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

Nghiêm Q. Huỳnh† Ngoc T. T. Nguyen‡

January 2, 2025

Abstract

How does economic development influence sex selection when parents face pressures from work, childcare, and son preference? We investigate this question in Vietnam using the 2001 trade liberalization. Our model integrates son preference into a quantity-quality framework with maternal childcare burdens to generate distinct predictions from competing theories. By exploiting tariff cuts across industries, we find that women in exposed industries have more male children, fewer births, and work more. The impacts stem from maternal exposure rather than fathers' industries or local markets, indicating that income, bargaining, or relative returns to daughters have little effect.

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

†University of Oklahoma, nqh@ou.edu

‡Lightcast, jade.ng.nguyen@gmail.com

1 Introduction

Since Becker and Lewis (1973), economists have studied the quantity-quality tradeoff parents face between family size and investment per child. In contexts where parents strongly prefer sons, however, this tradeoff takes on an additional dimension—they must balance their desire for a son against the rising costs of additional children. When mothers bear primary responsibility for childcare while participating in the labor market, improved economic opportunities can create particularly acute pressure on maternal time, affecting both fertility and sex selection decisions.

This focus on maternal time offers a new perspective on how economic development shapes sex selection, which existing research has approached through several other channels. Jayachandran (2017) provides indirect evidence for this mechanism by showing how fertility decline intensifies pressure for sex selection as parents face a stricter constraint on achieving a son. Other studies emphasize returns to daughters (Qian, 2008), while recent work highlights the role of household income (Almond et al., 2019). However, evidence remains mixed, partly because these channels are difficult to disentangle empirically. Moreover, these explanations largely overlook the asymmetric burden of childcare between mothers and fathers, a global regularity (Doepke et al., 2023).

Vietnam's experience provides an ideal setting to test these competing mechanisms. The country's sex ratio at birth rose from 107 boys per 100 girls in 1995 to 112 in 2015 amid rapid economic integration, mirroring a global phenomenon observed in many developing countries (Figure 1). The 2001 Bilateral Trade Agreement (BTA) with the United States generated large, industry-specific changes in labor market opportunities, allowing us to separately identify effects operating through maternal employment, paternal employment, household income, and returns to daughters. These distinct sources of variation motivate our theoretical framework that nests competing explanations while highlighting the role of maternal time costs.

We formalize this intuition through a standard quality-quantity (Q-Q) model that emphasizes families' desire for at least one son and mothers' work-childcare tradeoff (Jones et al., 2011). Our model shows how rising maternal market outcomes can increase sex selection even while improving women's economic status, contrasting with standard relative returns explanations. Crucially, our model predicts that *only* maternal wages affect sex selection, whereas household income and paternal wages influence fertility but not sex selection. These predictions diverge from income-based explanations like Almond et al. (2019), in which all income sources raise both fertility and sex selection. They also differ from bargaining models, which posit that paternal wages can shift bargaining power, thereby affecting sex selection.

Using detailed household surveys from 2002 to 2016, we exploit variation in BTA exposure across industries and provinces to identify the causal impact of trade liberalization on sex selection. We measure exposure through four channels: mother's industry-specific tariff cuts, father's industry-specific tariff cuts, local province-level economic effects, and local female labor demand. This approach allows us to disentangle how maternal opportunity costs, household income, and returns

to daughters separately influence sex selection.

The results generally support our model's emphasis on maternal time costs as the key mechanism linking trade liberalization to demographic outcomes. Women whose industries faced a median 7 percentage point tariff cut were 2 percentage points more likely to have male children (increasing sex ratio from 107 to 115 boys per 100 girls), worked 3 more hours per month, and had 0.07 percentage point lower birth probability. Consistent with our model, similar exposure through fathers' industries or local labor markets shows no impact on sex selection, though both increase fertility as predicted by income effects.

These baseline findings remain robust to panel analysis that accounts for unobserved individual preferences, with effects on sex selection roughly twice as large among women with births both before and after the policy change. Adding region-by-year fixed effects to account for infrastructure development, using different clustering approaches, and employing logit estimation yield similar results. Examining heterogeneity reveals effects are strongest in rural and high patrilocal areas, higher birth orders, and among less educated mothers–precisely the contexts where son preference and childcare burdens are most binding. Event study analysis shows no differential pre-trends across industries and sustained effects post-trade liberalization.

Our findings contribute to several literatures. First, we extend research on quantity-quality tradeoffs in fertility decisions. The classic literature examines how parents balance family size against investments in each child's human capital, such as education (Rosenzweig & Wolpin, 1980; Rosenzweig & Zhang, 2009). Recent work shows how son preference can skew these post-birth investments toward boys (Anukriti et al., 2022; Jayachandran & Kuziemko, 2011). We demonstrate that parents face an additional tradeoff even before birth: between family size and child gender composition. By incorporating mothers' asymmetric childcare burden into a Q-Q framework, we show how rising maternal opportunity costs intensify pressure to secure a son while having fewer children.

Second, we contribute to research on the determinants of sex selection and, more generally, gender inequality in developing countries (Jayachandran, 2015). Existing studies have provided evidence of several mechanisms, including fertility decline (Ebenstein, 2010), access to ultrasound technology (Chen et al., 2013; Lin et al., 2014), and the relative returns to daughters (Rosenzweig and Schultz, 1982). Our theoretical framework and empirical design allow us to test these competing explanations, showing how they interact through maternal time constraints. By emphasizing mothers' work-childcare tradeoff, we provide a new perspective on why economic development can exacerbate gender imbalances even as it increases women's economic opportunities.

Third, we extend the literature on how trade affects gender outcomes in developing countries. Previous studies examining trade's impact on sex selection in India (Anukriti & Kumler, 2019; Chakraborty, 2015) rely on decennial census data, requiring strong assumptions about the stability of parents' employment and location over long periods. Our analysis exploits detailed household surveys conducted every two years, allowing us to observe parents' industry and location at the mo-

ment fertility decisions are made, rather than years before or after. This data advantage, combined with our theoretical framework that delivers sharp predictions across outcomes, enables us to more precisely identify how trade liberalization affects parents' fertility choices.

2 Model

We adapt a standard Q-Q model from Jones et al. (2011) to examine the channels through which trade liberalization affects sex selection in Vietnam. Our model captures two key features of the Vietnamese context. First, households exhibit a strong preference for having at least one son. While they value the total number of children (n), their utility is significantly enhanced by the presence of a son. We model this preference by introducing the concept of an *effective number of sons*, defined as $Q = qn$, where $q \in [0, 1]$ represents the probability of a male child. This formulation implies that households prefer one son ($Q = 1$) to multiple daughters without a son ($Q = 0$).

Second, mothers bear the primary responsibility for childcare while also participating in the labor market. In Vietnam, the female employment rate is high, approximately 77% among prime-aged women (20-64 years old) in our household data. Furthermore, women spend over five hours per day on childcare compared to men's three hours, and this gap widens with larger family sizes (ActionAid-Vietnam, 2016). Each child requires γ units of the mother's time, reducing her available time for work.

A household consists of a mother (m) and a father (d). They jointly decide on their private consumption (c_m, c_d), leisure time (ℓ_m, ℓ_d), the number of children (n), and the probability of a child being male (q). The utility function for each parent $g \in \{m, d\}$ is:

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

where α_i are preference weights, and $Q = qn$ captures the household's son preference.

Each parent has a time endowment normalized to one, allocated between leisure and labor. Mothers incur an additional time cost γn for childcare. The time constraints imply that mothers' labor hours are $h_m = 1 - \ell_m - \gamma n$, while fathers' labor hours are $h_d = 1 - \ell_d$. The household's budget constraint is:

$$c_m + c_d + p_q Q \leq I + w_m h_m + w_d h_d, \tag{2}$$

where w_m and w_d are wages; I is non-labor income; p_q is the marginal cost of increasing the effective number of sons, Q, encompassing economic and psychological costs of sex selection.

The household maximizes a weighted sum of the parents' utilities:

$$\max_{\{c_m, c_d, \ell_m, \ell_d, n, q\}} \lambda_m U_m + \lambda_d U_d, \tag{3}$$

subject to the budget and time constraints, where $\lambda_m + \lambda_d = 1$.

The derivations of the optimal choices are provided in the online Appendix B. Our model yields twelve predictions regarding the effects of changes in wages and other parameters on sex selection (q), mother's labor supply (h_m), and fertility (n).

First, higher maternal wages (w_m) increase sex selection, $\frac{\partial q}{\partial w_m} > 0$, because higher w_m raises the opportunity cost of time spent on childcare, leading mothers to prefer having fewer children but ensuring the presence of a son. In contrast, paternal wages (w_d) and non-labor income (I) do not affect sex selection, since they do not alter the mother's time constraint or opportunity cost.

Second, the mother's labor supply responds positively to increases in her own wage ($\frac{\partial h_m}{\partial w_m} > 0$) to capitalize on higher wages. Conversely, higher paternal wages and non-labor income reduce the mother's labor supply ($\frac{\partial h_m}{\partial w_d} < 0$, $\frac{\partial h_m}{\partial I} < 0$) through income effects that make leisure more affordable.

Third, higher maternal wages reduce fertility ($\frac{\partial n}{\partial w_m} < 0$) due to the increased opportunity cost of childcare time. In contrast, higher paternal wages and non-labor income increase fertility ($\frac{\partial n}{\partial w_d} > 0$, $\frac{\partial n}{\partial I} > 0$) through pure income effects. The cost of sex selection does not affect fertility as it influences only the quality dimension (q).

These predictions, which are summarized in Table 1, contrast with alternative theories. Our model predicts that higher maternal wages lead to increased sex selection and labor supply but reduced fertility, due to the mother's dual role in childcare and work.

In contrast, models such as Almond et al. (2019) emphasize income effects, where higher non-labor income or paternal wages increase both fertility and sex selection. However, in our model, these factors do not affect the demand for sons; they only contribute to an increase in the overall number of children.

Furthermore, Qian (2008) suggests that higher returns to female labor increase the preference for daughters. Our model interprets this result in two ways. First, if higher returns to female labor affect only the mother's wage (w_m) relative to the father's, then—in the Vietnamese context—our model predicts the opposite: an increase in the mother's wage, holding the spouse's wage and non-labor income constant, leads to a higher demand for sons. This outcome arises from the trade-off mothers face between childcare and work.

Second, it is challenging to determine whether parents base their preferences on their own relative wages (i.e., father's wage versus mother's wage) or on the relative returns in the local labor market (which affect the relative returns to sons versus daughters). The latter is captured by the opportunity cost of sex selection (p_q), as higher returns to female labor suggest that daughters can contribute more to their parents' lifetime income and old-age benefits. Therefore, our empirical design accounts for these local-level opportunity costs (p_q) and the parents' relative wages.

Finally, we extend this framework in online Appendix B.1 to allow bargaining weights to depend on relative wages ($\lambda(w_m/w_d)$) and preferences to differ by gender for both fertility ($\alpha_{nm} \neq \alpha_{nd}$) and sons ($\alpha_{qm} \neq \alpha_{qd}$). When parents share identical preferences, the predictions remain similar to our baseline model even with endogenous bargaining weights, as the asymmetric childcare burden drives the results. The predictions change substantively only when mothers have a stronger relative

preference for sons ($\alpha_{qm}/\alpha_{nm} > \alpha_{qd}/\alpha_{nd}$) and prefer fewer children ($\alpha_{nm} < \alpha_{nd}$) than fathers. As shown in Panel C of Table 1, this bargaining model can generate most of our empirical findings– mother's wages increase sex selection and reduce fertility, while father's wages increase fertility. However, it also predicts that fathers' wages should reduce sex selection by shifting bargaining power away from mothers, contrary to our empirical finding of no effect. We return to discuss this extension after presenting our main results.

3 Data and Research Design

3.1 Trade Liberalization and Measures of Exposure

The 2001 US-Vietnam Bilateral Trade Agreement (BTA) provides an ideal setting to study how trade liberalization affects fertility and sex selection decisions. Before the BTA went into effect on December 10, 2001, Vietnamese exports to the US faced high tariffs under the Column 2 schedule applied to non-market economies. The BTA granted Vietnam Most Favored Nation (MFN) status, dramatically reducing US tariffs on Vietnamese exports while having minimal impact on Vietnam's import tariffs.[1]

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

First, we measure direct exposure through parents' industries following McCaig et al. (2022):

$$\tau_j = \log(1 + \text{Column } 2_j) - \log(1 + \text{MFN}_j) \tag{4}$$

where j denotes a parent's industry. A higher value of τ_j indicates a larger effective tariff cut for that industry. Since mothers and fathers often work in different industries, this measure allows us to capture gender-specific changes in labor market returns, as McCaig and Pavcnik (2018) and Fukase (2013) document that the BTA increased both formal employment and wages.

Second, we construct a province-level shift-share measure to capture broader local economic effects:

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

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

Third, we create a female-specific exposure measure to capture changes in market returns to

[1]While Vietnam was obligated to reduce tariffs mainly on agricultural and food products, these reductions were quite minimal, ranging from 0.03 to 2.7 percentage points. In addition to tariffs, Vietnam agreed to eliminate various import quotas. Most of these quotas had, however, been eliminated by the end of 2002 (STAR-Vietnam, 2003).

women's work:

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

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

3.2 Data

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

Our main sample uses repeated cross-sections of approximately 4,000 infants aged 0-1 years from each survey wave. For these infants, we observe their gender and their parents' contemporaneous industry codes and work hours. Although our first observation is from 2002, the survey's 12-month recall period means the 2002 data largely reflects pre-BTA conditions since the agreement took effect in December 2001.

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

In addition to the gender of each infant and mothers' work hours, we need a measure of fertility without data on completed birth history. From the cross-sectional data, we create a retrospective panel that tracks annual birth outcomes. Specifically, for each woman in a given survey year, we construct observations spanning from when she turned 20 until her survey year, with our key outcome being whether she gave birth in a given year. This approach allows us to control for

mother-specific fertility preferences through mother fixed effects, which is particularly important for fertility decisions that may reflect persistent individual characteristics.

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

3.3 Identification Strategy

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

$$\begin{aligned} y_{imdt} = \beta_1 \tau_m \times \text{Post}_t + \beta_2 \tau_d \times \text{Post}_t + \beta_3 \tau_p \times \text{Post}_t + \beta_4 \tau_p^w \times \text{Post}_t \\ + \alpha_t + \delta_p + \gamma_{j(m)} + \gamma_{j(d)} + X'_{imdt}\lambda + \delta S_p \times t + \varepsilon_{imdt} \end{aligned} \tag{5}$$

where the dependent variable y_{imdt} represents outcomes for individual i born to mother m and father d in year t. The indicator Post_t equals 1 for years after 2002. For child gender, y_{imdt} is an indicator for male births. For the mother's labor supply, the outcome variable y_{mdt} measures monthly work hours in the mother's main job. The sample here includes the mothers of the infants previously examined in the sex selection analysis. The treatment variables capture exposure through multiple channels: τ_m and τ_d measure industry-specific tariff cuts faced by mothers and fathers in their respective industries $j(m)$ and $j(d)$ as in (4), while τ_p and τ_p^w capture province-level and female-specific province exposure respectively. The specification includes province fixed effects δ_p, year fixed effects α_t, and industry fixed effects for both parents $\gamma_{j(m)}, \gamma_{j(d)}$.

For fertility outcomes, we construct a retrospective panel that tracks mothers' birth histories over time. We modify the baseline specification by replacing the outcome with an indicator variable Birth_{mt} for whether mother m had a birth in year t, and by substituting mother fixed effects μ_m for the province and industry fixed effects. This allows us to control for time-invariant fertility preferences.

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

($\beta_4 < 0$) and increase mother labor supply ($\beta_4 > 0$) but have no effect on fertility.

We also leverage a short and small panel of women aged 20-40 who gave birth both before and after 2002. Following McCaig and Pavcnik (2018), we assign each woman to her initial industry of employment (in 2001, prior to the BTA) when constructing tariff cuts. This approach helps mitigate potential bias from post-shock sorting—for example, if workers relocate from low-tariff-cut sectors (such as agriculture) to high-tariff-cut sectors (like textiles) in response to trade liberalization. Specifically, we estimate:

$$\begin{aligned} y_{iht} = \beta_1 \tau_i \times \text{Post}_t + \beta_2 \tau_h \times \text{Post}_t + \beta_3 \tau_p \times \text{Post}_t + \beta_4 \tau_p^w \times \text{Post}_t \\ + \alpha_t + \mu_i + X'_{iht}\lambda + \delta S_p \times t + \varepsilon_{iht}, \end{aligned} \tag{6}$$

where τ_i and τ_h are industry-specific tariff cuts for woman i and her husband (respectively) based on their 2001 industries. Individual fixed effects μ_i account for unobserved individual characteristics including fertility and sex ratio preferences. By fixing exposure to initial industries, we aim to isolate the effects of tariff cuts from subsequent shifts in employment choices.

Our baseline controls X include mothers' age and education, as well as indicators for urban locations and minority status. The term S_p represents a vector of 1999 employment shares in agriculture and manufacturing in province p. Their interactions with the time trend t account for differential trends across provinces with varying industrial compositions. All regressions are weighted by sampling weights, with standard errors clustered at the province level.

Our identification leverages cross-sectional variation in exposure to tariff cuts across industries and provinces, and temporal variation before versus after BTA implementation. The key identifying assumption is that outcomes would have evolved similarly across industries and provinces with different exposure levels absent the BTA. For the panel analysis, we assume women's pre-BTA industry choices are not endogenous to anticipated policy effects.

4 Results

4.1 Baseline Estimates

Table 2 presents our main difference-in-differences estimates. These results strongly support our model's main predictions over alternative explanations summarized in Table 1.

For mothers experiencing median exposure tariff cuts ($\tau_m = 0.07$), our coefficient estimate implies that the probability of having a male birth rises by 1.8 percentage points (t-stat = 2.5). This shift translates into an increase in the sex ratio at birth from 107 to roughly 115 boys per 100 girls. These mothers also work 2.6 more hours per month ($= 36.6 \times 0.07$), a 1.6% increase from the baseline of 159.6 hours, and reduce their likelihood of giving birth by 0.07 ($= -0.01 \times 0.07$) percentage points.[2]

[2]The sample of mothers here is smaller than that of children in the sex selection analysis because of maternity

These estimates are in line with findings from other studies. For instance, Anukriti (2018) discovered that an Indian program aimed at addressing the fertility-sex ratio trade-off resulted in a 1-2.3% increase in the probability of the first birth being male. Almond et al. (2019) found that land reform in China between 1978 and 1984 led to a 3 percentage point or 5.6% increase in the fraction of males following the birth of a first girl from a baseline sex ratio of about 112 for this group of children. However, the estimates from these studies are specific to birth parity and the gender composition of the first sibling, whereas our findings are estimated across all birth orders.

The effects of fathers' industry exposure (τ_d) show a different pattern, generally consistent with our model's predictions. While fathers' exposure has no significant effect on child sex ($-0.05, t$-stat $= -0.5$), it increases fertility ($0.02, t$-stat $= 3.6$) without affecting mothers' work hours ($19.6, t$-stat $= 1.3$).

The female-specific provincial exposure (τ_p^w) increases mothers' work hours ($20.4, t$-stat $= 2.1$) but shows no significant effect on sex selection or fertility, offering limited support for the female labor market exposure mechanism in this context. Furthermore, general provincial exposure (τ_p) which proxies for increases in nonlabor income reduces work hours ($-20.4, t$-stat $= -2.1$) and increases fertility ($0.01, t$-stat $= 4.5$) without any effect on sex selection ($0.03, t$-stat $= 0.8$).

Figure 2a supports our identification strategy by showing no pre-trends in the effect of all exposures on sex selection. For each row, we display year-by-year coefficients relative to 2001 for mother's industry exposure (τ_m), father's industry exposure (τ_d), province-level exposure (τ_p), and female-specific province exposure (τ_p^w) using the event-study version of regression (5).

The coefficients on mothers' industry exposure (top row) show no significant differences across industries with varying exposure between 2000 and 2003. This pattern is consistent with fertility being a slow-adjusting outcome. Effects emerge after 2004 and remain stable, with magnitudes stabilizing around 0.3-0.5 percentage points higher probability of male birth per percentage point of tariff reduction. This persistence suggests a permanent rather than temporary shift in sex selection behavior.

The remaining exposure measures show minimal effects. Father's industry exposure exhibits brief deviations from zero—positive in 2004 and negative in 2011—but no consistent pattern. Both province-level measures remain statistically insignificant, reinforcing our model's main predictions.

4.2 Heterogeneity

We examine how the effect of maternal industry exposure on sex selection varies across subgroups. For each subgroup, we estimate our baseline specification (5) separately. Figure 2b reveals five key patterns of heterogeneity. First, the effect of mother's industry exposure on male births is stronger in rural areas compared to urban areas. This difference aligns with evidence of stronger son preference in rural areas.

leave. Since the 2010 wave, VHLSS has asked about the number of work hours within the last 30 days. If a mother is on maternity, she will drop out of the sample, although her child appears in the infant sample.

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

Third, birth order matters. While the effect is positive yet insignificant for first births, it increases for second and higher-order births.[3] This pattern aligns with intensified sex selection when parents lack a son. Additionally, the effects are concentrated among mothers with less than nine years of education, indicating that more educated mothers may be less likely to engage in sex selection due to economic incentives.

Finally, there is minimal difference in results between mothers above and below age 30. While the effect appears slightly stronger among younger mothers (less than 30), the differences are not statistically significant.

While these patterns align with our expectations about where son preference might be strongest, we cannot reject that the effects are statistically equal across subgroups. This finding suggests that the impact of trade liberalization on sex selection operates similarly across different populations, consistent with our simple model's prediction that maternal wage increases affect sex selection independently of the strength of son preference.

4.3 Robustness

Region-Year FE A potential concern is that improvements in transportation infrastructure could affect access to sex-selective services independently of trade exposure. Vietnam's infrastructure development typically follows eight administrative regions: Northeast, Northwest, Red River Delta, North Central Coast, South Central Coast, Central Highlands, Southeast, and Mekong River Delta. To account for this, we add region-by-year fixed effects to control for any region-specific time-varying factors, including infrastructure development.

As shown in Table A4, our main results remain stable with these additional controls. The effect of maternal industry exposure on the probability of male births, mother's labor supply, and fertility are similar in magnitude and significance to our baseline estimates. This suggests our findings are not driven by concurrent regional infrastructure improvements that could facilitate access to sex-selection services.

Clustering We examine robustness to alternative clustering approaches. When we cluster standard errors at the mother's industry level instead of the province level, our key finding–that maternal exposure to tariff cuts increases sex selection– becomes more statistically significant ($t = 3.3$) as

[3]Birth of order 3 or higher is rare in the sample. By 2000, the total fertility rate is 2.01 in Vietnam.

shown in Table A5. However, this alternative clustering reduces the statistical significance of the effects on maternal labor supply and fertility decisions, while all other findings maintain similar significance level. Nonetheless, since this approach fails to account for important geographic correlation and province-level treatment components, we maintain the more conservative province-level clustering in our baseline specifications.

Weighting We also examine sensitivity to the use of survey weights and report the results in Table A6. When excluding sampling weights from our baseline specification, we find that all main results remain qualitatively and quantitatively similar at comparable levels of statistical significance. This robustness is reassuring and unsurprising. As discussed in Solon et al. (2015), unweighted estimates can be as good or even better than weighted ones when sampling is not strongly correlated with the heterogeneity of treatment effects. In our context, while the VHLSS oversamples certain geographic areas, the effects of trade exposure on sex selection appear relatively stable across subgroups (Figure 2b). Nevertheless, we maintain the use of sampling weights in our preferred specifications since they ensure our estimates are representative of the population-level impacts of trade liberalization.

Logit For our main binary outcome–whether a birth is male–we examine robustness to using logit estimation instead of OLS. Due to the number of fixed effects in our specification, we employ logit rather than probit estimation. As reported in Table A7, the logit estimates yield similar results to our OLS specifications and continue to support our model's predictions regarding the impact of maternal industry exposure on sex selection. This robustness suggests that our main findings are not driven by the linear probability model specification.

Panel Analysis Finally, we estimate the panel specification in (6), which uses *initial* industry affiliation to mitigate potential endogenous sorting and incorporates individual fixed effects to account for unobserved preferences. Comparing the panel estimates with the cross-sectional estimates offers insight into whether sorting biases our findings. If women with stronger son preferences systematically moved into industries with higher tariff cuts, the cross-sectional estimates would reflect both "true" treatment effects and selection. In contrast, the panel approach assigns each woman to her 2001 industry, so any subsequent movement to a high-exposure sector does not inflate her measured exposure. Consequently, if sorting were driving the cross-sectional results, we would expect *smaller* treatment effects in the panel.

Instead, we find that the panel estimates are at least as large as, or even larger than, the cross-sectional estimates. As shown in Table 3, women initially employed in industries with greater tariff cuts are significantly more likely to have male children post-BTA, with coefficients of 1.1–1.4 percentage points per one-percentage-point increase in tariff cuts. These coefficients exceed our cross-sectional estimates, suggesting that sorting does not substantially inflate the impact of trade liberalization on sex selection. Moreover, the larger magnitude likely reflects that our panel sample

consists of women who had at least one child before 2002—a group under more pressure to deliver a son, given the heightened work–childcare tradeoff.

This panel analysis, which follows the same women over time, allows us to examine how prior son status affects the impact of trade liberalization. Column (3) of Table 3 includes an indicator for whether a woman already had a son before her current birth. Consistent with the model's predictions, having a prior son substantially dampens the effect of trade liberalization on sex selection–the coefficient on maternal industry exposure ($\tau_m \times Post$) drops by about 40 percent (from 1.4 to 0.81) and becomes statistically insignificant. The coefficient on having a prior son is large and highly significant at -0.53, indicating these women are much less likely to have another son. The result suggests that once a family achieves their desired goal of having a son, their sex selection behavior becomes less responsive to economic incentives from trade shocks.

For labor supply, the positive coefficient on mother's industry exposure (30-33 hours per month) aligns with our cross-sectional findings, though the estimates are less precise. We also find stronger positive effects of female-specific province exposure (42-50 hours) on work hours, consistent with improved labor market conditions for women. These labor supply responses are particularly notable given that these women maintained both employment and childbearing through the policy change.

For fertility outcomes, where we expand our sample to include all women aged 20-40 (not just those with births in both periods), we find small and statistically insignificant coefficients that differ from the negative effects in our cross-sectional analysis. The larger sample size of more than 10,000 observations reflects this broader inclusion criterion. However, the short panel spanning only 2002-2006 captures fertility decisions over just 4 years, while the cross-sectional analysis spanning 2002-2016 better reflects completed fertility choices. This difference in time horizon likely explains why we find muted fertility responses in the panel.

The results reinforce our core finding that greater exposure to trade liberalization through maternal employment leads to increased sex selection, longer work hours, and reduced fertility—a pattern consistent with an intensified quantity-quality tradeoff in child-rearing decisions.

5 Discussion

Our empirical results strongly support the baseline quantity-quality model over alternative explanations. The key evidence comes from the asymmetric effects of parental exposure to trade liberalization: while maternal industry exposure affects both sex selection and fertility, paternal exposure influences only fertility. This pattern is difficult to reconcile with standard theories of son preference that emphasize income effects or relative returns to daughters

The results also present challenges for bargaining-based explanations. While a model with endogenous bargaining weights could match our findings if mothers prefer fewer children and have stronger relative son preference, such a model would predict that fathers' wages reduce sex selection by shifting bargaining power away from mothers–contrary to our finding of no effect.

The patterns from heterogeneous effects and panel analysis align with our childcare burden mechanism. Together, these results suggest Vietnam's trade liberalization affected demographic outcomes primarily through mothers' opportunity cost of time rather than through bargaining power or returns to daughters.

While our analysis focuses on Vietnam's 2001 trade liberalization, the mechanisms we identify likely apply more broadly. Many countries combine strong son preference with highly gendered childcare responsibilities, even in developed countries (Doepke et al., 2023). Our results suggest that rapid economic changes that increase female labor market opportunities without addressing this underlying asymmetry could produce similar demographic responses.

However, specific institutional features may matter. Vietnam's high female labor force participation and strong state capacity for family planning services, including abortion, may have facilitated rapid responses to changing economic incentives. The effects might differ in contexts with lower baseline female employment or weaker family planning infrastructure. Understanding these interactions between institutional context and economic incentives remains important for future research.

6 Conclusion

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

Our findings highlight how trade liberalization can create economic opportunities for women while exacerbating gender imbalances in son-preferring societies. This suggests that policies aimed at promoting female labor force participation in developing economies should be coupled with efforts to address underlying cultural preferences for sons. Future research could explore how various policy tools, such as childcare subsidies, might help mitigate these unintended demographic consequences of economic development.

References

ActionAid-Vietnam. (2016). *Unpaid Care Work: Make a House becomes a Home* (tech. rep.). ActionAid Vietnam. Hanoi. Retrieved October 7, 2022, from https://vietnam.actionaid.org/en/publications/2016/policy-brief-unpaid-care-work-make-house-becomes-home

Almond, D., Li, H., & Zhang, S. (2019). Land reform and sex selection in china. *Journal of Political Economy*, *127*(2), 560–585.

Anukriti, S. (2018). Financial incentives and the fertility-sex ratio trade-off. *American Economic Journal: Applied Economics*, *10*(2), 27–57.

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

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

Autor, D., Dorn, D., & Hanson, G. (2019). When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men. *American Economic Review: Insights*, *1*(2), 161–178.

Becker, G. S., & Lewis, H. G. (1973). On the Interaction between the Quantity and Quality of Children. *Journal of Political Economy*, *81*(2, Part 2), S279–S288.

Chakraborty, T. (2015). Trade Liberalization in a Traditional Society: Implications for Relative Female Survival. *World Development*, *74*, 158–170.

Chao, F., Gerland, P., Cook, A. R., & Alkema, L. (2019). Systematic assessment of the sex ratio at birth for all countries and estimation of national imbalances and regional reference levels. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(19), 9303–9311.

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

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

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

Ebenstein, A. (2010). The "Missing Girls" of China and the Unintended Consequences of the One Child Policy. *Journal of Human Resources*, *45*(1), 87–115.

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

Fukase, E. (2013). Export Liberalization, Job Creation, and the Skill Premium: Evidence from the US-Vietnam Bilateral Trade Agreement (BTA). *World Development*, *41*(1), 317–337.

Jayachandran, S. (2015). The Roots of Gender Inequality in Developing Countries. *Annual Review of Economics*, *7*(1), 63–88.

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

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

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

Lin, M. J., Liu, J. T., & Qian, N. (2014). More missing women, fewer dying girls: The impact of sex-selective abortion on sex at birth and relative female mortality in Taiwan. *Journal of the European Economic Association*, *12*(4), 899–926.

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

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

Qian, N. (2008). Missing women and the price of tea in China: The effect of sex-specific earnings on sex imbalance. *Quarterly Journal of Economics*, *123*(3), 1251–1285.

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

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

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

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

STAR-Vietnam. (2003). *An Assessment of the Economic Impact of the United States - Vietnam Bilateral Trade Agreement* (tech. rep.). The National Political Publishing House. Hanoi. Retrieved October 7, 2022, from https://pdf.usaid.gov/pdf_docs/Pnadc703.pdf

Figure 1: Sex Ratio at Birth Trends Across Countries

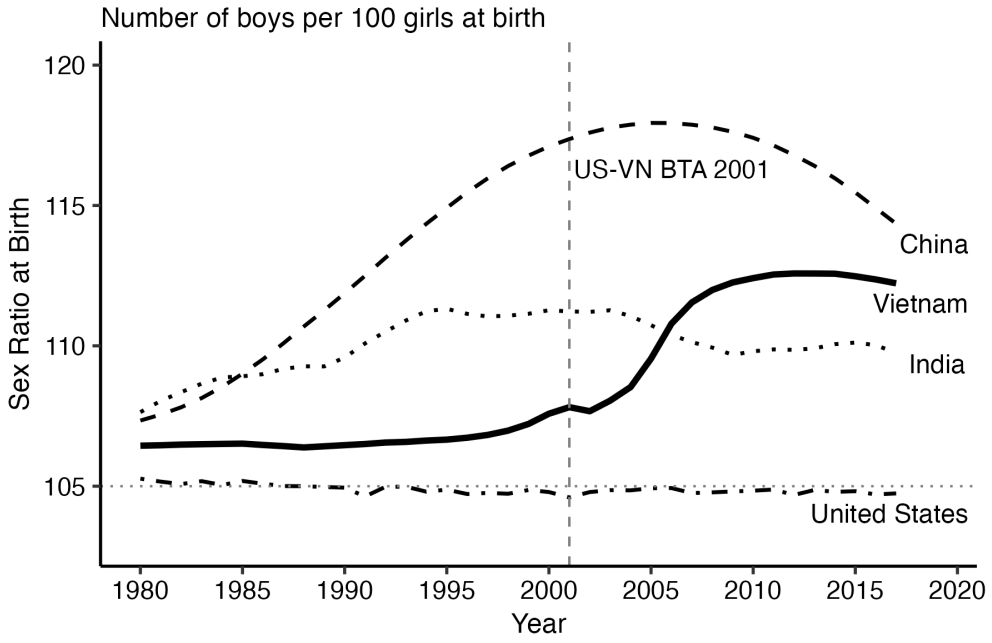

Data source: Chao et al. (2019)
Notes: The figure shows the sex ratio at birth (boys per 100 girls) for several countries. The vertical dashed line indicates the year when the bilateral trade agreement was signed. The horizontal dotted line is the biological theoretical sex ratio.

Figure 2: Effects of Trade Liberalization on Male Births

(a) Event Study of Mother's Industry Exposure to BTA on Male Birth

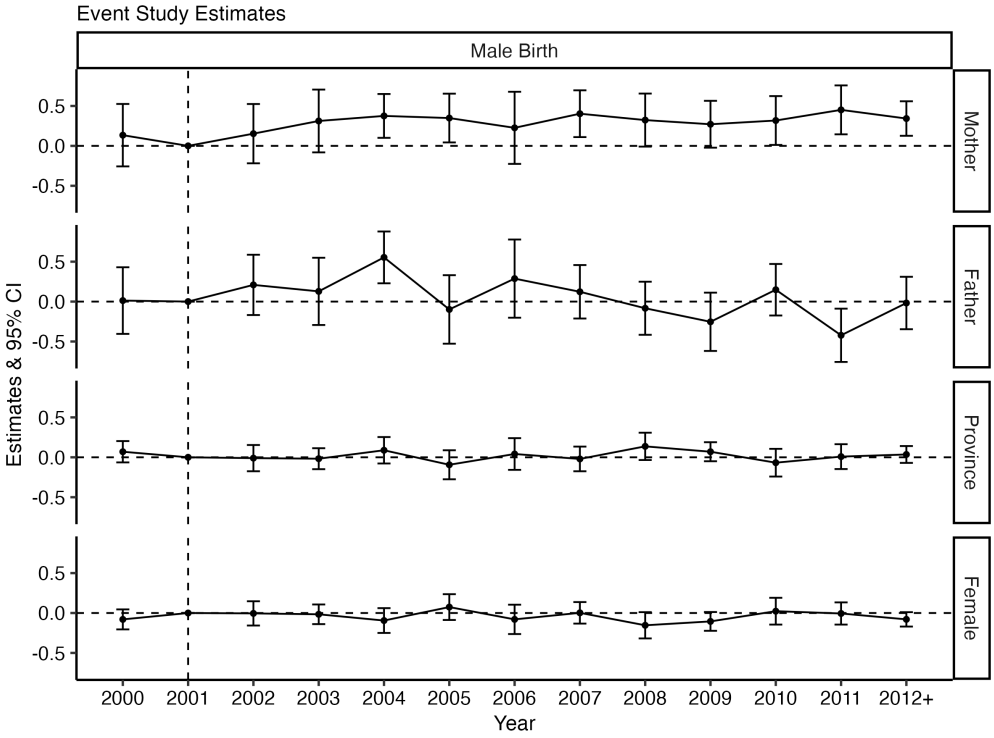

(b) Heterogeneous Effects of Mother's Industry Exposure

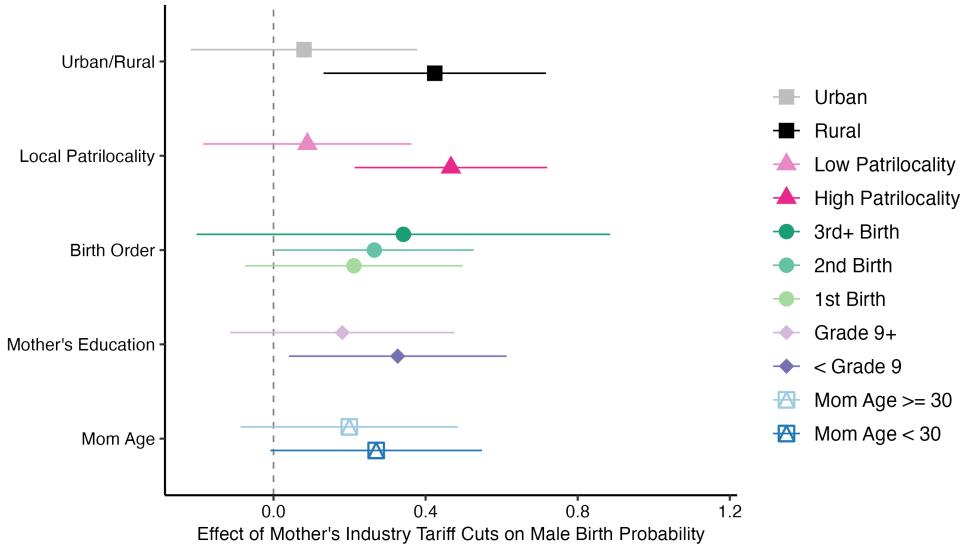

Notes: Panel (a) plots the coefficients from event study regressions showing how the effect of maternal industry exposure to trade liberalization evolves over time, with 2001 as the reference year. Panel (b) shows how the effect of maternal industry exposure varies across different groups. For each subgroup, we estimate the coefficient on maternal industry exposure ($\tau_m \times Post$) from our baseline specification (5). Both panels report 95% confidence intervals using standard errors clustered at the province level. Sample includes children aged 0-1 from VHLSS 2002-2016.

Table 1: Key Predictions of Our Model and Alternative Theories

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

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

Table 2: Effects of Trade Liberalization on Sex Selection, Labor Supply, and Fertility

	Male Birth		Mom Work Hrs		Any Birth	
	(1)	(2)	(3)	(4)	(5)	(6)
$\tau_m \times$ Post	0.26	0.26	34.9	36.6	-0.01	-0.01
	(2.5)	(2.4)	(2.3)	(2.3)	(-3.4)	(-3.1)
$\tau_d \times$ Post	-0.05	-0.05	18.9	19.6	0.02	0.02
	(-0.45)	(-0.47)	(1.2)	(1.3)	(3.6)	(3.6)
$\tau_p \times$ Post	0.03	0.02	-17.8	-20.4	0.004	0.01
	(0.83)	(0.51)	(-1.7)	(-2.1)	(0.83)	(3.7)
$\tau_p^w \times$ Post	-0.03	-0.03	22.2	20.4	-0.003	-0.005
	(-0.98)	(-0.79)	(2.2)	(2.1)	(-0.60)	(-1.7)
R^2	0.009	0.010	0.33	0.34	0.28	0.28
Observations	28,390	28,389	26,129	26,128	1,911,566	1,911,566
Control Mean	0.58	0.58	157.9	157.9	0.01	0.01
Mom FE					✓	✓
Birth Year FE	✓	✓			✓	✓
Survey Year FE			✓	✓		
Province, Industry FEs	✓	✓	✓	✓		
Controls		✓		✓		✓

Notes: Table reports difference-in-differences estimates from (5). Dependent variables are: indicator for male child (columns 1-2), mother's monthly work hours (columns 3-4), and indicator for birth in past two years (columns 5-6). τ_m, τ_d are mother's and father's industry tariff cuts, τ_p is province exposure, and τ_p^w is female-specific province exposure. T-statistics based on standard errors clustered by province in parentheses. Controls include mothers' age and education, indicators for urban location and minority status, and linear time trends of 1999 agriculture and manufacturing shares. Data from VHLSS 2002-2016.

Table 3: Panel Evidence on Trade Liberalization and Sex Selection

	Male Birth			Work Hrs		Any Birth	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\tau_m \times$ Post	1.1	1.4	0.81	32.9	30.1	0.03	0.03
	(2.0)	(2.6)	(1.6)	(0.52)	(0.44)	(0.20)	(0.22)
$\tau_d \times$ Post	0.52	0.56	0.34	-22.5	-16.8	0.06	0.05
	(0.84)	(0.87)	(0.60)	(-0.21)	(-0.16)	(0.38)	(0.35)
$\tau_p \times$ Post	-0.17	-0.28	-0.12	-39.3	-47.9	-0.006	0.006
	(-0.83)	(-1.4)	(-0.65)	(-1.3)	(-1.4)	(-0.18)	(0.12)
$\tau_p^w \times$ Post	0.14	0.18	0.05	50.2	42.5	0.02	0.005
	(0.71)	(1.0)	(0.32)	(1.7)	(1.4)	(0.50)	(0.15)
Had a son before			-0.53				
			(-9.0)				
R^2	0.82	0.84	0.86	0.76	0.76	0.43	0.43
Observations	2,470	2,470	2,470	2,388	2,388	10,572	10,572
Control Mean	0.56	0.56	0.56	165.1	165.1	0.04	0.04
Individual FE	✓	✓	✓	✓	✓	✓	✓
Birth Year FE	✓	✓	✓				
Survey Year FE	✓	✓	✓	✓	✓	✓	✓
Controls		✓	✓		✓		✓

Notes: Table reports estimates from panel regressions (6) using women's and their husbands' pre-BTA industries in 2001. Dependent variables are an indicator for male birth (columns 1-3), monthly work hours (columns 4-5), and an indicator for birth in the past 2 years (columns 6-7). Independent variables and controls are described in Table 2. T-statistics based on standard errors clustered by province in parentheses. Sample: Women 20-40 years old; Data: VHLSS Panel 2002-2006