equation 7, where lower leisure ℓ_m means higher total non-leisure time $1 - \ell_m = h_m + \gamma n$). Because fertility n also falls, market hours $h_m = 1 - \ell_m - \gamma n$ unambiguously increase.[9]

[7] We treat bargaining weights as exogenous in the baseline model. Appendix B.1 extends the model to allow bargaining power to depend on relative wages, generating contrasting predictions discussed in Section 7.

[8] The constraint $q \leq 1$ would bind only if the right-hand side implied $q \geq 1$. In the empirically relevant range (q between 0.51 and 0.54), the interior solution is innocuous.

[9] Substituting Equations (7) and (8) into the time constraint yields $h_m = 1 - (\lambda_m \alpha_\ell + \alpha_n) W / w_m$. Differentiating, $\partial h_m / \partial w_m = (\lambda_m \alpha_\ell + \alpha_n)(I + w_d)/w_m^2 > 0$, since the substitution effect (higher opportunity cost of all non-market time) dominates the income effect. For paternal wages and non-labor income, $\partial h_m / \partial w_d = \partial h_m / \partial I = -(\lambda_m \alpha_\ell + \alpha_n)/w_m < 0$,

Paternal wages w_d and non-labor income I operate differently. Neither income source appears in Equation (6). When the father's wage rises or the household receives a transfer, total income W increases but the mother's time constraint remains unchanged. These pure income effects increase fertility and maternal leisure, but they do not affect the sex composition of births. The household becomes richer and has more children, yet the male-female balance remains unchanged because fathers bear no direct time cost of childrearing.

Table 1 summarizes the comparative statics from these three income sources across three outcomes. A positive sign indicates that an increase in the income source raises the outcome, a negative sign indicates a decrease, and zero indicates no effect. Consider the first row as an example. When the mother's wage rises, sex selection increases (column 1, positive) because the opportunity cost of childcare time rises. Maternal labor supply also increases (column 2, positive) as mothers substitute toward market work. Fertility falls (column 3, negative) because children have become more costly in terms of foregone maternal earnings. The remaining rows follow similar logic, with each sign determined by whether the income source tightens the maternal time constraint.

The pattern of signs in Table 1 distinguishes our framework from models where sex selection is simply a normal good that rises with household income. In models like that of Almond et al. (2019), any income increase would raise both fertility and sex selection. Higher paternal wages and non-labor income would increase sex selection just as maternal wages do. Our model predicts otherwise. Because only mothers bear childcare time costs, paternal wages and non-labor income have zero effect on sex selection (column 1, rows 2-3) even though they increase fertility (column 3, rows 2-3). The contrast between the first row and the remaining rows within the 3×3 block provides a direct way to distinguish the maternal time-cost mechanism from income-based explanations.

The sign pattern in Table 1 distinguishes our framework from income-based accounts, but raises a further question about relative-returns frameworks. Rising maternal wages might seem to already capture the mechanism of Rosenzweig and Schultz (1982) and Qian (2008), in which higher female labor market returns raise the value of daughters and reduce son preference. In those frameworks, higher female wages signal better economic prospects for daughters, shifting parental preferences toward girls. Our model predicts the opposite. Higher maternal wages raise sex selection by tightening the mother's childcare time constraint, not by raising daughters' economic value. The distinction is between individual wage gains and local labor market conditions for women as a group. A mother directly exposed to the trade shock sees her own wage rise and her childcare time become more costly, driving sex selection up. A mother in an unaffected industry but in a province where female-intensive production has broadly expanded faces no personal wage change, yet daughters in her local economy now command better earnings prospects.

The cost parameter p_q provides a way to capture the relative returns channel. When female-

a pure income effect that raises both leisure and fertility, crowding out market hours.

Table 1: Model Predictions of How Income Sources Affect Household Outcomes

	Outcome		
	(1)	(2)	(3)
Increase in:	Sex Ratio	Maternal Labor	Fertility
	q (male share)	h_m	n
Mother's wage (w_m)	+	+	−
Father's wage (w_d)	0	−	+
Non-labor income (I)	0	−	+
Local female labor demand (p_q)	−	0	0

Notes: The first three rows show effects of income sources that differ in whether they tighten the maternal time constraint, corresponding to the treatment channels (M_i, F_i, P_p) in Equation (11). The fourth row captures the relative-returns mechanism, in which local female labor demand raises the economic value of daughters and reduces sex selection, tested using a female-weighted provincial exposure measure P_p^w in Section 7. See Equations (6)–(8) for derivations.

intensive industries expand in a province, daughters become economically more valuable in that local labor market. This provincial-level variation raises the opportunity cost of sex-selective practices by increasing what families forgo when they select against daughters. The model predicts that higher p_q reduces sex selection (Table 1, row 4), working in the opposite direction to maternal wage effects. This distinction will become clearer when we discuss the trade shock, which provides variation in both individual wages (through industry exposure) and provincial labor market conditions (through local industry composition).

The baseline model abstracts from endogenous bargaining power and heterogeneous preferences across spouses. Appendix B.1 develops an extension that relaxes these assumptions. In that extension, each spouse's weight in household decisions depends on their outside option, which improves with their wage. The direction of the bargaining effect on sex selection depends on whose son preference is stronger. If mothers have weaker son preference than fathers, then a shift in bargaining power toward the mother reduces the household's demand for son selection. Because the mother's bargaining weight rises with her own wage, the extension would predict a negative relationship between maternal wages and sex selection, directly contradicting the baseline model's prediction. However, the main discriminating prediction used to evaluate the bargaining channel empirically concerns the paternal-wage effect on sex selection. We return to this alternative mechanism in Section 7 after presenting the empirical results.

3 Data and Measures

We implement the theoretical framework using Vietnam's 2001 Bilateral Trade Agreement (BTA) shock. This section describes how we construct exposure measures corresponding to the theoretical channels in Table 1, and then presents the household survey and census data used to observe parents' industries and birth outcomes.

3.1 The 2001 BTA and Exposure Measures

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.[10] Following the agreement, Vietnam's exports to the US surged from 0.7 billion dollars in 2000 to nearly 10 billion dollars in 2008.

We construct two complementary measures to capture how these differential tariff cuts transmit to fertility and sex selection through the labor market channels described in the model. The first captures direct exposure through parents' industries:

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

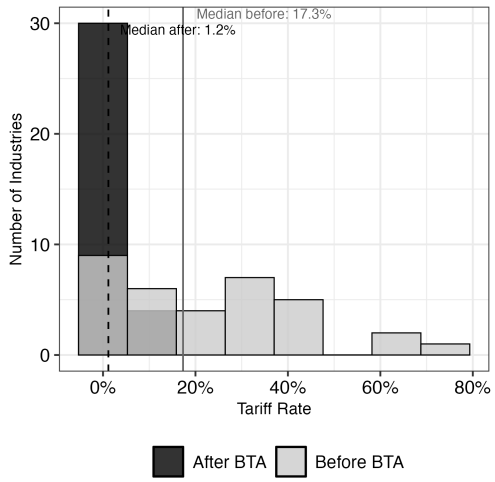
where j denotes a parent's industry. A higher value of τ_j indicates a larger effective industry tariff cut. Because mothers and fathers tend to work in different industries, we construct separate measures for each parent to capture relative changes in labor market returns.[11] These industry-level measures proxy for the maternal and paternal wage channels (w_m and w_d) in Table 1. The measure captures genuine labor market improvements in the Vietnamese context because, using the same construction, McCaig and Pavcnik (2018) document large-scale reallocation from informal to formal employment following the BTA, and McCaig et al. (2025) show significant growth in both foreign and domestic firm entry in more exposed industries. We confirm this formalization channel directly in our own sample of parents at the time of fertility decisions below. Because female labor force participation was already approximately 70 percent before the BTA, the shock's labor market effects on employed mothers operated primarily through wages and hours rather than through new labor force entry.[12]

[10] While Vietnam was obligated to reduce tariffs mainly on agricultural and food products, these reductions were 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).

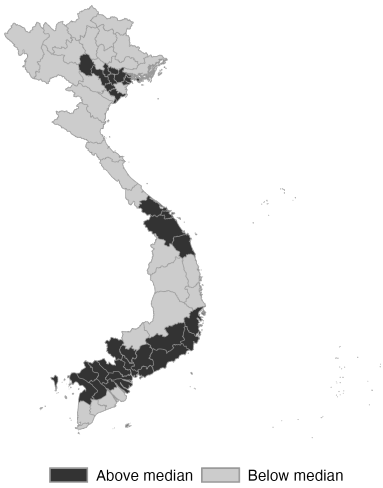
[11] The correlation between maternal and paternal tariff exposure in our sample is consistently low, ranging from 0.24 to 0.36 across household survey years, described below, confirming that parents face different levels of exposure. The correlation between individual parental exposure and the provincial shift-share τ_p is even lower, remaining below 0.15 throughout the sample period. Appendix Figure C1 displays the full pairwise correlation structure by wave.

[12] Huynh and Ku (2025) use a province-level shift-share exposure to document that the BTA reduced female labor force participation, particularly among married and wealthier women. Our paper, by contrast, uses individual industry-level exposure and focuses on demographic consequences for mothers who remained employed.

Figure 2: Variation in Trade Exposure Across Industries and Provinces



(a) Industry Tariff Changes



(b) Province-Level Exposure

Notes: Panel (a) shows the distribution of US tariff rates on Vietnamese exports before (Column 2) and after (MFN) the 2001 BTA across two-digit ISIC tradeable industries using data from McCaig (2011). Panel (b) shades provinces by whether their shift-share exposure $\tau_p = \sum_j s_{jp}\tau_j$ exceeds the sample median, where s_{jp} is industry j's share of employment in province p from the 1999 Vietnam Census. Darker shading indicates above-median exposure.

The second measure is a province-level shift-share that captures broader local economic effects:

$$\tau_p = \sum_j s_{jp}\tau_j \tag{10}$$

where s_{jp} is the share of province p's employment in tradeable industry j in the 1999 Census. Following Kovak (2013), employment shares are computed over all tradeable industries. This provincial measure proxies for the non-labor income channel (I) in Table 1 and captures broader local economic effects that individual industry measures miss. Export expansion raises returns to land and agricultural assets, and income gains spread through intra-household transfers within extended families. Consistent with these channels, McCaig (2011) uses this province-level exposure measure to show that the BTA significantly reduced poverty rates across Vietnamese provinces, and McCaig et al. (2022) find that district-level variation in the same BTA shock led to improvements in children's human capital investment. The provincial measure τ_p therefore captures the reduced-form local effect of the trade shock, including gains that flow through household and community networks rather than individual labor markets alone.

Figure 2 illustrates the variation in both measures. Panel (a) shows the distribution of US tariff rates before and after the BTA across two-digit industries using data from McCaig (2011). The tariff reductions were large. The median Column 2 rate was 17.3 percentage points, falling to a

median MFN rate of just 1.2 percentage points, a near-complete elimination of tariff protection. In log-difference terms, the industry tariff cut τ_j has a median of 0.168, reaching a maximum of 0.438 in the most exposed sector. A parent in the most exposed industry faced a tariff reduction roughly 2.6 times the median, implying substantially different incentives for occupational upgrading and wage growth across the treatment distribution.

Panel (b) maps the resulting province-level exposure. The measure varies by a factor of 3.6 across 61 provinces, from 0.067 in the least exposed to 0.239 in the most exposed (mean 0.089, median 0.081). The geographic pattern is economically intuitive. Above-median exposure is concentrated in southern provinces around Ho Chi Minh City and along the central and southern coast, areas with established export infrastructure and a higher pre-existing share of employment in manufacturing. Northern inland and mountainous provinces, where employment was concentrated in agriculture, show the lowest exposure.

For the main analysis, we convert these continuous measures into binary treatment indicators by splitting at their sample medians. We define $M_i = \mathbf{1}[\tau_{j(m)} > \tilde{\tau}]$, equal to one if the mother's industry tariff cut exceeds the sample median $\tilde{\tau}$, and zero otherwise, and define F_i analogously for the father. We define $P_p = \mathbf{1}[\tau_p > \tilde{\tau}_p]$, equal to one for provinces with above-median shift-share exposure.[13] Binary indicators admit a straightforward interpretation. The treatment effects compare high-exposure and low-exposure groups. Standard two-way fixed effects regressions with a continuous treatment variable can assign difficult-to-interpret and potentially negative weights to heterogeneous treatment effects, and require a substantially stronger parallel trends assumption that rules out selection on potential outcome gains across treatment doses (Callaway et al., 2024). The binary specification identifies a well-defined high-versus-low-dose contrast under the standard parallel trends assumption, without this additional restriction. Section 6 shows that continuous-treatment specifications produce consistent results.

3.2 Household Surveys and Census Data

To study how trade exposure affects sex ratio at birth, the ideal dataset would record each child's sex alongside both parents' industry, employment, and location around conception, allowing us to assign the trade shock separately to mothers, fathers, and provinces. The sex-selection literature typically relies on Demographic and Health Surveys (DHS) or population censuses, but each source introduces an identification failure that prevents separate channel identification.[14] We therefore use the Vietnam Household Living Standards Surveys (VHLSS) (General Statistics Office of Vietnam, 2010), a biennial nationally representative survey spanning 2002 through 2010

[13] Roughly 21 percent of parents work in industries where the Column 2 and MFN rates were both zero, yielding $\tau_j = 0$. These workers, concentrated in non-export service and agricultural sectors, fall in the low-treatment group.

[14] DHS surveys provide birth histories, but they contain limited information on fathers and are conducted only intermittently. Population censuses are even more infrequent and come with a different problem: treatment must be assigned using industry and location at the census date rather than at conception. If workers switch industries or move after the shock, treatment is mismeasured and estimated effects are attenuated toward zero.

that records industry, location, and employment for both parents at each survey wave. Because the VHLSS does not contain complete birth histories, we work with the flow of births in each wave, treating each cross-section as a sample of infants aged 0–1 observed alongside their parents. We supplement the VHLSS with the 1999 Population Census from IPUMS International (Ruggles et al., 2025), whose province-level employment shares serve as weights for the shift-share exposure measure and whose birth records extend the pre-treatment baseline.

The Census provides births from 1997–1999, contributing approximately 49,000 infant observations to the pre-treatment baseline. The VHLSS waves from 2002 through 2010 contribute repeated cross-sections of infants aged 0–1, yielding approximately 17,700 infant-parent observations across five waves. The pre- and post-treatment periods differ across outcomes. For sex selection and fertility, we define post-treatment as birthyear $>$ 2001, so that Census births from 1997–1999 are entirely pre-treatment while children born in 2002 or later are post-treatment. For maternal labor supply, we define post-treatment as survey year $>$ 2002, making the 2002 VHLSS the sole pre-treatment baseline for that outcome and waves from 2004 onward post-treatment. Because the Census does not record work hours or wage employment status, only industry codes, it contributes to the sex selection and fertility analysis but not to the labor supply analysis. For each infant we observe the child's sex and the parents' contemporaneous industry codes and, where available, work hours and job types.

A key measurement challenge in this setting is identifying parent-child relationships within extended family households, where multiple generations often live together and survey relationships are recorded relative to the household head rather than within nuclear families. We develop an algorithm that identifies both couple and parent-child relationships within extended households, validated against known parent-child links in the 2014 and 2016 VHLSS waves, achieving 93 percent accuracy for mothers and 91 percent for fathers. Appendix A.1 details this procedure.

Table 2 presents summary statistics for Census 1999 and the five VHLSS waves. Mothers' baseline characteristics are comparable across sources, with average ages around 27–28 years and education ranging from 6.5 to 7.4 years across waves. The proportion of urban residents rises over time, consistent with Vietnam's urbanization trend. Male birth shares differ between the Census and the VHLSS. Census 1999 records 51.8 percent male births, corresponding to a sex ratio at birth of approximately 107 boys per 100 girls, while the 2002 VHLSS records 53.2 percent (SRB $\approx$ 113). The gap in sex ratios (approximately six SRB points, corresponding to a 1.4 percentage point difference in male birth shares) reflects the different time periods covered. Census births span 1997–1999, while the 2002 VHLSS captures births from 2000 through 2002, years in which the secular rise in sex selection was already underway, consistent with the upward shift documented in Figure 1. Across the 2004–2008 waves, the sex ratio remains elevated (SRB 110–112).

The model predicts that maternal and paternal wages have opposing effects on fertility (Table 1, column 3), yet testing these predictions requires observing multiple births per woman over time. Our main analysis uses infants aged 0–1, capturing at most one birth per woman per wave

Table 2: Summary Statistics for Census 1999 and VHLSS Survey Waves

	Census	VHLSS				
	1999	2002	2004	2006	2008	2010
Male birth	0.518	0.532	0.529	0.524	0.527	0.511
	(0.500)	(0.499)	(0.499)	(0.499)	(0.499)	(0.500)
Mother's age	27.00	28.10	28.39	27.88	28.09	28.24
	(5.05)	(5.19)	(5.29)	(5.27)	(5.13)	(5.04)
Mother's educ. (yrs)	6.46	6.72	7.17	7.42	8.01	8.22
	(3.43)	(3.60)	(3.71)	(3.69)	(3.71)	(3.66)
Ethnic minority	0.201	0.206	0.218	0.215	0.204	0.213
	(0.401)	(0.404)	(0.413)	(0.411)	(0.403)	(0.409)
Urban	0.131	0.175	0.211	0.231	0.266	0.270
	(0.337)	(0.380)	(0.408)	(0.421)	(0.442)	(0.444)
Birth order	1.93	2.07	2.02	1.90	1.88	1.89
	(0.81)	(0.80)	(0.80)	(0.80)	(0.78)	(0.76)
Mother hrs/month	—	138.16	136.16	137.16	138.14	168.15
		(65.99)	(66.24)	(68.50)	(67.77)	(71.86)
Father hrs/month	—	144.79	149.30	150.70	155.21	173.42
		(60.15)	(63.42)	(64.66)	(65.43)	(62.86)
Mother wage earner	—	0.202	0.243	0.275	0.333	0.366
		(0.402)	(0.429)	(0.446)	(0.471)	(0.482)
Father wage earner	—	0.372	0.397	0.420	0.486	0.512
		(0.483)	(0.489)	(0.494)	(0.500)	(0.500)
N	48,897	5,934	3,159	3,034	3,164	2,374

Notes: Standard deviations in parentheses. Census 1999 includes births from 1997–1999. VHLSS waves include children aged 0–1 in each survey year. The 2010 VHLSS wave captures only children born before 2010 because births occurring in 2010 are incompletely recorded in that survey wave. Dashes indicate variables unavailable in Census 1999.

and providing no within-woman fertility history. To supplement the main analysis, we reconstruct annual birth outcomes from the household roster. Because the roster records the birth years of all children currently living with the mother, we can infer in which years she had a birth. For each woman in a given survey year, we create one observation per year from when she turned 20 until the survey date, coding the outcome as one if a child in the roster has the corresponding birth year. The resulting sample contains approximately two million woman-year observations and supports mother fixed effects, which absorb time-invariant fertility preferences. The main limitation of the retrospective measure is that children who have left the household through marriage, migration, or death are absent from the roster, so earlier births are systematically missing for older women.

The measure nonetheless provides a useful test of the fertility channel using the same treatment variation as the main analysis.

While migration across provinces does not contaminate the industry-level variation we exploit, workers may nonetheless switch industries in response to the trade shock, so that their current industry in the cross-section reflects post-shock sorting rather than pre-shock exposure. To address this selection problem, we exploit the VHLSS's rotating panel feature, which tracks a subset of households across three consecutive waves (2002–2004–2006). From these households we draw women aged 20–40 as our analysis sample. By fixing each parent's industry affiliation at the 2001 baseline, before the BTA could have affected labor market sorting, we isolate trade exposure from compositional changes driven by post-shock worker reallocation. The panel structure also allows us to control for time-invariant unobservables, including fertility preferences and son preference, through individual fixed effects. Appendix Table C1 compares the panel sub-sample to the main cross-section at the 2001 baseline. Mothers' age, education, minority status, urban share, and province-level tariff exposure are closely matched across the two samples, suggesting that selection into the rotating panel is unlikely to explain differences between panel and cross-sectional estimates. The panel sub-sample is smaller (879 versus 5,934) because it requires women to have had at least one birth before the BTA shock and at least one birth after, so that within-mother comparisons can identify the treatment effect. Given this requirement and the short panel length, we rely on the cross-sectional variation for our main analysis and use the panel estimates as a robustness check.

4 Empirical Strategy

Testing the comparative statics in Table 1 requires an empirical design that separately identifies changes in maternal wages, paternal wages, and household non-labor income from a single trade shock. The binary treatment indicators M_i, F_i, and P_p defined in Section 3.1 provide this variation, with each indicator mapping to a distinct row of Table 1 and proxying for a different channel of the BTA's influence on household decisions. We embed all three indicators in a difference-in-differences framework that exploits cross-sectional variation in exposure generated by the BTA's uneven tariff cuts across industries and provinces, comparing outcomes before and after the December 2001 implementation. Estimating all three channels jointly allows each coefficient to be identified conditional on the others, so that the resulting pattern across outcomes can discriminate among the competing mechanisms summarized in Table 1.

Following Section 3.2, we set $\text{Post}_t = \mathbf{1}[\text{birthyear} > 2001]$ for sex selection and fertility, and $\text{Post}_t = \mathbf{1}[\text{survey year} > 2002]$ for labor market outcomes. Because the BTA took effect in December 2001 and gestation takes approximately nine months, births in 2001 and early 2002 largely reflect pre-agreement conceptions, so the pre-treatment baseline extends into the formally post-treatment window.

Our baseline specification estimates a linear probability model for sex selection and an OLS regression for maternal labor supply, both taking the form:

$$y_{ipt} = \beta_1(M_i \times \text{Post}_t) + \beta_2(F_i \times \text{Post}_t) + \beta_3(P_p \times \text{Post}_t) + \alpha_t + \delta_p + \gamma_{j(m)} + \gamma_{j(d)} + X'_{ipt}\phi + \varepsilon_{ipt}. \quad (11)$$

For sex selection, y_{ipt} is an indicator equal to one if child i born in province p in year t is male; for maternal labor supply, it is the log of the mother's monthly work hours. The specification includes year fixed effects α_t, province fixed effects δ_p, and ISIC industry fixed effects for both the mother $\gamma_{j(m)}$ and the father $\gamma_{j(d)}$. The control vector X'_{ipt} includes the mother's age, education, urban residence, minority status, and birth parity, as well as interactions of pre-BTA provincial employment shares in agriculture and manufacturing with year fixed effects, which absorb differential trends arising from preexisting sectoral composition. Standard errors are clustered at the province level.[15] The coefficients β_1, β_2, and β_3 capture the post-BTA differential between high- and low-exposure groups, net of all fixed effects listed above.

For fertility, the specification replaces province and industry fixed effects with mother fixed effects. The retrospective birth panel described in Section 3.2 yields approximately two million woman-year observations, with the outcome AnyBirth_{mt} equal to one if mother m had a birth in year t:

$$\text{AnyBirth}_{mt} = \beta_1(M_m \times \text{Post}_t) + \beta_2(F_m \times \text{Post}_t) + \beta_3(P_p \times \text{Post}_t) + \alpha_t + \mu_m + X'_{mt}\phi + \varepsilon_{mt}. \quad (12)$$

Mother fixed effects μ_m absorb all time-invariant heterogeneity, including underlying son preference, baseline fertility ideals, and permanent differences in industry and location. The control vector X'_{mt} includes the mother's age at birth and interactions of pre-BTA provincial employment shares in agriculture and manufacturing with year fixed effects. Identification comes from within-mother variation in birth propensity before and after the BTA, comparing mothers with different levels of exposure.

Two threats to the cross-sectional estimates motivate the panel specification. The primary concern is industry switching. The cross-section records each parent's current industry rather than pre-shock affiliation, so workers who move into high-tariff-cut industries after the BTA can contaminate M_i and F_i with post-shock selection. Migration across provinces is a secondary concern. If the trade shock draws households with stronger son preference into high-exposure provinces, the composition of the high-P_p group shifts in a way that mechanically inflates β_3.

To address these concerns, we exploit the VHLSS rotating panel described in Section 3.2 and lock each parent's treatment status to her pre-BTA industry affiliation, following McCaig and

[15] All specifications use 61 province clusters. Section 6 reports robustness to clustering at the industry level and to two-way clustering by province and industry.

Pavcnik (2018):

$$y_{it} = \beta_1(M_i \times \text{Post}_t) + \beta_2(F_i \times \text{Post}_t) + \beta_3(P_p \times \text{Post}_t) + \alpha_t + \mu_i + X_{it}'\phi + \varepsilon_{it}. \tag{13}$$

In the panel specification, M_i and F_i are assigned based on each parent's industry reported in the 2002 VHLSS wave, the earliest available survey. Because the 2002 wave was fielded within months of the December 2001 BTA implementation and the VHLSS industry question refers to the respondent's current or most recent occupation, the industry recorded in 2002 effectively captures pre-BTA affiliation before meaningful occupational reallocation could have occurred. Fixing exposure at this initial wave ensures that post-BTA industry switching cannot alter assigned treatment. Individual fixed effects μ_i absorb all time-invariant heterogeneity, including son preference and baseline fertility ideals, and the design mirrors an intent-to-treat experiment. The panel sample consists of women aged 20–40 observed across the 2002, 2004, and 2006 VHLSS waves; for sex selection and labor supply, we restrict to women who gave birth both before and after the BTA, so that within-woman comparisons can identify the treatment effect. If endogenous sorting were inflating the cross-sectional estimates, fixing treatment to pre-shock industries should attenuate the panel coefficients relative to the cross-sectional results.

The key identifying assumption is that, absent the BTA, outcomes would have evolved similarly across high- and low-exposure groups. Parallel trends is untestable in the post-reform period, but we can assess it directly in the pre-reform data. The 1999 Census births (1997–1999) and VHLSS births in 2000–2001 provide pre-treatment observations spanning multiple years before the shock. The event study in Section 5 shows no differential trends between high- and low-exposure groups across these years, supporting the assumption. For the panel specification, we additionally assume that women's 2001 industry choices were not made in anticipation of the BTA's differential tariff effects across sectors.

The coefficients β_1, β_2, and β_3 map directly to the comparative statics in Table 1. The maternal time-cost mechanism predicts $\beta_1 > 0$ for sex selection and maternal labor supply and $\beta_1 < 0$ for fertility, because a rise in the mother's wage tightens her time constraint and intensifies the tradeoff between work and childcare. Paternal exposure (β_2) and provincial income (β_3) relax the budget constraint without altering the mother's opportunity cost of childrearing, so the model predicts $\beta_2 = \beta_3 = 0$ for sex selection while $\beta_2 > 0$ and $\beta_3 > 0$ for fertility. As Table 1 summarizes, the joint pattern $\beta_1 > 0$, $\beta_2 = 0$, $\beta_3 = 0$ for sex selection alongside $\beta_1 < 0$, $\beta_2 > 0$, $\beta_3 > 0$ for fertility is uniquely consistent with the maternal time-cost mechanism, and distinguishes it from the income-based, bargaining, and relative-returns alternatives discussed in Section 2.

5 Main Results

5.1 Effects on Sex Selection

Mothers in high-exposure industries are significantly more likely to have sons after the BTA, consistent with the time-cost prediction in column (1) of Table 1. Table 3 reports estimates from progressively richer specifications of Equation (11), where the dependent variable is an indicator equal to one if the child is male.

Columns (1) and (2) introduce the maternal and paternal exposure channels individually, while columns (3) and (4) combine them, with column (4) adding the full control set. Columns (5) through (7) introduce and combine the province-level channel. When entered alone in column (1), maternal industry exposure ($M_i \times \text{Post}_t$) raises male birth probability by 6.5 percentage points (s.e. 0.025), significant at the 5 percent level. Against a dependent variable mean of 0.52, this corresponds to a 13 percent relative increase in the probability that a birth is male. For a mother whose industry received the sample median tariff cut of approximately 7 percentage points, the continuous treatment specification in Appendix Table C7 implies a 1.4 percentage point increase in male birth probability, raising the sex ratio from approximately 108 to 114 boys per 100 girls. This magnitude is consistent with the aggregate trend in Figure 1, which shows Vietnam's sex ratio rising from approximately 107 to 112 boys per 100 girls over the study period. The comparison is informative because it uses a setting-representative tariff cut: approximately half of employed mothers in our sample received a cut at or above the median, so the median-exposure estimate captures the experience of a typical rather than marginal worker.

The father's industry exposure ($F_i \times \text{Post}_t$) enters with a positive but imprecisely estimated coefficient of 0.04 in column (2). In the combined specifications of columns (3) through (7), the paternal coefficient falls to 0.02 and remains statistically indistinguishable from zero, consistent with the model prediction that $\beta_2 = 0$, since paternal wages raise household income without tightening the maternal time constraint. Across these same specifications, the maternal coefficient moves only modestly, from 0.065 in column (1) to 0.058 when the father is added and recovering to 0.065 with the full control set in column (4). The stability of the maternal coefficient across specifications suggests that the estimate is not driven by omitted correlates of industry affiliation or demographic composition.

Province-level exposure ($P_p \times \text{Post}_t$) is near zero and insignificant throughout, entering at 0.013 in column (5) and 0.018 in column (7). The provincial measure is a noisy aggregate that bundles local labor demand, income effects, and infrastructure changes affecting all households in a province simultaneously. The null is therefore unsurprising and difficult to interpret on its own. The sharper test of whether income alone drives sex selection comes from the paternal channel, which isolates household-level income gains without tightening the maternal time constraint. The paternal null on sex selection, combined with the significant maternal effect, is the pattern that distinguishes the time-cost mechanism from income-based alternatives (Section 7).

Table 3: Effects of BTA Trade Exposure on Male Birth Probability

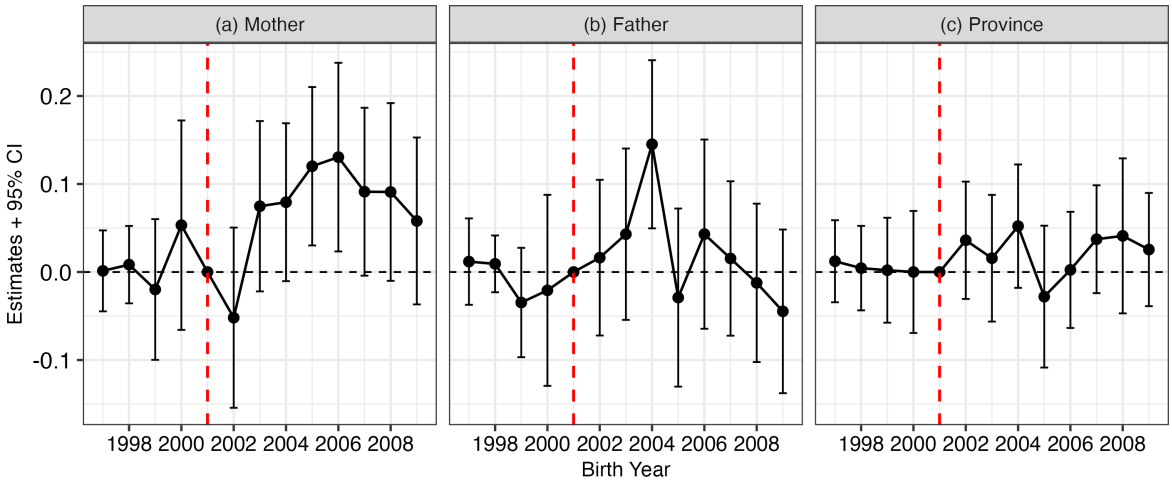
	MaleBirth						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$M_i \times$ Post	0.065**		0.058**	0.065**		0.057**	0.063**
	(0.025)		(0.025)	(0.025)		(0.024)	(0.024)
$F_i \times$ Post		0.040	0.020	0.023		0.020	0.021
		(0.031)	(0.031)	(0.031)		(0.031)	(0.032)
$P_p \times$ Post					0.013	0.003	0.018
					(0.013)	(0.012)	(0.017)
R^2	0.015	0.015	0.015	0.016	0.014	0.015	0.016
Observations	63,894	63,894	63,894	63,894	63,894	63,894	63,894
Dep. Var. Mean	0.518	0.518	0.518	0.518	0.518	0.518	0.518
Controls				✓			✓

Notes: The dependent variable is an indicator equal to 1 if the child is male. M_i equals 1 if the mother's 2-digit industry received a tariff cut above the sample median (maternal high-exposure). F_i equals 1 if the father's 2-digit industry received a tariff cut above the sample median (paternal high-exposure). P_p equals 1 if province p's employment-weighted average tariff cut exceeds the cross-province median (province high-exposure). All specifications include birth year, province, and parents' 2-digit ISIC fixed effects. Controls in columns (4) and (7) include mother's age, education, urban location, minority status, birth parity, and pre-BTA provincial employment shares in agriculture and manufacturing interacted with survey-year dummies. Standard errors clustered at the province level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Column (7) presents the fully controlled specification with all three exposure channels entering simultaneously. The maternal coefficient of 0.063 ($p < 0.05$) is our preferred estimate. For a more interpretable magnitude, the continuous specification in Appendix Table C7 implies a 1.4 percentage point increase in male birth probability for a mother whose industry received the sample median tariff cut of approximately 7 percentage points, the appropriate basis for cross-study comparison. Anukriti (2018) finds that first births become 0.5 to 1.2 percentage points more likely to be male after the Devi Rupak fertility incentive in India, and Almond et al. (2019) estimate a cumulative increase of approximately 2.8 percentage points over four years for second births following a firstborn girl in China. Our population-average estimate of 1.4 percentage points falls within the range documented by Anukriti (2018) and below the benchmark from Almond et al. (2019), consistent with the fact that the latter conditions on the subgroup facing the strongest son preference pressure while our estimate pools across all birth orders and sibling compositions. The heterogeneity analysis below explores whether the maternal exposure effect varies across subgroups by demographic characteristics.

Figure 3 traces these effects over time, plotting event study estimates from the full specification where all three exposure channels enter simultaneously with the complete set of controls.

Figure 3: Event Study Estimates of BTA Exposure on Male Birth Probability



Notes: Each panel plots event study coefficients from a single regression of Equation (11) that includes all three exposure channels simultaneously with the full set of controls. Panel (a) shows maternal industry exposure (M_i), panel (b) shows paternal industry exposure (F_i), and panel (c) shows province-level exposure (P_p). Exposure measures are binary indicators equal to one if the relevant tariff cut exceeds the sample median. The omitted reference year is 2001. The vertical dashed line marks the BTA implementation in late 2001. Error bars represent 95% confidence intervals based on standard errors clustered at the province level.

Appendix Figure C2 presents individual event studies for each channel estimated separately, confirming that the patterns are not artifacts of joint estimation.

Panel (a) displays the maternal exposure coefficients. Pre-BTA estimates for the 1997–1999 Census cohorts and the 2000–2001 VHLSS births cluster tightly around zero, confirming that high- and low-exposure groups followed parallel trends well before the shock.[16] The coefficients remain indistinguishable from zero in 2002 and 2003. The BTA took effect in December 2001, so most births recorded in 2002 reflect conceptions that predate the agreement, and the 2002 VHLSS recall period straddles the policy change. The null coefficients through 2003 therefore extend the parallel trends validation into the post-announcement period. From 2004 onward, the maternal coefficients become positive and remain elevated. The two-to-three year delay between the BTA and the onset of effects is consistent with the time-cost channel. Although US tariff cuts took effect immediately, firms needed time to expand hiring in treated industries and women needed to observe the resulting wage improvement before adjusting conception timing. The nine-month gestation lag then pushes affected births into the 2004 wave and beyond.

[16] A reverse causality concern could arise because the design is cross-sectional. The mother's industry is measured contemporaneously rather than before the birth, so a mother who had a son could in principle sort into a high-treatment industry afterward. Two pieces of evidence rule this out. First, the flat pre-BTA coefficients show no pre-existing differential trend between high- and low-exposure groups, ruling out pre-existing compositional differences correlated with treatment. Second, if having a son led parents to work more and sort into higher-wage industries, the same mechanism should produce a positive paternal coefficient, but panel (b) shows no such effect.

Panels (b) and (c) display paternal and province-level exposure. Neither channel exhibits a consistent post-BTA pattern. Paternal industry exposure fluctuates around zero in most years, with the exception of an isolated 2004 spike that is a sampling artifact.[17] Province-level exposure is statistically insignificant throughout the study period. Pre-trend coefficients are flat for both channels, ruling out the concern that differential pre-existing trends could explain the maternal results through correlated shocks. The contrast across the three panels is the central identification result. The maternal time-cost mechanism predicts $\beta_1 > 0$ with $\beta_2 \approx 0$ and $\beta_3 \approx 0$ for sex selection (Table 1, column 1). The event study confirms that the positive pattern is unique to maternal exposure and emerges only after the BTA could plausibly have affected labor market conditions. Because the main difference-in-differences specification defines Post as birth year after 2001, the pooled estimate includes 2002–2003 births whose conceptions largely predate the agreement. The null event-study coefficients for these transition cohorts imply that the reported DiD estimate is a conservative average that includes cohorts not yet affected by the BTA.

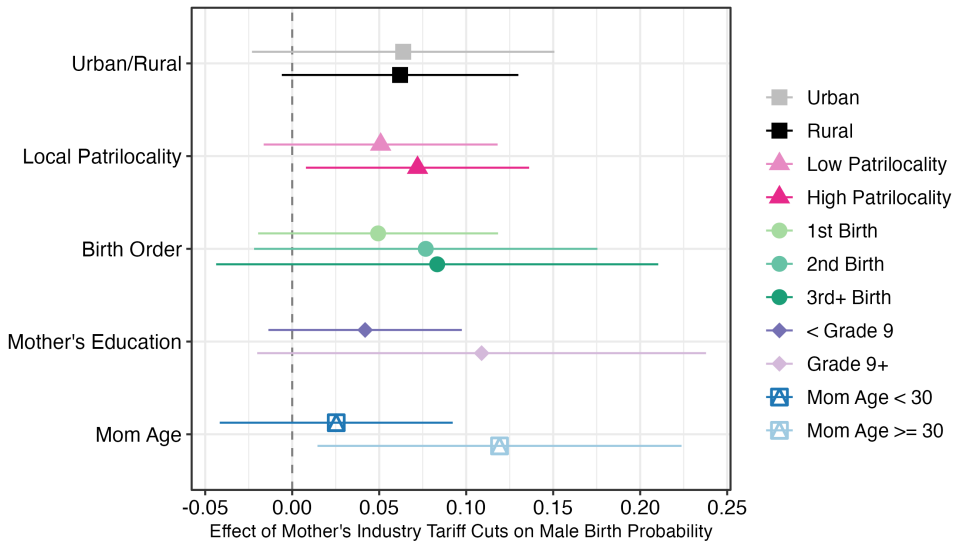
Heterogeneity. We examine how the maternal exposure effect varies across subgroups by estimating Equation (11) separately within each group. Figure 4 displays the coefficient on $M_i \times \text{Post}_t$ and its 95 percent confidence interval for each subgroup.

The maternal exposure effect is similar in rural and urban areas, with point estimates of 0.062 and 0.064, respectively, and neither is statistically distinguishable from zero at conventional levels. The absence of an urban-rural gradient is informative. If the mechanism operated primarily through differential access to sex-selective technology, which is more concentrated in cities, we would expect a larger urban effect. The near-equality instead points to a time-cost channel operating through industry-level labor demand rather than through geographic access to medical services.

The effect concentrates in provinces with high patrilocality. We classify provinces above the median share of couples residing with the husband's parents as high patrilocality, using the 1999 Census. This residence pattern has traditionally reinforced son preference by linking property inheritance and old-age support to male heirs (den Boer and Hudson, 2017; Ebenstein, 2014). The point estimate in high-patrilocality provinces is 0.068 and statistically significant ($p = 0.028$), while the estimate in low-patrilocality provinces is 0.051 and imprecise. The stronger effect in patrilocal areas is consistent with the time-cost mechanism interacting with underlying son preference. Where son preference is more institutionally entrenched, the marginal value of sex selection is higher when the opportunity cost of childcare rises. This finding points to a pattern of particular importance in the literature on sex selection. The cultural intensity of son preference and the current household composition jointly determine how strongly economic incentives translate

[17] The high-father-treatment group for 2004 births comprises 198 observations spread across 146 province–industry cells, with 97 percent of cells containing at most 3 observations. The all-male draws in these tiny cells produce the spike. Adjacent years with similar cell structures do not show the same pattern, and the individual event study for the paternal channel confirms the spike is not an artifact of joint estimation.

Figure 4: Heterogeneous Effects of Maternal Industry Exposure on Male Birth Probability



Notes: Each point estimate and its 95 percent confidence interval represent the coefficient on $M_i \times \text{Post}_t$ from a separate regression of Equation (11) estimated within the indicated subgroup. Subgroups are defined by urban/rural residence, high/low patrilocality (above/below median share of couples residing with the husband's parents, from the 1999 Census), birth order (1st, 2nd, 3rd+), maternal education (below or at/above grade 9), and maternal age at birth (below 30 or 30 and above). All regressions include the full set of controls and fixed effects as specified in Equation (11).

into sex selection. The patrilocality dimension captures the former, and the birth order gradient discussed below captures the latter.

Birth order shows a monotone gradient, with the coefficient on maternal industry exposure rising from 0.049 at first parity to 0.077 at second parity and 0.083 at third and higher parities.[18] None of the parity-specific estimates is statistically significant at conventional levels, but the monotone pattern is directionally consistent with the time-cost mechanism. Higher-parity parents have fewer remaining births over which to fulfill son preference, so the marginal incentive to sex select at each birth is larger.

Mothers with higher education (ten or more years of schooling) show a larger point estimate (0.109) than those with less schooling (0.042), with the high-education estimate marginally significant ($p = 0.103$) and about 2.6 times as large. More educated mothers typically command higher wages in the labor market, so a given improvement in industry conditions implies a larger increase in their opportunity cost of time and a stronger incentive to economize on time-intensive childcare through sex selection. Neither estimate reaches the 5 percent threshold, so the education gradient should be read as suggestive rather than definitive.

Maternal age at birth generates the sharpest heterogeneity. Mothers aged 30 and above show

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

a large and significant effect (0.119, $p = 0.029$), while mothers under 30 show no discernible response (0.025, $p = 0.462$). Older mothers are more likely to be at higher parities without yet having a son, have accumulated work experience and face a higher opportunity cost of childcare, and confront a shorter remaining fertility horizon. Each of these factors reinforces the incentive to sex select and the pressure to produce a son before the fertility window closes.

The overall pattern is broadly consistent with the time-cost mechanism. Significant effects emerge precisely in the subgroups where institutional son preference and the opportunity cost of childcare are jointly highest, namely high-patrilocality provinces and older mothers facing a narrowing fertility horizon. Where neither condition is particularly acute, the estimates are imprecise and no clear concentration is detectable.

5.2 Effects on Maternal Labor Market Outcomes

The BTA increased both the probability of maternal wage employment and the intensity of work, consistent with the time-cost prediction (Table 1, column 2, row 1). McCaig and Pavcnik (2018) document that the BTA increased formal-sector employment across the broader Vietnamese workforce. Table 4 shows that the same shift holds for the specific subgroup of mothers with young children in our data, with columns (1)–(3) reporting effects on wage employment probability and columns (4)–(6) on log monthly hours worked.[19]

On the extensive margin, column (1) shows that mothers in high-tariff-cut industries are 16 percentage points more likely to hold a wage job after the BTA ($p < 0.01$), from a baseline of 25 percent. The maternal exposure effect is stable at 0.16 across columns (2) and (3) as additional exposure channels and controls are added. Paternal and province-level exposure effects are small and statistically insignificant once the maternal channel is included. The wage employment result also provides evidence for the model's underlying premise that the BTA raised maternal wages. Because only one in four mothers held a wage job before the BTA, measuring wage levels directly would leave the majority of the sample, who work in agriculture or self-employment, without an observed market wage. The transition into wage employment from informal or agricultural work is itself a signal of improved labor market conditions, and the shift from the informal to the formal sector is precisely the reallocation that McCaig and Pavcnik (2018) document at the aggregate level. The increase in wage employment probability therefore confirms that the BTA genuinely improved mothers' labor market position, validating the premise on which the time-cost mechanism rests.

On the intensive margin, column (4) shows that mothers in the same high-tariff-cut industries work 0.13 log-points more monthly hours after the BTA, significant at the 1 percent level. Exponentiating gives a 14 percent increase in monthly hours. With the sample mean at approximately

[19] The sample of mothers here is smaller than the sample 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 leave, she will drop out of the hours sample, although her child still appears in the infant sample.

Table 4: Effects of BTA Exposures on Maternal Labor Market Outcomes

	Wage Employment			Log Monthly Hours		
	(1)	(2)	(3)	(4)	(5)	(6)
$M_i \times$ Post	0.16***	0.16***	0.16***	0.13***	0.11**	0.10**
	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)
$F_i \times$ Post		0.004	0.004		0.01	0.01
		(0.02)	(0.01)		(0.03)	(0.03)
$P_p \times$ Post			−0.001			0.06
			(0.02)			(0.04)
R^2	0.60	0.61	0.61	0.25	0.25	0.25
Observations	17,389	17,389	17,389	17,198	17,198	17,198
Dep. Var. Mean	0.25	0.25	0.25	4.7	4.7	4.7
Controls		✓	✓		✓	✓

Notes: The dependent variable in columns (1)–(3) is an indicator equal to 1 if the mother has any wage employment. In columns (4)–(6) it is the log of monthly hours worked in her main job. Controls in columns (2), (3), (5), and (6) match those in Table 3. All specifications include survey-year, province, and parents' 2-digit ISIC fixed effects. Standard errors, clustered at the province level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

110 hours per month, this effect translates into roughly 15 additional hours per month. The estimate attenuates only modestly, falling to 0.10 log-points (11 percent, or about 12 hours per month) in column (6) once all exposure channels and controls are included. The direction and stability of the maternal coefficient are consistent with the model prediction that higher maternal wages increase labor supply (Table 1, column 2, row 1).

Paternal industry exposure ($F_i \times \text{Post}_t$) on hours is near zero once maternal exposure is included, remaining statistically insignificant through column (6). The model predicts a negative income effect of paternal wages on maternal labor supply ($\beta_2 < 0$ in Table 1, column 2, row 2), but the point estimate provides no support for a quantitatively important income channel on maternal hours. The null finding is consistent with two offsetting forces operating through the mother's time constraint $h_m = 1 - \ell_m - \gamma n$. Paternal wages raise maternal leisure through an income effect (reducing h_m) and raise fertility (confirmed empirically at $+0.5$ percentage points in Table 5), which increases childcare time γn and further reduces h_m. Both forces push in the same direction, but the magnitudes are small because paternal exposure affects the mother's time allocation only indirectly through household income. The null estimate is therefore consistent with a small negative effect that the data lacks power to detect.

Similarly, province-level exposure ($P_p \times \text{Post}_t$) enters at 0.06 log-points in column (6) but falls below conventional significance levels. The attenuation after adding controls suggests that some of the provincial association reflects differential trends in economic structure rather than a direct

Table 5: Effects of BTA Exposures on Fertility (Probability of Giving Birth)

	AnyBirth					
	(1)	(2)	(3)	(4)	(5)	(6)
$M_i \times$ Post	-0.004***		-0.006***	-0.006***	-0.006***	-0.006***
	(0.001)		(0.001)	(0.001)	(0.001)	(0.001)
$F_i \times$ Post		0.003*	0.005***	0.005***	0.005***	0.005***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$P_p \times$ Post					-0.0001	-0.002
					(0.002)	(0.002)
R^2	0.29	0.29	0.29	0.29	0.29	0.29
Observations	2,011,963	2,011,963	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
Controls				✓		✓

Notes: The dependent variable is an indicator equal to 1 if the mother had any birth in the past 12 months. Columns show progressive addition of treatment variables and controls. Controls in columns (4) and (5) include mother's age at birth and province-level manufacturing and agricultural employment shares from the 1999 census, each interacted with birth-year dummies. All specifications include mother and birth-year fixed effects. Standard errors, clustered at the province level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

labor demand channel.

The model also has predictions for fathers. Paternal wages generate an income effect in the household but do not tighten the maternal time constraint, so paternal industry exposure should drive paternal labor supply without affecting sex selection. Appendix Table C2 examines the same specifications for fathers. Maternal industry exposure ($M_i \times \text{Post}_t$) is near zero for both of father's outcomes and falls further toward zero once paternal exposure is included. In contrast, paternal industry exposure ($F_i \times \text{Post}_t$) enters at 11 percentage points for wage employment and 0.09 log-points for monthly hours, both significant at the 1 percent level. The pattern confirms that the BTA affected each parent's labor supply through their own industry conditions, reinforcing that the sex selection response operated through the maternal time-cost channel rather than through a generalized household income effect. Taken together, the labor market evidence provides reassurance that the industry-level trade shocks are valid proxies for actual changes in labor market conditions, and that the sex selection results documented earlier reflect genuine shifts in maternal opportunity costs rather than spurious correlations between industry tariff cuts and fertility behavior.

5.3 Effects on Fertility

The same time-cost mechanism that drives sex selection and maternal labor supply also predicts reduced fertility. When the opportunity cost of time rises, mothers reduce childbearing. Table 5 reports estimates from Equation (12), where the dependent variable is an indicator equal to one if the mother had any birth in the past 12 months. The sample is a retrospective panel of woman-year observations with mother fixed effects, so identification comes from within-mother changes in birth propensity before and after the BTA.

Maternal industry exposure reduces the probability of giving birth. Column (1) yields a coefficient of -0.004 ($p < 0.01$) when maternal exposure enters alone, and the estimate stabilizes at -0.006 ($p < 0.01$) across all combined specifications in columns (3) through (5), confirming that the result does not depend on controlling for paternal or provincial income channels. Because the dependent variable mean is only 0.02, a reduction of 0.6 percentage points represents a 30 percent decline in the annual birth probability. Expressed differently, a mother in the high-exposure group would have roughly 0.06 fewer children over a ten-year horizon relative to a comparable mother in the low-exposure group. This absolute magnitude is consistent with the demographic scale of the sex selection response documented above. The treatment combines an extensive-margin shift into wage employment and an intensive-margin increase in hours worked (Table 4), both of which create direct time conflicts with childbearing. The sign and magnitude are consistent with the model prediction that higher maternal wages reduce fertility by raising the time cost of children (Table 1, column 3, row 1).

Paternal industry exposure moves in the opposite direction. The father's coefficient is 0.003 ($p < 0.10$) when entered alone in column (2) and rises to 0.005 ($p < 0.01$) once maternal exposure is included in columns (3) through (5). The 0.005 estimate implies a 25 percent increase in annual birth probability relative to the mean. The positive sign is consistent with the model prediction that paternal wages raise fertility through a pure income effect without tightening the maternal time constraint (Table 1, column 3, row 2). The increase from column (2) to column (3) reflects that maternal and paternal exposures are positively correlated, so omitting maternal exposure attenuates the paternal estimate downward. The distinction between the two channels is thus not a theoretical assumption but an empirical regularity. Once each parent's exposure is held constant, the signs diverge sharply and in the direction the model predicts.

Province-level exposure ($P_p \times \text{Post}_t$) is near zero and statistically insignificant in column (5). This result is imperfectly consistent with the model. If the provincial Bartik variable cleanly captured non-labor income, the model would predict a positive coefficient through the income effect, paralleling the positive paternal estimate. The near-zero result is therefore not a prediction success. The provincial measure is a noisy aggregate that conflates local labor demand shocks with income effects and affects all households in a province simultaneously, including those without direct industry exposure. It does not map onto the household-level non-labor income channel the model has in mind, and we do not interpret the null as a sharp test of that channel. The more in-

formative test of the income effect is the paternal exposure estimate in column (3), which isolates household-level income variation without tightening the maternal time constraint. The provincial null does, at a minimum, rule out the possibility that broad local business cycle movements are driving the fertility result.

The opposing signs of maternal and paternal exposure on fertility, combined with the null for provincial exposure, are difficult to reconcile with a model where any positive income shock uniformly affects childbearing in the same direction. The pattern instead supports the structural distinction at the core of our framework. Maternal wages tighten the time constraint, reducing births and raising the incentive for sex selection. Paternal wages relax the budget constraint, increasing births without affecting the sex composition of children.

One qualification is that the 12-month birth indicator captures annual birth hazard rather than completed fertility, so a portion of the estimated reduction may reflect birth timing rather than a permanent decrease in desired family size. The heterogeneity estimates in Figure 4, which show concentrated effects among older mothers facing a narrowing fertility window where postponement is least plausible, provide evidence against a pure timing interpretation. The fertility results close the loop on the time-cost mechanism, showing that the same industry-level shock that increases maternal labor supply simultaneously reduces births, as the model predicts.[20]

5.4 Panel Evidence

A remaining concern with the cross-sectional estimates is that women with stronger son preference may have sorted into high-tariff-cut industries after the BTA, inflating the estimated maternal exposure coefficient. Table 6 addresses this concern by moving from the cross-sectional specification to a within-mother panel design. The progression is informative but not mechanical: column (2) changes both the sample and treatment assignment, column (3) replaces location and industry fixed effects with mother fixed effects, and columns (4)–(5) add additional channels and composition controls. We interpret the panel results as reducing concern that time-invariant son preference or contemporaneous industry switching explains the cross-sectional estimate, while recognizing that they do not by themselves eliminate dynamic selection or other time-varying confounding.

Column (1) estimates the maternal exposure effect on the 2002–2006 cross-sectional sample. The maternal coefficient is 0.06 ($p < 0.05$), matching the full-sample estimate in Table 3. Column (2) then makes two substantive changes at once: it restricts the sample to panel-eligible mothers and assigns treatment based on the mother's 2002-wave industry rather than her contemporaneous industry. Among women working in 2002, approximately 23% changed their 2-digit ISIC industry by 2004 or 2006, so freezing treatment at the 2002 assignment reduces concern that

[20] Appendix Tables C10 and C11 report the same specifications with standard errors clustered at the mother level rather than the province level, to account for within-mother serial correlation in the retrospective panel. The results are unchanged.

Table 6: Panel Evidence on Within-Mother Identification

	MaleBirth				
	(1)	(2)	(3)	(4)	(5)
$M_i \times$ Post	0.06**	0.17**	0.34***	0.30***	0.17
	(0.03)	(0.07)	(0.08)	(0.11)	(0.11)
$F_i \times$ Post				0.06	0.07
				(0.16)	(0.14)
$P_p \times$ Post				0.03	0.02
				(0.06)	(0.06)
Prior son					-0.55***
					(0.06)
R^2	0.01	0.09	0.83	0.82	0.86
Observations	12,127	2,684	2,788	2,615	2,615
Mothers		1,723	1,723	1,723	1,723
Dep. var. mean	0.53	0.52	0.52	0.53	0.53
Individual FE			✓	✓	✓
Industry FE	✓	✓			
Province FE	✓	✓			

Notes: All columns include birth year fixed effects and controls for age, education, urban residence, and minority status, with province-level clustering. Column (1) uses the full 2002–2006 cross-section with treatment assigned from the mother's contemporaneous industry. Columns (2)–(5) restrict to panel-eligible mothers with treatment assigned from the mother's 2002-wave industry. Column (3) replaces province and industry fixed effects with individual mother fixed effects. Paternal and provincial channels enter from column (4). Column (5) adds a prior-son indicator. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

contemporaneous post-BTA occupation changes mechanically generate the treatment effect. The maternal coefficient rises to 0.17 ($p < 0.05$). Table C1 shows that panel-eligible mothers are broadly similar to the broader cross-section on education, ethnicity, urban residence, and province-level tariff exposure, with the main difference being that they are about one year older on average. We therefore view column (2) as moving the design closer to pre-determined treatment assignment while staying close to the main sample on observables.

Column (3) replaces province and industry fixed effects with individual mother fixed effects, absorbing time-invariant maternal heterogeneity including baseline son preference and permanent industry sorting. The maternal coefficient rises to 0.34 ($p < 0.01$). This larger within-mother estimate is consistent with attenuation in the cross-section rather than amplification by simple sorting. Three features of the panel sample help explain the magnitude. First, the requirement that mothers have births both before and after 2001 selects high-fertility women facing a narrow-

ing fertility window, a group for whom the heterogeneity analysis (Figure 4) already shows larger treatment effects. Second, if women with lower baseline son preference sorted into export industries, the cross-sectional estimate is attenuated by this selection, and mother fixed effects remove the attenuation. Third, the panel sample is small, so individual mothers carry substantial leverage and the estimate is informative about direction but imprecise about magnitude. At the same time, the comparison does not by itself rule out dynamic selection, changes in which births are observed over time, or other time-varying confounding within the panel sample. The increase in observations relative to column (2) reflects that the mother fixed-effects specification no longer requires a contemporaneous industry fixed-effect code for each birth record. We therefore read column (3) as strengthening the case against a narrow sorting story, not as a definitive causal resolution.

Column (4) adds paternal and provincial exposure channels and demographic controls. The maternal coefficient is essentially unchanged at 0.30 ($p < 0.01$). The stability confirms that M_i does not proxy for paternal or provincial income shocks. The paternal and provincial channels are individually small and statistically insignificant, consistent with the main results in Table 3.

Column (5) adds a prior-son indicator. The prior-son coefficient is large, negative, and precisely estimated at -0.55 ($p < 0.01$), confirming that prior family composition is a powerful predictor of subsequent male-birth probability. The maternal exposure coefficient falls to 0.17 and loses statistical significance. This attenuation is descriptively consistent with treatment effects being stronger when the realized sex composition of earlier births is less male-biased. We do not interpret column (5) as a clean causal decomposition, however, because prior son status may itself lie on the treatment path for some births. The specification is therefore best read as suggestive composition evidence rather than as a definitive mechanism test. The composition-contingent pattern also helps account for why Almond et al. (2019)'s estimates for second births following a firstborn girl substantially exceed our population-average effects, since their conditioning selects the subgroup facing the strongest composition pressure.

6 Robustness Checks

This section examines the sensitivity of our main results to alternative estimation choices. We consider alternative clustering schemes, region-by-year fixed effects, nonlinear specifications, and continuous treatment measures.

Alternative Clustering. Our baseline specifications cluster standard errors at the province level (61 clusters), the natural choice given that the BTA shock propagates through local labor markets and households. However, because treatment is defined by the tariff cut assigned to each parent's 2-digit ISIC industry, unobserved shocks may also correlate within industries. We address this concern with two alternative clustering schemes. Table C3 clusters by 59 2-digit ISIC industries alone, directly testing whether within-industry error correlation inflates our baseline inference.

Table C4 applies two-way clustering by province and industry simultaneously, the most conservative approach as it accounts for both geographic and sectoral correlation in the error term.

Under ISIC-only clustering (Table C3), the maternal exposure coefficients on log hours (0.10), wage-earner probability (0.16), and fertility (-0.006) all remain significant at the 1 percent level, with standard errors close to the baseline. The maternal effect on sex selection retains its magnitude (0.06) and is significant at the 5 percent level, though the standard error rises from 0.02 to 0.03. With 59 industry clusters comparable to the 61 province clusters, few-cluster imprecision is not a concern. This modest attenuation instead reflects genuine within-industry error correlation arising from industry-level treatment, which is the scenario industry clustering is designed to detect. Under two-way clustering (Table C4), all four maternal exposure effects are significant at the 1 percent level, with standard errors virtually identical to those in the province-only baseline. Across both approaches, the paternal exposure coefficients remain small and insignificant for sex selection and maternal labor outcomes, though the paternal fertility effect (AnyBirth 0.005) retains significance at the 5 percent level. The provincial exposure coefficient for labor supply is also insignificant under two-way clustering, though it reaches marginal significance under ISIC-only clustering alone. The overall pattern is consistent with the baseline conclusion that maternal industry exposure drives the main findings.

Region-Year Fixed Effects. A potential concern is that regional infrastructure improvements could independently affect access to sex-selective technologies (Chen et al., 2013) 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), and shocks to healthcare access or childcare provision may vary across these regions over time. Since these regions correspond to official planning units, region-by-year fixed effects absorb region-level policy variation and differential investment in healthcare or childcare infrastructure. We incorporate region-by-year fixed effects into our baseline specifications to absorb any such time-varying regional confounders.

Table C5 shows that including region-by-year fixed effects leaves all four outcomes essentially unchanged. The maternal exposure coefficients retain their baseline magnitudes and significance levels across sex selection (0.06, $p < 0.05$), hours (0.10, $p < 0.05$), wage employment (0.15, $p < 0.01$), and fertility (-0.006, $p < 0.01$). Paternal and provincial exposure measures remain small or statistically insignificant throughout. The results are unlikely to be driven by omitted regional shocks to infrastructure or healthcare access.

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 C6 reports the logit estimates. The maternal

exposure coefficient ($M_i \times$ Post) is 0.22 without controls and 0.21 with controls, both significant at the 5 percent level. Given the mean male birth probability of approximately 52 percent, the implied odds ratio is $e^{0.21} \approx 1.23$, indicating a 23 percent increase in the odds of a male birth for mothers in high-exposure industries. The implied average marginal effect of approximately 0.05 is close to the linear probability model estimate of 0.06, confirming that the linear probability approximation is accurate at the sample mean.

The paternal and provincial exposure coefficients remain small and statistically insignificant across both specifications, confirming that maternal industry exposure drives the sex-selection response. The agreement between logit and linear probability model results confirms that the main estimates are not an artifact of the linear functional form.

Continuous Treatment Exposures. Our main analysis uses binary indicators that classify workers as high-exposure or low-exposure based on whether their industry's tariff cut exceeds the sample median. This binary split discards variation in exposure intensity within each group. We now replicate the analysis using continuous tariff-cut measures (τ_m, τ_d, τ_p from Section 3.1), which exploit the full distribution of tariff reductions across industries. Tables C7, C8, and C9 present estimates for sex selection, maternal labor supply, and fertility.

The sex-selection result is robust to the continuous specification. Table C7 shows that the maternal exposure coefficient on male birth ($\tau_m \times$ Post) is positive and significant at the 5 percent level in the specifications where it is included (columns 1 and 3–6), ranging from 0.20 to 0.22. Column (2) isolates the paternal channel and does not include τ_m. Paternal and provincial exposure coefficients are small and insignificant, as in the binary results.

The maternal labor supply results in Table C8 confirm the wage employment effect and are broadly consistent on hours. The maternal exposure coefficient on wage employment is 0.51 ($p < 0.01$), stable across all three columns. The hours effect is positive (0.15 log-points) but is not precisely estimated under continuous treatment, reflecting that labor supply responses may be more concentrated among the highest-exposure workers and thus attenuated when the continuous measure spreads the effect across the full tariff distribution. The direction remains consistent with the binary result. Paternal exposure on wage employment and provincial exposure on hours each show a marginally significant negative coefficient in some specifications, but the effects are modest relative to the maternal channel and do not alter the conclusion that the primary hours response operates through the maternal time-cost channel.

For fertility, Table C9 confirms the negative maternal exposure effect. When maternal and paternal treatments are jointly included, the $\tau_m \times$ Post coefficient is -0.008 ($p < 0.05$), consistent with the binary estimate of -0.006 ($p < 0.01$). The attenuation of the unconditional maternal estimate (-0.003, column 1) toward zero reflects positive correlation between maternal and paternal exposure: paternal exposure raises fertility through an income effect, partially offsetting the maternal time-cost effect in a single-exposure regression. Controlling for paternal exposure sharp-

ens the maternal estimate, producing coefficients between -0.007 and -0.009 that are statistically significant in columns (3)–(6). Paternal exposure ($\tau_d \times$ Post) consistently raises fertility, with coefficients ranging from 0.021 to 0.023 ($p < 0.01$) across all combined specifications, consistent with the binary result that $F_i \times$ Post increases birth probability.

Across all three outcomes, the continuous specifications confirm that each parent's labor supply responds to their own industry conditions, and that the sex-selection effect operates through the maternal time-cost mechanism rather than a generalized household income channel. These continuous-treatment estimates should nonetheless be interpreted with caution: Callaway et al. (2024) show that standard TWFE regressions with a continuous treatment variable can assign negative implicit weights to some dose groups and require a substantially stronger parallel trends assumption that rules out selection on potential outcome gains across treatment doses. The directional consistency with the binary results is reassuring, but the binary estimates reported in the main tables remain the preferred specifications.

7 Alternative Mechanisms

The pattern of results across outcomes and exposure channels allows us to distinguish the maternal opportunity-cost mechanism from three competing explanations. Each alternative generates predictions that differ from our model on at least one outcome-channel combination, and Table 1 in Section 2 summarizes the relevant comparative statics.

Consider first the income-effects explanation, in which sex selection is a normal good that rises with household resources (Almond et al., 2019). Under this explanation, any income source should increase sex selection, including paternal wages and provincial income shocks. The data reject this prediction. Across all specifications, paternal exposure coefficients for sex selection are small and statistically insignificant (Tables 3 and C7), while provincial exposure likewise shows no significant effect on the sex ratio. The strongest evidence against income-driven sex selection is the source-of-income prediction. Under a pure income mechanism, paternal and provincial exposure should also raise sex selection, since they raise household resources. Neither does. The joint pattern of positive sex selection, increased labor supply, and reduced fertility in response to maternal exposure, combined with the absence of sex-selection effects from paternal or provincial exposure, is difficult to reconcile with a model in which income alone drives sex selection. The fertility decline under maternal exposure is consistent with this picture but does not by itself distinguish the income and opportunity-cost channels, since maternal wage shocks combine an income effect (which would raise fertility) with a substitution effect on time (which lowers it). The fact that the substitution effect dominates confirms the centrality of the time-cost channel.

Next consider the relative-returns explanation of Qian (2008), building on earlier evidence in Rosenzweig and Schultz (1982), in which higher returns to female labor increase the economic value of daughters, reducing sex selection. Under this channel, expanding opportunities in female-

intensive industries should weaken son preference by making daughters more valuable to the household. Thus, maternal exposure to the BTA should decrease sex selection. Our estimates show the opposite. The maternal exposure coefficient ($M_i \times \text{Post}_t$) is positive and significant at the 5 percent level across all specifications in Tables 3 and C7, indicating that mothers in BTA-exposed industries are more likely to have sons, not less. This pattern contradicts the central prediction of the relative-returns mechanism.

The preceding test interprets the relative-returns channel at the level of individual wages. If a mother's own industry expands, her returns rise and daughters should become more valuable to the household. But the mechanism's original argument operates at a different level. It is not about whether a particular mother earns more today, but whether daughters born now will face better economic prospects as adults in the local labor market. When female-intensive industries expand in a province, all families in that province have reason to expect that daughters will find more rewarding work when they grow up, regardless of the mother's own industry. In the framework of Section 2, this channel enters through the cost parameter p_q. The forgone economic value of daughters increases as local female labor demand rises, raising the cost of sex selection and reducing son preference (Table 1, row 4).

To test this local-market reading directly, we construct a female-weighted provincial BTA exposure measure $\tilde{P}_p^w = \sum_j f_j \cdot s_{jp} \cdot \tau_j$, where f_j is the pre-BTA female employment share in industry j, s_{jp} is the pre-BTA employment share of industry j in province p, and τ_j is the industry-level tariff cut. Because $f_j \cdot s_{jp}$ need not sum to one across industries, $\tilde{P}_p^w$ captures both the intensity of tariff cuts in female-intensive industries and the baseline importance of female employment in each province. This is the relevant object for testing the relative-returns mechanism, since provinces where female-intensive industries are both larger and more exposed to tariff cuts are precisely those where daughters' future labor market prospects improve most. We define $P_p^w = \mathbf{1}[\tilde{P}_p^w > \text{median}]$ as a binary indicator for above-median female-weighted provincial exposure. Provinces where $P_p^w = 1$ are precisely those where the BTA most expanded demand in female-intensive industries, generating the variation in daughters' local economic value that the relative-returns mechanism emphasizes. Under this account, $P_p^w \times \text{Post}_t$ should reduce sex selection. We estimate:

$$y_{ipt} = \beta_1(M_i \times \text{Post}_t) + \beta_2(F_i \times \text{Post}_t) + \beta_3(P_p \times \text{Post}_t) + \beta_4(P_p^w \times \text{Post}_t) + \text{FE} + X'_{ipt}\phi + \varepsilon_{ipt}, \quad (14)$$

where the relative-returns prediction is $\beta_4 < 0$ for sex selection. Table C12 reports the results across all four outcomes, entering all exposure channels simultaneously. The female-weighted provincial exposure coefficient is small and statistically insignificant for sex selection and fertility. The one outcome where it is large and significant is maternal wage employment, where $P_p^w \times \text{Post}_t$ raises the probability that mothers hold wage jobs by 6 percentage points (column 3, $p < 0.05$). This reflects higher returns to female labor in provinces most exposed to female-intensive indus-

try expansion, which is precisely the local labor-market condition that raises the opportunity-cost component of p_q through which the relative-returns mechanism is supposed to operate. The mechanism's first stage is thus present in the data: female labor market returns rose where the mechanism predicts they should. Yet higher returns do not reduce son preference. Neither the local-labor-market reading nor the individual-wage reading of the relative-returns mechanism accounts for the sex-selection patterns we observe.

The bargaining explanation presents a subtler challenge. Appendix B.1 extends the baseline model by allowing bargaining weights to depend on relative wages, $\lambda = \lambda(w_m/w_d)$ with $\lambda'(\cdot) > 0$, and by permitting heterogeneous preferences across spouses, α_{qg} and α_{ng} for $g \in \{m, d\}$. Under this extension, a maternal wage increase shifts bargaining power toward the mother, and the direction of the effect on sex selection depends on whose son preference is stronger. The model predicts $\partial q/\partial w_m > 0$ if mothers have a stronger relative preference for sons than fathers, $\alpha_{qm}/\alpha_{nm} > \alpha_{qd}/\alpha_{nd}$. Higher maternal wages reduce fertility under an additional condition, that mothers prefer fewer children than fathers ($\alpha_{nm} < \alpha_{nd}$), which also implies that paternal wages raise fertility. Both predictions align with our estimates.

However, the bargaining model generates a testable prediction that the data reject. Under the same parameter restrictions that match the maternal wage effects, the model predicts that higher paternal wages should reduce sex selection ($\partial q/\partial w_d < 0$) by shifting bargaining power toward the father, who has weaker relative son preference. Paternal BTA exposure does produce meaningful economic effects. Fathers in high-exposure industries are 11 percentage points more likely to hold wage jobs and work 9 percent more hours (Table C2), and paternal exposure raises the annual birth probability by 0.5 percentage points (Tables 5 and C12). These patterns confirm that the paternal income channel is operative and that paternal wages rose in ways consistent with the bargaining mechanism activating. Yet the paternal exposure coefficient for sex selection is small and statistically insignificant across the binary (Table 3), continuous (Table C7), and panel (Table 6) specifications. Higher paternal wages raised income and fertility, as the baseline model predicts, but did not shift son preference in the direction the bargaining extension requires.

A simpler test reaches the same conclusion without imposing bargaining structure. If the positive maternal effect on sex selection reflected mothers exercising son preference as their bargaining power grew, then mothers who already hold a stronger bargaining position should select more sons even before the trade shock. We test this prediction using the 1999 Census, before the BTA took effect. Among births in 1997 and 1998, mothers with strictly more education than their husbands, roughly one quarter of the sample, are no more likely to have a son (Table C13). The absence of any pre-treatment association between relative parental education and sex selection undermines the premise that maternal bargaining power drives son preference.

The specific prediction that fails is the link from paternal wage gains to reduced son preference. The null sex-selection effect is not simply underpowered, since we observe significant paternal effects on labor supply and fertility. The baseline model accounts for this pattern without requiring

parameter restrictions on spousal preferences.

8 Conclusion

The 2001 Bilateral Trade Agreement between the United States and Vietnam raised sex selection among mothers in trade-exposed industries. Using industry-specific tariff cuts as a natural experiment, and contemporaneous household survey data that record parental employment at the time of birth, we find that women in high-exposure industries are approximately 6 percentage points more likely to have male children, work approximately 14 percent more hours per month, and have a 0.6 percentage point lower annual birth probability. These results are maternal effects. Paternal and provincial exposure channels are silent on sex selection. Panel analysis with individual fixed effects produces a larger estimate than the cross-sectional estimate, reducing concern that simple sorting stories drive the result. Treatment effects also attenuate substantially once a family has a son, a pattern consistent with maternal opportunity costs interacting with prior child sex composition.

Tests of competing mechanisms reveal a consistent pattern. Each alternative mechanism has an operative first-stage effect on wage employment and labor supply, but none transmits to sex selection. Paternal BTA exposure raises earnings and fertility, as the income channel predicts, yet leaves sex selection unchanged. Female-weighted provincial exposure raises maternal wage employment but does not reduce son preference, and higher paternal wages fail to shift sex selection in the direction the bargaining model requires. The maternal time-cost mechanism remains the only channel consistent with the full set of results.

The mechanism we identify is not unique to Vietnam. Albania, Montenegro, and Georgia show similar trajectories around their own trade integration episodes (Figure 1), and the same structural conditions are present in many other developing economies with son preference, patrilocal family structures, and rapid expansion of female employment, including parts of Indonesia and sub-Saharan Africa. The key moderating conditions are not trade itself but the conjunction of maternal time costs and son preference. Where mothers bear the primary childcare burden and families desire sons, trade-induced female employment opportunities can simultaneously empower women economically and intensify demographic imbalances. One limitation of our analysis is that we observe children aged 0–1 who are alive and co-resident at the time of the survey, so we cannot fully distinguish prenatal sex selection from postnatal mechanisms such as sex-differential infant mortality. Vietnam's low infant mortality rate, approximately 18 per 1,000 in 2002 with a small male-female gap, limits the scope for postnatal channels to generate effects of the magnitude we document. Administrative birth records linking ultrasound visits to maternal employment would provide a definitive test.

Addressing the link between female employment and sex selection requires policies that reduce the maternal time constraint directly. Income transfers or wage subsidies alone will not

suffice, since the underlying mechanism runs through time costs rather than income. What matters is expanding access to institutional childcare and other supports that decouple female labor market participation from the opportunity cost of raising children.

References

ActionAid-Vietnam (2016). Unpaid Care Work: Make a House becomes a Home. Technical report, ActionAid Vietnam, Hanoi.

Almond, D., Li, H., and 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. and Kumler, T. J. (2019). Women's worth: Trade, female income, and fertility in India. *Economic Development and Cultural Change*, 67(3):687–724.

Belanger, D. (2002). Son preference in a rural village in north Vietnam. *Studies in Family Planning*, 33(4):321–334.

Borker, G., Eeckhout, J., Luke, N., Minz, S., Munshi, K., and Swaminathan, S. (2024). Wealth, Marriage, and Sex Selection. Working paper.

Callaway, B., Goodman-Bacon, A., and 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., and Alkema, L. (2019). Systematic assessment of the sex ratio at birth for all countries and estimation of national imbalances and regional reference levels. *Proceedings of the National Academy of Sciences of the United States of America*, 116(19):9303–9311.

Chen, Y., Li, H., and Meng, L. (2013). Prenatal sex selection and missing girls in China: Evidence from the diffusion of diagnostic ultrasound. *Journal of Human Resources*, 48(1):36–70.

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

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

Drozdoff, N., McCaig, B., and Pavcnik, N. (2025). Foreign manufacturing jobs and women: Impacts on work, marriage, and childbirth. Working Paper.

Duflo, E. (2003). Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa. *World Bank Economic Review*, 17(1):1–25.

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

Gaddis, I. and Pieters, J. (2017). The Gendered Labor Market Impacts of Trade Liberalization: Evidence from Brazil. *The Journal of Human Resources*, 52(2):457–490. Publisher: [University of Wisconsin Press, Board of Regents of the University of Wisconsin System].

General Statistics Office of Vietnam (2002–2010). Vietnam household living standards survey (vhlss), 2002–2010. Household survey microdata. Waves used in this paper: 2002, 2004, 2006, 2008, and 2010.

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

Huynh, Q. and Ku, H. (2025). Gender Norms and Female Labor Supply: Evidence from Export Shocks in Vietnam. IZA Discussion Paper.

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

Jayachandran, S. and Voena, A. (2026). Women's power in the household. Forthcoming, Journal of Economic Literature. Draft dated January 14, 2026.

Jones, L. E., Schoonbroodt, A., and Tertilt, M. (2011). Fertility Theories: Can They Explain the Negative Fertility-Income Relationship? In Shoven, J. B., editor, *Demography and the Economy*, National Bureau of Economic Research Conference Report, pages 43–100. University of Chicago Press, Chicago, IL.

Kovak, B. K. (2013). Regional effects of trade reform: What is the correct measure of liberalization? *American Economic Review*, 103(5):1960–1976.

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., Nguyen, M., and Kaestner, R. (2022). Export expansion and investment in children's human capital: Evidence from the u.s.-vietnam bilateral trade agreement. NBER Working Paper 29830, National Bureau of Economic Research.

McCaig, B. and Pavcnik, N. (2018). Export markets and labor allocation in a low-income country. *American Economic Review*, 108(7):1899–1941.

McCaig, B., Pavcnik, N., and Wong, W. F. (2025). Foreign and Domestic Firms: Long-Run Employment Effects of Trade Liberalization. *Review of Economics and Statistics*. Forthcoming.

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. and Schultz, T. P. (1982). Market opportunities, genetic endowments, and intrafamily resource distribution: child survival in rural India. *American Economic Review*, 72(4):803–815.

Ruggles, S., Cleveland, L., Lovaton, R., Sarkar, S., Sobek, M., Burk, D., Ehrlich, D., Heimann, Q., Lee, J., and Merrill, N. (2025). Integrated public use microdata series, international: Version 7.6. Dataset.

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

STAR-Vietnam (2003). An Assessment of the Economic Impact of the United States - Vietnam Bilateral Trade Agreement. Technical report, The National Political Publishing House, Hanoi.

A Data Appendix

A.1 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 and 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 for a given individual, the second rule no longer applies to that 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. The algorithm finds both parents correctly for 88.95% of the children, the correct mother (regardless of father match) for 93.23% (88.95 + 4.28), and the correct father (regardless of mother match) for 90.99%. 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.

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.

B Model Appendix

This appendix develops an endogenous bargaining extension of the baseline model in Section 2 that generates testable predictions distinct from the time-cost mechanism.

B.1 Endogenous Bargaining Extension

Rising female wages may shift bargaining power within the household, affecting fertility decisions through a channel distinct from the time-cost mechanism. If mothers and fathers have different preferences over fertility and child gender, then changes in bargaining weights could generate maternal wage effects on sex selection even without asymmetric childcare costs.

Each parent $g \in \{m, d\}$ has potentially different preferences over fertility and sons:

$$U_g = \alpha_c \log(c_g) + \alpha_\ell \log(\ell_g) + \alpha_{ng} \log(n) + \alpha_{qg} \log(Q). \tag{15}$$

The household maximizes a weighted sum of utilities where bargaining weights depend on relative wages. The household's problem is:[21]

$$\max_{\{c_m, c_d, n, q\}} \lambda(w_m/w_d)U_m + (1 - \lambda(w_m/w_d))U_d \tag{16}$$

subject to the budget constraint:

$$c_m + c_d + p_q Q \leq I + w_d(1 - \ell_d) + w_m\left(1 - \ell_m - \gamma n\right), \tag{17}$$

where $\lambda'(\cdot) > 0$ so higher relative maternal wages increase mother's bargaining power. Time constraints remain as in the baseline model.

This model yields three key equations. The sex ratio at birth is:

$$\log q = \log\left(\frac{\lambda\alpha_{qm} + (1-\lambda)\alpha_{qd}}{\lambda\alpha_{nm} + (1-\lambda)\alpha_{nd}}\right) + \log\gamma - \log p_q + \log w_m. \tag{18}$$

Mother's leisure choice is:

$$\log \ell_m = \log\alpha_\ell + \log\lambda - \log w_m + \log W. \tag{19}$$

Total fertility is:

$$\log n = \log\left(\frac{\lambda\alpha_{nm} + (1-\lambda)\alpha_{nd}}{\gamma}\right) + \log W - \log w_m. \tag{20}$$

The comparative statics reveal that this extended model can match some but not all of our empirical findings. Mother's wages increase sex selection ($\partial q/\partial w_m > 0$) if mothers have stronger

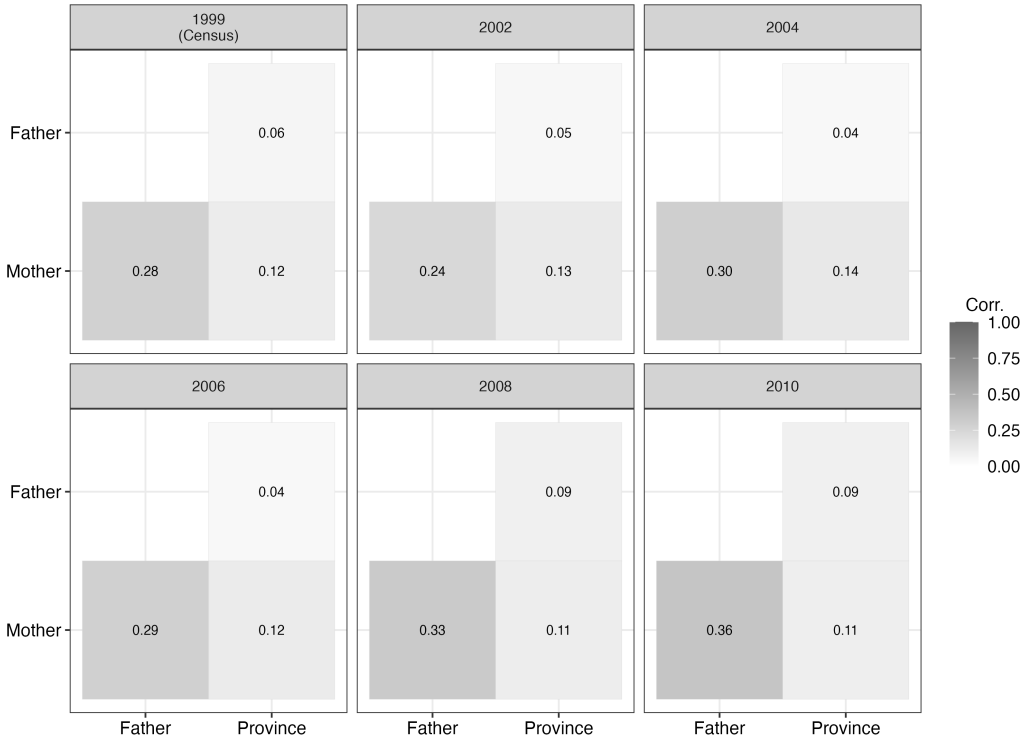
[21] The Lagrangian formulation and detailed first-order conditions are available upon request.

relative preference for sons than fathers ($\alpha_{qm}/\alpha_{nm} > \alpha_{qd}/\alpha_{nd}$). Under this condition, maternal wages also reduce fertility ($\partial n/\partial w_m < 0$). The fertility equation simplifies to $\log n = \log[(\lambda\alpha_{nm} + (1-\lambda)\alpha_{nd})/\gamma] + \log W - \log w_m$, where the weighted fertility preference replaces the common preference parameter of the unitary model. The direct opportunity-cost effect of w_m (the $-\log w_m$ term) dominates. If mothers prefer fewer children than fathers ($\alpha_{nm} < \alpha_{nd}$), then father's wages increase fertility ($\partial n/\partial w_d > 0$) by shifting bargaining weight toward the higher-fertility parent. These predictions align with our empirical results.

However, under these same parameter restrictions, the model predicts father's wages would reduce sex selection ($\partial q/\partial w_d < 0$), contrary to our finding of no effect (Table 3). The intuition is that higher paternal wages shift bargaining power toward fathers, who have weaker relative preference for sons by assumption. This counterfactual prediction indicates that bargaining cannot be the primary driver of the asymmetric pattern we observe, in which maternal wages raise sex selection but paternal wages do not. The baseline model's time-cost mechanism better matches the full pattern of empirical results, particularly the null effect of paternal wages on sex selection.

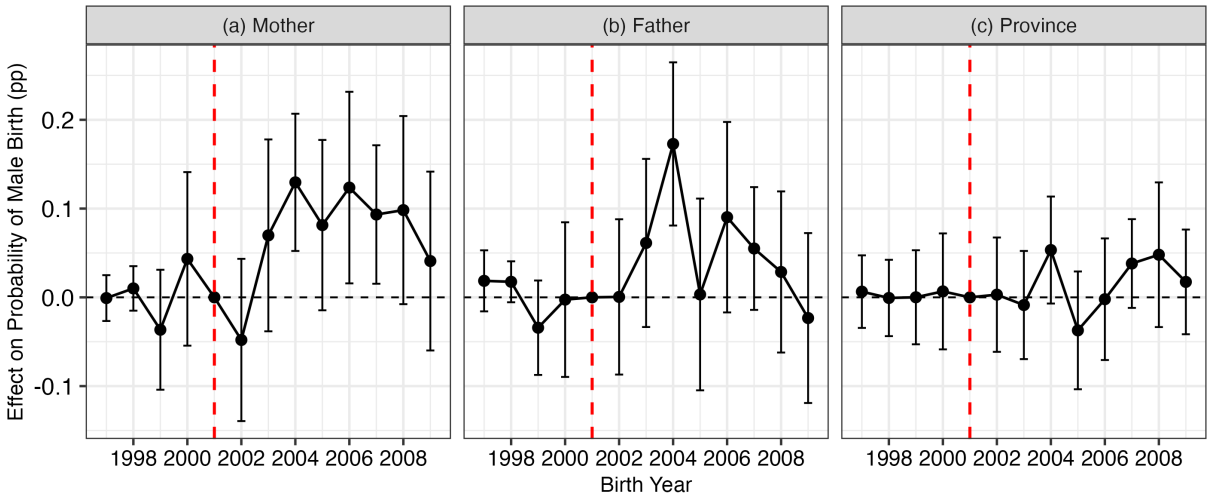
C Additional Tables and Figures

Figure C1: Correlations between Treatment Variables by Survey Wave



Notes: Each panel shows the pairwise correlation matrix between tariff exposure measures for one wave, using the 1999 Population Census (pre-treatment baseline) and the Vietnam Household Living Standards Survey (2002–2010). Mother (τ_m) denotes maternal industry tariff exposure, Father (τ_d) the paternal, and Province (τ_p) the shift-share provincial exposure.

Figure C2: Individual Event Study Estimates of BTA Exposure Effects on Male Births



Notes: This figure displays event study estimates examining the effects of BTA exposure on male births from separate regressions for each exposure channel. The dependent variable is an indicator for male birth, and the independent variables are year dummies interacted with the respective binary exposure indicator, controlling for province, year, and parental industry fixed effects. The panels show coefficients for mothers' industry tariff exposure (left), fathers' industry tariff exposure (center), and province-level tariff exposure (right). The binary indicators compare high tariff exposure (above median) to low exposure (below median). 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.

Table C1: Comparison of Panel Sub-sample and Main Cross-section (2002 Baseline)

	Cross-section (2002 VHLSS)	Panel sub-sample (2002–2004–2006)
Mother's age	28.04	29.05
	(5.07)	(5.08)
Mother's education (yrs)	6.61	6.49
	(3.62)	(3.60)
Ethnic minority	0.230	0.233
	(0.421)	(0.423)
Urban	0.173	0.163
	(0.378)	(0.369)
Province exposure	0.09	0.09
	(0.02)	(0.02)
N	5,934	879

Notes: Standard deviations in parentheses. Cross-section uses the 2002 VHLSS wave of the main kid sample (mothers of infants aged 0–1). Panel sub-sample uses the 2002 baseline wave of the rotating panel restricted to women aged 20–40 with a stable residential location and at least one birth before the BTA shock and at least one birth after, ensuring within-mother variation across the pre- and post-treatment periods.

Table C2: Effects of BTA Exposures on Paternal Labor Market Outcomes

	Wage Employment			Log Monthly Hours		
	(1)	(2)	(3)	(4)	(5)	(6)
$M_i \times$ Post	0.04	0.00	0.00	0.07**	0.02	0.01
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
$F_i \times$ Post		0.11***	0.11***		0.09***	0.09***
		(0.03)	(0.03)		(0.02)	(0.02)
$P_p \times$ Post			−0.01			0.05*
			(0.04)			(0.03)
R^2	0.53	0.53	0.53	0.23	0.24	0.24
Observations	17,389	17,389	17,389	17,223	17,223	17,223
Dep. Var. Mean	0.41	0.41	0.41	4.9	4.9	4.9
Controls		✓	✓		✓	✓

Notes: The dependent variable in columns (1)–(3) is an indicator equal to 1 if the father has wage employment. In columns (4)–(6) it is the log of monthly hours worked in his main job. Controls in columns (2), (3), (5), and (6) match those in Table 3. All specifications include survey-year, province, and parents' 2-digit ISIC fixed effects. Standard errors, clustered at the province level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C3: Robustness to ISIC Industry Clustering

	MaleBirth	log(MomWorkHrs)	MomWageEmp	AnyBirth
	(1)	(2)	(3)	(4)
$M_i \times$ Post	0.06**	0.10***	0.16***	-0.006***
	(0.03)	(0.03)	(0.03)	(0.002)
$F_i \times$ Post	0.02	0.01	0.004	0.005**
	(0.03)	(0.04)	(0.02)	(0.002)
$P_p \times$ Post	0.02	0.06**	-0.001	-0.002
	(0.02)	(0.03)	(0.01)	(0.0009)
R^2	0.02	0.25	0.61	0.29
Observations	63,894	17,198	17,389	2,011,963
Dep. Var. Mean	0.52	4.7	0.25	0.02
Mom FE				✓
Survey Year FE		✓	✓	
Birth Year FE	✓			✓
Dad ISIC FE	✓	✓	✓	
Mom ISIC FE	✓	✓	✓	
Province FE	✓	✓	✓	
Controls	✓	✓	✓	✓

Notes: Controlled specifications with binary treatment indicators (M_i, F_i, P_p). Standard errors are clustered by 2-digit ISIC industry only, replacing the baseline province clustering. Controls include household size, parental ages at birth, parental education, urban status, and minority status. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C4: Robustness to Province and ISIC Two-Way Clustering

	MaleBirth	log(MomWorkHrs)	MomWageEmp	AnyBirth
	(1)	(2)	(3)	(4)
$M_i \times$ Post	0.06***	0.10***	0.16***	-0.006***
	(0.02)	(0.04)	(0.04)	(0.002)
$F_i \times$ Post	0.02	0.01	0.004	0.005**
	(0.03)	(0.03)	(0.01)	(0.002)
$P_p \times$ Post	0.02	0.06	-0.001	-0.002
	(0.01)	(0.04)	(0.02)	(0.001)
R^2	0.02	0.25	0.61	0.29
Observations	63,894	17,198	17,389	2,011,963
Dep. Var. Mean	0.52	4.7	0.25	0.02
Mom FE				✓
Survey Year FE		✓	✓	
Birth Year FE	✓			✓
Dad ISIC FE	✓	✓	✓	
Mom ISIC FE	✓	✓	✓	
Province FE	✓	✓	✓	
Controls	✓	✓	✓	✓

Notes: Controlled specifications with binary treatment indicators (M_i, F_i, P_p). Standard errors are computed with two-way clustering by province and 2-digit ISIC industry simultaneously, accounting for both geographic and industry-level error correlation. Controls include household size, parental ages at birth, parental education, urban status, and minority status. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C5: Robustness to Region-by-Year Fixed Effects

	MaleBirth	log(MomWorkHrs)	MomWageEmp	AnyBirth
	(1)	(2)	(3)	(4)
$M_i \times$ Post	0.06**	0.10**	0.15***	-0.006***
	(0.02)	(0.04)	(0.03)	(0.001)
$F_i \times$ Post	0.03	0.04	-0.01	0.005***
	(0.03)	(0.03)	(0.01)	(0.001)
$P_p \times$ Post	0.02	0.05	-0.002	-0.002
	(0.02)	(0.06)	(0.02)	(0.002)
R^2	0.02	0.26	0.61	0.29
Observations	63,894	17,198	17,389	2,011,963
Dep. Var. Mean	0.52	4.7	0.25	0.02
Mom FE				✓
Region-Year FE	✓	✓	✓	✓
Dad ISIC FE	✓	✓	✓	
Mom ISIC FE	✓	✓	✓	
Province FE	✓	✓	✓	
Controls	✓	✓	✓	✓

Notes: Each specification uses binary treatment indicators (M_i, F_i, P_p) and adds region-by-year fixed effects to absorb region-specific, time-varying confounders. MaleBirth, log monthly hours, and wage-earner specifications include province, mother's ISIC, and father's ISIC fixed effects; MaleBirth adds birth-year fixed effects while labor supply specifications add survey-year fixed effects. The AnyBirth specification uses individual mother fixed effects in place of province and industry fixed effects. All specifications include controls for household size, parental ages at birth, parental education, urban status, and minority status. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C6: Robustness to Logit Estimation for Male Birth Outcome

	MaleBirth	
	(1)	(2)
$M_i \times$ Post	0.22**	0.21**
	(0.10)	(0.10)
$F_i \times$ Post	0.11	0.11
	(0.12)	(0.12)
$P_p \times$ Post	0.006	-0.006
	(0.05)	(0.05)
Observations	66,554	64,162
Dep. Var. Mean	0.52	0.52
Birth Year FE	✓	✓
Dad ISIC FE	✓	✓
Mom ISIC FE	✓	✓
Province FE	✓	✓
Controls		✓

Notes: Logit estimates of the effect of parental and province-level exposures on the probability of a male birth using binary treatment indicators (M_i, F_i, P_p) for workers in above-median tariff reduction industries. Both specifications include birth-year, province, mother's ISIC, and father's ISIC fixed effects. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C7: Robustness to Continuous Tariff Exposure Measures

	MaleBirth					
	(1)	(2)	(3)	(4)	(5)	(6)
$\tau_m \times$ Post	0.216**		0.205**	0.206**	0.212**	0.204**
	(0.093)		(0.090)	(0.086)	(0.086)	(0.088)
$\tau_d \times$ Post		0.100	0.049	0.052	0.046	0.050
		(0.127)	(0.126)	(0.129)	(0.127)	(0.129)
$\tau_p \times$ Post					-0.003	0.005
					(0.008)	(0.012)
R^2	0.015	0.014	0.015	0.015	0.015	0.015
Observations	66,562	66,562	66,562	64,170	66,562	64,170
Dep. Var. Mean	0.518	0.518	0.518	0.518	0.518	0.518
Controls				✓		✓

Notes: The dependent variable is an indicator equal to 1 if the child is male. The continuous treatment variables (τ_m, τ_d, τ_p) measure the actual tariff reduction exposure for mother's industry, father's industry, and province-level exposure respectively. Columns show progressive addition of treatment variables and controls, following the same structure as Table 3. All regressions include the same fixed effects as the main specification. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C8: Robustness to Continuous Tariff Exposure on Maternal Labor Market Outcomes

	MomWageEmp			log(MomWorkHrs)		
	(1)	(2)	(3)	(4)	(5)	(6)
$\tau_m \times$ Post	0.509***	0.507***	0.508***	0.146	0.044	0.068
	(0.146)	(0.155)	(0.155)	(0.150)	(0.159)	(0.157)
$\tau_d \times$ Post		-0.101*	-0.101*		0.161	0.165
		(0.055)	(0.055)		(0.142)	(0.142)
$\tau_p \times$ Post			-0.002			-0.082**
			(0.023)			(0.039)
R^2	0.600	0.603	0.603	0.246	0.253	0.253
Observations	17,665	17,665	17,665	17,472	17,472	17,472
Dep. Var. Mean	0.248	0.248	0.248	4.66	4.66	4.66
Controls		✓	✓		✓	✓

Notes: The dependent variables are an indicator equal to 1 if the mother has wage employment (columns (1)–(3)) and the log of monthly hours worked (columns (4)–(6)). The continuous treatment variables (τ_m, τ_d, τ_p) measure actual tariff reduction exposure for mother's industry, father's industry, and province respectively. Columns show progressive addition of treatments and controls. All regressions include survey-year, province, and parents' 2-digit ISIC fixed effects. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C9: Robustness to Continuous Tariff Exposure on Fertility

	AnyBirth					
	(1)	(2)	(3)	(4)	(5)	(6)
$\tau_m \times$ Post	-0.003		-0.008**	-0.007**	-0.009***	-0.008**
	(0.003)		(0.004)	(0.003)	(0.003)	(0.003)
$\tau_d \times$ Post		0.021***	0.023***	0.023***	0.023***	0.023***
		(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$\tau_p \times$ Post					0.0009***	0.004***
					(0.0003)	(0.001)
R^2	0.286	0.286	0.286	0.287	0.286	0.287
Observations	2,011,978	2,011,992	2,011,963	2,011,963	2,011,963	2,011,963
Dep. Var. Mean	0.018	0.018	0.018	0.018	0.018	0.018
Controls				✓		✓

Notes: The dependent variable is an indicator equal to 1 if the mother had any birth in the past 12 months. The continuous treatment variables (τ_m, τ_d, τ_p) measure actual tariff reduction exposure for mother's industry, father's industry, and province respectively. Columns show progressive addition of treatments following the same structure as Table C7. All specifications include mother and birth-year fixed effects. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C10: Fertility Results with Mother-Level Clustering (Binary Treatment)

	AnyBirth					
	(1)	(2)	(3)	(4)	(5)	(6)
$M_i \times$ Post	-0.004***		-0.006***	-0.006***	-0.006***	-0.006***
	(0.0008)		(0.0009)	(0.0009)	(0.0009)	(0.0009)
$F_i \times$ Post		0.003***	0.005***	0.005***	0.005***	0.005***
		(0.0010)	(0.001)	(0.001)	(0.001)	(0.001)
$P_p \times$ Post					-0.0001	-0.002**
					(0.0006)	(0.0007)
R^2	0.29	0.29	0.29	0.29	0.29	0.29
Observations	2,011,963	2,011,963	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
Controls				✓		✓

Notes: Same specification as Table 5. Standard errors are clustered at the mother level instead of the province level to account for within-mother serial correlation in the retrospective panel. All specifications include mother and birth-year fixed effects. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C11: Fertility Results with Mother-Level Clustering (Continuous Treatment)

	AnyBirth					
	(1)	(2)	(3)	(4)	(5)	(6)
$\tau_m \times$ Post	-0.003		-0.008**	-0.007**	-0.009***	-0.008**
	(0.003)		(0.003)	(0.003)	(0.003)	(0.003)
$\tau_d \times$ Post		0.021***	0.023***	0.023***	0.023***	0.023***
		(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
$\tau_p \times$ Post					0.0009***	0.004***
					(0.0003)	(0.0009)
R^2	0.286	0.286	0.286	0.287	0.286	0.287
Observations	2,011,978	2,011,992	2,011,963	2,011,963	2,011,963	2,011,963
Dep. Var. Mean	0.018	0.018	0.018	0.018	0.018	0.018
Controls				✓		✓

Notes: Same specification as Table C9. Standard errors are clustered at the mother level instead of the province level to account for within-mother serial correlation in the retrospective panel. All specifications include mother and birth-year fixed effects. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C12: Relative-Returns Test for All Outcomes with All Four Exposure Channels

	MaleBirth	log(MomWorkHrs)	MomWageEmp	AnyBirth
	(1)	(2)	(3)	(4)
$M_i \times$ Post	0.06**	0.10**	0.15***	-0.006***
	(0.02)	(0.04)	(0.03)	(0.001)
$F_i \times$ Post	0.02	0.02	-0.0007	0.005***
	(0.03)	(0.03)	(0.01)	(0.001)
$P_p \times$ Post	0.009	0.12*	-0.04	-0.002
	(0.01)	(0.07)	(0.03)	(0.002)
$P_p^w \times$ Post	0.01	-0.09	0.06**	0.0008
	(0.01)	(0.07)	(0.03)	(0.002)
R^2	0.02	0.25	0.61	0.29
Observations	63,894	17,198	17,389	2,011,963
Dep. Var. Mean	0.52	4.7	0.25	0.02
Mom FE				✓
Dad ISIC FE	✓	✓	✓	
Mom ISIC FE	✓	✓	✓	
Province FE	✓	✓	✓	

Notes: Each column reports the controlled specification for one dependent variable. M_i and F_i are binary indicators for above-median maternal and paternal industry BTA exposure. P_p is a binary indicator for above-median province-level shift-share BTA exposure. P_p^w is a binary indicator for above-median female-weighted provincial BTA exposure ($\sum_j f_j \cdot s_{jp} \cdot \tau_j$), where f_j is the pre-BTA female employment share in industry j. Post is defined as birth year after 2001 for MaleBirth and AnyBirth, and survey year after 2002 for labor outcomes. Columns (1)–(3) include province, parents' 2-digit ISIC, and time fixed effects. Column (4) includes mother and birth-year fixed effects. Controls match those in Table 3. Standard errors, clustered at the province level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C13: Pre-Treatment Bargaining Power and Sex Selection, 1999 Census

	Mom $\leq$ Dad education	Mom $>$ Dad education
Proportion male	0.5125	0.5145
	(0.0029)	(0.0049)
N	30,548	10,619
Regression. Dep. var. = Male birth		
Mom $>$ Dad education	0.0026	
	(0.0048)	
N	41,167	
Province FE	Yes	

Notes: Sample consists of children born in 1997–1998 observed in the 1999 Vietnam Population Census (IPUMS). "Mom $>$ Dad education" equals one if the mother's education level is strictly higher than the father's. The top panel reports group means with standard errors of the mean in parentheses. The bottom panel reports the OLS coefficient from regressing a male-birth indicator on the education-comparison variable with province fixed effects. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.