Essays in Labor and International Economics

by

Ngan Nguyen Phuong Tran

A dissertation accepted and approved in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
in Economics

Dissertation Committee:
Anca Cristea, Chair
Kathleen Mullen, Core Member
Woan Foong Wong, Core Member
Tuong Vu, Institutional Representative

University of Oregon
Spring 2026

©2026 Ngan Nguyen Phuong Tran

All rights reserved.

DISSERTATION ABSTRACT

Ngan Nguyen Phuong Tran

Doctor of Philosophy in Economics

Title: Essays in Labor and International Economics

This dissertation examines three related topics at the intersection of international trade policy and labor market outcomes, with particular attention to gender inequality and distributional effects.

The first chapter, *Unequal Gains: The Gendered Impact of Trade Liberalization in Vietnam*, studies the labor market consequences of Vietnam's 2001 Bilateral Trade Agreement with the United States. Exploiting cross-provincial variation in exposure to tariff reductions, I find that increased access to export markets widened the gender wage gap, with effects concentrated among high-skilled, married, older and rural women.

The second chapter, *From Trade War to Job Growth: Evidence from Vietnam*, examines a later episode of global trade reallocation during the 2018–2019 U.S.–China trade war. Using firm-level data, we show that trade diversion toward Vietnam generated employment growth, particularly among foreign-owned firms and private firms with prior export experience, highlighting the role of firm heterogeneity in shaping labor market responses to trade shocks.

The third chapter, *Assessing Bias in Hypothetical Ability Ratings*, shifts focus from trade shocks to the role of information in shaping inequality. Using data from the 2024 American Work Capacity and Ability Survey, we investigate whether perceptions of ability are systematically biased by gender, race, and ethnicity. The results show little evidence of bias when performance is clearly observed, suggesting that evaluators rely primarily on performance signals rather than demographic characteristics.

Together, these chapters contribute to a broader understanding of how economic shocks and information environments shape inequality. They show that while trade integration can generate substantial economic opportunities, the distribution of these gains depends on worker characteris-

tics, firm structure, and institutional context.

The dissertation includes previously unpublished co-authored research.

ACKNOWLEDGMENTS

I am deeply grateful to my dissertation committee – Anca Cristea, Kathleen Mullen, Woan Foong Wong, and Tuong Vu – for their invaluable support and guidance throughout this process. Special thanks are due to my chair, Anca Cristea, who has been a steadfast mentor since my second year in the program, offering both intellectual guidance and personal encouragement at every stage of this journey.

I am also grateful to the Department of Economics at the University of Oregon for providing a supportive and collegial environment in which I develop as a researcher. I particularly thank Michael Kuhn, Alfredo Burlando, Ed Rubin, Jonathan Davis, Laura Bakkensen, Emmett Reynier, Phillip Economides, Kyutaro Matsuzawa, Jaichung Lee and Owen Jetton for generously offering their time and constructive feedback on my research.

This dissertation was supported in part by the Southeast Asian Studies Graduate Scholarship, the Global Oregon Graduate Research Award, and the Oregon Sasakawa Young Leaders Fellowship Fund (SYLFF) Graduate Fellowship for International Research. I gratefully acknowledge this financial support.

Finally, I am deeply grateful to my parents and my brother for their unconditional love and encouragement, which have sustained me throughout this journey, even from afar since I left my home country, Vietnam. I am especially thankful to my husband, Blair, whose unwavering support, patience, humor, and belief in me have been my anchor and greatest source of strength. Thank you for everything.

To my parents and Huy, whose love made this journey possible.

To Blair and Rocky, my home, my strength, and my joy.

TABLE OF CONTENTS

Chapter Page

1 Introduction 13

2 Unequal Gains: The Gendered Impact of Trade Liberalization in Vietnam 15

2.1 Introduction 15

2.2 Theory and Existing Empirical Evidence 19

2.2.1 Impact of Trade on the Raw Gender Wage Gap 19

2.2.2 Impact of Trade on the Residual Gender Wage Gap 20

2.2.3 Empirical Evidence 22

2.3 Institutional and Economic Context 25

2.3.1 Overview of the 2001 U.S. - Vietnam Bilateral Trade Agreement 25

2.3.2 Gender Disparities in Vietnam's Labor Market 27

2.4 Data 31

2.5 Empirical Methodology 34

2.5.1 Wage Regression with Trade Shock at the Industry Level 34

2.5.2 Measuring Trade Exposure at the Provincial Level 36

2.5.3 Provincial-level Analysis 39

2.5.4 Individual-level Analysis 41

2.6 Empirical Results 42

2.6.1 Provincial-level Analysis 42

2.6.2 Individual Analysis 44

2.7 Robustness of Results 47

2.8 Mechanisms and Channels of Adjustment 51

2.8.1 Heterogeneity in Trade Effects 51

2.8.2 Decomposition of the Provincial Gender Wage Gap 55

2.9 Conclusion 58

3 From Trade War to Job Growth: Evidence from Vietnam 61
3.1 Introduction 62
3.2 Institutional Background 65
3.2.1 Overview of the U.S.-China Trade War 65
3.2.2 Vietnam's Trade Patterns and Stylized Facts 66
3.3 Conceptual Framework and Empirical Evidence 72
3.3.1 Employment Effect: Theoretical Framework 72
3.3.2 Employment Effect of the Trade War: Empirical Evidence 74
3.4 Data 75
3.4.1 Vietnam firm-level data 76
3.4.2 Trade data 78
3.4.3 Trade War Tariffs 79
3.5 Empirical Methodology 79
3.5.1 Measuring Trade War Exposure at the Industry Level 79
3.5.2 Employment Effects of the U.S.-China Trade War on Vietnamese Firms . . 80
3.6 Empirical Results 82
3.6.1 Employment Effect 82
3.6.2 Worker Composition and Wage Effect 87
3.7 Conclusion 89

4 Assessing Bias in Hypothetical Ability Ratings 92
4.1 Introduction 93
4.2 Related Literature 95
4.3 Data 97
4.4 Empirical Strategy 99
4.4.1 Detecting gender-based and race-based biases separately 100
4.4.2 Testing intersectionality of gender and race 102
4.4.3 Testing if respondents' characteristics affect biases 103

4.5 Results . . . 105
4.5.1 Descriptive Results . . . 105
4.5.2 Gender-Based Bias . . . 107
4.5.3 Race-Based Bias . . . 110
4.5.4 Intersectionality . . . 114
4.5.5 Matching of gender . . . 115
4.5.6 Matching of race/ethnicity . . . 117
4.5.7 Heterogeneity due to incorrect assessment . . . 120
4.6 Conclusion . . . 123

A Chapter II Appendix . . . 126
A.1 Supplementary Tables and Figures . . . 126
A.2 Data Appendix . . . 134
A.2.1 Changes Made to Industry Codes in 1999 Population and Housing Census (Accessed Through IPUMS–International) . . . 134

B Chapter III Appendix . . . 135

C Supplementary Tables & Figures . . . 135

D Chapter IV Appendix . . . 139

E Appendix . . . 139

REFERENCES . . . 142

LIST OF FIGURES

Figure Page

2.1 Vietnam's Trade with the United States 26
2.2 Industry-Level Tariff Cuts Sorted by Female Employment Share (Pre-BTA) 28
2.3 Employment Rate and Percentage of Wage Workers by Gender 29
2.4 Provincial BTA Tariff Exposure 38
2.5 Estimated Effects of Trade Exposure on Women's Wages (Relative to Men) 46
3.1 Progression of U.S. Import Tariffs on China during 2018-2019 Trade War 65
3.2 Distribution of Cumulative Tariff Increases at HS6 Level 67
3.3 Vietnam's Aggregate Trade 69
3.4 Vietnam's Aggregate Trade 70
3.5 Changes in Vietnam's Exports by Sector 71
3.6 Changes in Vietnam's Imports by Sector 72
3.7 Event-study estimates of trade war exposure on firm employment 86
4.1 Average Ability Ratings by Randomized Performance and Gender 106
4.2 Average Ability Ratings by Randomized Performance and Race/ Ethnicity 106
A.1 Vietnam's Export 126
A.2 Wage Distribution by Gender 127
A.3 Provincial Trade Exposure vs. Population (1999) 128
A.4 Provincial Trade Exposure vs. Industrial Gross Output (1997) 128
A.5 Distribution of leave-one-province-out estimates of the coefficient on *Trade_exposure*× *Post_BTA*×*Female*. 129
C.1 Distribution of Cumulative Tariff Increases at ISIC4 Level 136
C.2 Vietnam's Exports by Destination 136
C.3 Vietnam's Imports by Destination 137
E.1 Distribution of Ratings by Gender (Unconditional) 139

E.2 Distribution of Ratings by Race (Unconditional) 141

LIST OF TABLES

Table Page

2.1 Empirical Evidence on Trade Liberalization and Gender Outcomes 23
2.2 Average Wages and Log Wage Gaps by Gender and Year 30
2.3 Estimated Effects of Industry-Level Trade Shocks on Individual Wages 35
2.4 Estimated Effects of Provincial Trade Shocks on Provincial Wage Premiums for Female Workers 43
2.5 Estimated Effects of Provincial Trade Shocks on Provincial Wage Premiums for Male Workers 44
2.6 Estimated Effects of Trade Exposure on Individual Log Real Wages 45
2.7 Heterogeneity in the Effect of Trade Exposure on Women's Wages 52
2.8 Heterogeneity in the Effect of Trade Exposure on Women's Wages by Employer Type 54
2.9 Estimated Effects of Provincial Trade Shocks on Components of the Gender Wage Gap 57
3.1 Summary Statistics by Ownership and Year 78
3.2 Effect of Trade War Tariffs on Firm Employment 82
3.3 Effect of Trade War Tariffs on Firm Employment by Ownership 84
3.4 Effect on Female Employment 88
3.5 Effect on Average Wages 89
4.1 Effects of Female Identity on Perceived Ability 108
4.2 Pooled Regression Results with Ability and Respondent Fixed Effects 110
4.3 Effects of Race on Perceived Ability 111
4.4 Pooled Regression Results by Race/Ethnicity 113

4.5 Intersectionality in ability ratings: fully interacted specification 114
4.6 Regression results for gender matching effects (fully interacted specification) 116
4.7 Regression results for race/ethnicity matching effects (single-race sample, fully interacted specification) 118
4.8 Regression results for race/ethnicity matching effects (full sample with flexible same-race indicator) 120
4.9 Self-Assessment Accuracy by Gender and Task 121
4.10 Miscalibration and gender differences in ability ratings 122
A.1 Summary Statistics by Gender and Year 130
A.2 Estimated Effects of Trade Exposure on Individual Log Real Wages: Accounting for Survey Weights 131
A.3 Estimated Effects of Trade Exposure: Adding Industry-by-year FEs 131
A.4 Estimated Effects of Trade Exposure: Binary Treatment Specification 132
A.5 Estimated Effects of Trade Exposure: Alternative Wage Measure (Hourly Wage) . . 133
A.6 Industry Code Conversion and Number of Observations 134
C.1 Timeline of U.S. and Chinese Trade Policy Measures 135
C.2 Effect of Trade War Exposure on Employment - Unbalanced Panel 137
C.3 Effect of Trade War Tariffs on Female Employment Share by Ownership 138
C.4 Effect of Trade War Tariffs on Average Wages by Ownership 138
E.1 Names by Gender and Race/Ethnicity 139
E.2 Tail Analysis: Linear Probability Models for High and Low Ratings 140
E.3 Distribution of Reported Race/Ethnicity Indicators 140

Chapter I: Introduction

Understanding the sources of labor market inequality — who gains from economic opportunity, who is left behind, and why — is one of the central questions in economics. This dissertation contributes to that understanding from two complementary angles: the distributional consequences of trade liberalization and the role of ability perceptions in shaping how workers are evaluated.

The first chapter studies the gendered consequences of Vietnam's 2001 U.S.–Vietnam Bilateral Trade Agreement (BTA), which triggered a sharp and plausibly exogenous reduction in U.S. tariffs on Vietnamese exports. Exploiting cross-provincial variation in exposure to the BTA, I find that provinces more exposed to the trade shock experienced a modest but statistically significant widening of the gender wage gap. Women faced an additional wage penalty relative to men, concentrated among those in state-owned enterprises and among high-skilled, married, older, and rural women. Oaxaca–Blinder decompositions indicate that this penalty operates through the unexplained component of the wage gap, consistent with differences in returns to characteristics rather than observable endowments, pointing to the role of entrenched gender norms and occupational sorting in mediating the labor market consequences of trade. The findings contribute to a literature that has documented mixed effects of globalization on gender inequality, and highlight that export-expanding trade shocks do not automatically translate into broader gains for women.

The second chapter, co-authored with Anca Cristea and Hoang Pham, turns to a more recent trade disruption: the 2018–2019 U.S.–China trade war. As escalating U.S. tariffs on Chinese goods redirected demand away from Chinese exporters, Vietnam emerged as one of the principal beneficiaries of trade diversion. Using firm-level data from eight waves of the Vietnam Enterprise Survey, we find that firms in more exposed industries experienced meaningful employment growth, with effects concentrated among foreign-owned firms and private domestic firms with prior international trading experience. The employment expansion was not gender-neutral:

exposed firms disproportionately hired female workers, increasing both the level and share of female employment. At the same time, average wages declined among the most exposed firms, consistent with compositional shifts toward lower-wage workers at the margin of expansion. Together, the first two chapters paint a nuanced picture of how trade shapes labor market outcomes in Vietnam — generating aggregate gains while distributing them unevenly across firms, workers, and gender lines.

The third chapter steps back from trade to examine a different source of labor market inequality: biased perceptions of ability. Using data from the 2024 American Work Capacity and Ability Survey, co-authored with Italo Lopez Garcia, Nicole Maestas, and Kathleen Mullen, we test whether ability ratings assigned to hypothetical individuals vary systematically with gender, race, or ethnicity, holding observed performance constant. The experimental design which independently randomizes both identity-signaling names and performance scores allows us to isolate the causal effect of perceived demographic identity on ability assessments. We find little systematic evidence of bias: ratings respond strongly to performance, and this relationship is highly similar across demographic groups. These null results are consistent with models of statistical discrimination, which predict that group-level stereotypes should attenuate when precise individual performance signals are available. The findings suggest that biased ability perceptions may be most consequential in settings where performance is noisier, less comparable, or more subject to interpretation.

Overall, this dissertation contributes to our understanding of labor market inequality by documenting how external shocks, institutional structures, and individual beliefs shape who gains and who loses in the labor market. The evidence points to a consistent theme: observable signals and market opportunities can reduce certain inequalities, but structural factors – gender norms, and informational frictions – determine how broadly those gains are shared.

Chapter II: Unequal Gains: The Gendered Impact of Trade Liberalization in Vietnam

Ngan Tran

Abstract

This paper examines how trade liberalization affected gender wage inequality in Vietnam following the 2001 U.S.-Vietnam Bilateral Trade Agreement (BTA), which triggered a sharp and plausibly exogenous reduction in U.S. tariffs on Vietnamese exports. Exploiting cross-provincial variation in exposure to the BTA, constructed from pre-liberalization industry composition and industry-specific tariff cuts, and estimating a triple difference-in-differences specification, I find that provinces more exposed to the BTA experienced a modest but statistically significant widening of the gender wage gap. A one-standard-deviation increase in trade exposure led to an additional 0.8 percent decline in women's wages relative to men's. This gender penalty is concentrated among women in state-owned enterprises and among high skilled, married, older, and rural women, underscoring the role of social norms and gender role attitudes in shaping women's labor market outcomes in response to economic opportunities from trade liberalization.

2.1 Introduction

The distributional consequences of globalization has generated a large body of research; however, the focus has been placed on the comparison between workers across different skills or educational levels (Dollar and Kraay, 2004; Goldberg and Pavcnik, 2007; Caliendo et al., 2019). In contrast, less is known about how trade shocks shape labor market outcomes across gender lines, and existing findings remain mixed and context-dependent (Seguino, 2000; Berik et al., 2004; Vijaya and Kaltani, 2007; Braunstein and Brenner, 2007; Menon and Van der Meulen Rodgers, 2009; Wolszczak-Derlacz, 2013). This paper contributes to the literature by examining the gendered consequences of trade liberalization in Vietnam, leveraging the quasi-experimental setting created by the 2001 U.S.–Vietnam Bilateral Trade Agreement (BTA).

The BTA provides a unique opportunity to study the effect of a trade shock in the form of increased access to a major export market. First, upon entry into force of the agreement, Vietnam switched from Column 2 tariffs on imports to Column 1 (i.e. Most Favored Nation) tariffs, resulting in tariff reductions for Vietnam's exports to the U.S. from an average of 23% to 3.5%. The movement between pre-existing sets of tariff rates creates exogenous variation that are less susceptible to political economy concerns. Second, Vietnam had already granted the U.S. Most Favored Nation (MFN) status prior to the enactment of the BTA, suggesting that the most significant change that the BTA brings about is export opportunities without introducing import competition. Moreover, the tariff cuts were immediate as opposed to being phased in over a number of years, and effects of the BTA have been found in the short run as in McCaig (2011)[1].

Following the local labor market approach pioneered by (Topalova, 2007), I construct a Bartik-style measure of trade exposure at the provincial level, combining industry-level tariff reductions with pre-liberalization employment shares. Using five waves of the Vietnam Household Living Standards Survey (2002–2010), I combine province-level and individual-level analyses to examine how this trade shock affected both overall wage levels and the relative earnings of men and women. At the provincial level, I employ a two-stage procedure to recover gender-specific provincial wage premiums and relate them to trade exposure. At the individual level, I estimate a triple difference-in-differences model to identify gender-differentiated effects of trade liberalization while controlling for province and year fixed effects and a rich set of worker characteristics.

Results reveal substantial heterogeneity in how provinces adjusted to the BTA. Provinces more exposed to tariff reductions experienced a signficant compression of provincial wage premiums. An important nuance, however, is that this convergence is driven entirely by the informal sector; no comparable wage compression appears in state-owned, private, or foreign-invested firms. This pattern is consistent with BTA-induced transitions from informal to formal employment (McCaig and Pavcnik, 2018) and the negative selection of workers who remain informal,

[1]McCaig (2011) finds that increased access to the U.S. market decreased incidence of poverty two years after the BTA came into effect

contributing to lower average wages in that segment of the labor market.

Turning to gender differences, I find that women experienced an additional wage penalty within the same province: a one–standard-deviation increase in trade exposure lowered women's wages by roughly 0.8 percent relative to men's. Oaxaca–Blinder decompositions indicate that this effect operates primarily through the unexplained portion of the wage gap, consistent with larger differences in returns to characteristics rather than differences in observed endowments. Moreover, the penalty is not evenly distributed: it is concentrated among women working in state-owned enterprises and among married, older, and rural women, who face greater domestic obligations and mobility constraints. By contrast, no gender differentials are found in private or foreign-invested firms, suggesting that firms more exposed to market competition and international standards operate under more gender-neutral wage-setting rules. These findings highlight the central role of prevailing institutional factors, particularly entrenched gender roles and occupational sorting, in shaping unequal gains from trade in Vietnam.

This paper makes three main contributions. First, it provides new evidence to a literature that finds mixed effects of globalization on gender inequality.[2] In contrast to most work that focuses on import competition, this paper examines an export-expanding shock — an understudied margin of trade liberalization — and is the first to provide systematic evidence from Vietnam. My findings also highlight the importance of local context: the effects of the BTA vary substantially across provinces, suggesting that institutional environments and labor market structures condition how trade shocks translate into wage adjustments.

Second, this paper contributes to the emerging literature on the gendered impacts of the BTA. Using household survey data, Sakakibara (2023) shows that the BTA expanded female employment in manufacturing, especially in wearing apparel, yet these gains did not translate into greater intra-household bargaining power, as reflected in unchanged spending on female-preferred goods. Complementing this evidence, Hoang et al. (2024) use firm-level data to document a decline in

[2]The next section synthesizes existing studies on trade and gendered labor market outcomes.

wage distortions against women in manufacturing, especially in the private sector, suggesting that expanding export opportunities can induce more gender-neutral labor practices among firms. At the same time, household-level responses appear more constrained. Huynh and Ku (2025) find that the BTA reduced women's labor market participation, especially among married women in wealthier households, with stronger effects in the South than in the North, consistent with regional differences in gender role attitudes. Huynh and Nguyen (2025) further show that greater industry exposure increased women's likelihood of having male children and working longer hours, with effects driven solely by maternal exposure. Together, this line of work suggests a divergence between firm-side adjustments that may benefit women and household-level dynamics shaped by deeply entrenched gender norms, which can ultimately reinforce rather than alleviate gender inequality.

Finally, this paper adds to the growing literature on the labor market consequences of trade reforms in Vietnam. Previous studies have shown that the BTA boosted job growth, reduced poverty, and triggered structural shifts from agriculture to manufacturing (McCaig, 2011; McCaig and Pavcnik, 2018; McCaig et al., 2022). Yet the distributional implications by gender have received little attention. By filling this gap, my study offers new insights into the complex interplay between trade liberalization and gender inequality in a rapidly developing economy.

The remainder of the paper is organized as follows. Section 2 discusses the theoretical channels linking trade and gender wage gaps and reviews the related empirical literature. Section 3 provides contextual background on the BTA and gender disparities in Vietnam's labor market. Section 4 describes the data, and Section 5 outlines the empirical strategy. Section 6 presents the main results, and Section 7 reports robustness checks. Section 8 examines mechanisms and channels of adjustment, and Section 9 concludes.

2.2 Theory and Existing Empirical Evidence

Before exploring the theoretical links between trade liberalization and the gender wage gap, it is important to distinguish between two concepts: the raw gender wage gap, which reflects the overall difference in average wages between men and women, and the residual wage gap, which remains after accounting for observable characteristics such as education and experience (Wolszczak-Derlacz, 2013). Trade liberalization can influence the raw gap through shifts in the composition of employment (i.e. the expansion or contraction of industries with different factor intensities) or through technological upgrading among exporting firms.

With respect to the residual wage gap, existing literature has focused on one key source: gender-based differences in returns to characteristics, often interpreted as discrimination.[3] Greater exposure to global markets may intensify competitive pressures, but theories in economics and political economy offer competing predictions about whether increased competition alleviates or amplifies gender-based wage disparities. The following subsections outline the principal channels through which these effects may arise.

2.2.1 Impact of Trade on the Raw Gender Wage Gap

Trade liberalization may affect the raw wage gap through two primary mechanisms:

Sectoral reallocation: Reductions in trade barriers can shift the composition of industries, reallocating labor across sectors with different gender intensities. If men and women are imperfect substitutes[4], or if persistent occupational segregation exists, as documented in Anker et al. (2003); Borrowman and Klasen (2020), such reallocation can alter gender-specific labor demand. For example, in developing countries, women are often concentrated in labor-intensive, export-oriented manufacturing (Çağatay and Özler, 1995). Given this type of gender segregation, trade

[3] Broader sources of the residual wage gap include differences in occupation and industry (through occupational sorting or compensating differentials related to long-hour requirements and flexibility), differences in psychological attributes or noncognitive skills, and heterogeneity in firm-specific pay premiums (Blau and Kahn, 2017; Card et al., 2016).

[4] As modeled by Galor and Weil (1996); Sauré and Zoabi (2014); Acemoglu et al. (2004); Do et al. (2016)

liberalization is expected to improve women's labor market outcomes relative to men's if a country has a comparative advantage in female intensive sectors, and vice versa (Gaddis and Pieters, 2017).

Technological upgrading: In heterogeneous firm models such as Melitz (2003), tariff reductions encourage more productive firms to export and force out the least productive. Evidence from Verhoogen (2008) and (Bustos, 2011) shows that exporting firms also upgrade in quality and technology. Technological changes can have differential effects by gender due to male-female differences in endowments of brain versus brawn (Gaddis and Pieters, 2017; Aguayo-Tellez et al., 2014; Weinberg, 2000; Rendall, 2017). More specifically, technological upgrading may increase the demand for female labor by being complementary to cognitive skills and reducing the demand for physical skills. Using establishment level data from Mexico and exploiting tariff reductions from NAFTA, Juhn et al. (2014) show that liberalization-induced technological upgrading improve the relative wage and employment of women in blue-collar occupations, but not in white-collar occupations.

2.2.2 Impact of Trade on the Residual Gender Wage Gap

The literature on the impact of trade liberalization on the residual wage gap is divided into three main theoretical perspective.

Neoclassical theory posits that wage discrimination is costly and should be eliminated under competitive pressure (Becker, 1957). As trade increases market competition, non-discriminating firms that are more cost-efficient will expand while discriminating firms will contract and eventually die off. Consequently, rising competitiveness from international trade is expected to eliminate any discriminatory wage differentials between male and female labor in the long run, closing the residual gender wage gap (Berik et al., 2004). Empirical support for this view comes from Black and Brainerd (2004), who find that U.S. manufacturing industries more exposed to international competition experienced larger reductions in gender-based wage disparities.

Alternative institutional and bargaining theories challange this prediction. Williams (1987) argues that discrimination may persist due to entrenched institutional structures that disadvantage women. Political economy theories also highlight the role of bargaining power in wage determination, especially under incomplete contracts (Vijaya and Kaltani, 2007; Reich, 2017; Rodrik, 1999). Building on these theories, Wolszczak-Derlacz (2013); Menon and Van der Meulen Rodgers (2009); Vijaya and Kaltani (2007) argue that increasing competition introduced by trade liberalization and Foreign Direct Investment (FDI) may weaken the position of female workers by lowering their bargaining power, especially in countries where women are segregated into lower-paying, lower-status jobs. Empirical support for this argument comes from Seguino (2000), who compares Taiwan and South Korea from 1981 to 1992 and documents divergent trends in gender wage differentials: the gap widened in Taiwan's manufacturing sector but narrowed in Korea. She argues that in Taiwan, highly mobile capital in female-intensive industries weakened women's bargaining power, lowering their relative wages, whereas Korea's more constrained capital mobility pushed firms to compete through technological upgrading instead of relocating.

Monopsonistic wage theory, pioneered by Robinson (1969), has received renewed attention in recent work. A notable contribution by Hirsch (2010) revives a form of monopsonistic discrimination first proposed by Robinson (1933), offering a compelling alternative to Becker's (1957) taste-based model. While Becker's framework attributes discrimination to employer prejudice, it struggles to explain persistent wage gaps under perfect competition, where such behavior should be eliminated over time. In contrast, Robinson's model allows firms to profit from wage-setting discrimination, making such outcomes more sustainable in the long run. Whereas Becker treats wages as exogenous, Robinson's framework makes wages endogenous to the firm's hiring decisions. Building on this perspective, Hirsch et al. (2010) and Sharma (2023) find that women's labor supply to the firm is less elastic than men's, using matched employer-employee data from Germany and Brazil, respectively.

2.2.3 Empirical Evidence

A growing body of empirical work examines how trade liberalization shapes gender wage outcomes across a variety of contexts. As pointed out by Dix-Carneiro and Kovak (2023), existing evidence is fragmented across contexts, outcomes, and types of trade shocks. Table 2.1 synthesizes key contributions, outlining the trade shocks studied, main findings, and proposed mechanisms. Although the evidence is mixed, several patterns emerge. Competition-enhancing reforms such as tariff reductions that increase import penetration often compress gender gaps in more competitive industries by limiting the scope for taste-based discrimination, yet they can exacerbate disparities where institutional frictions, labor segmentation, or weak bargaining power persist. In contrast, export-expanding reforms often induce firms to upgrade technologies and production processes, which can raise female productivity when these changes complement cognitive or non-physical skills, thereby improving women's labor market outcomes. Overall, the heterogeneity in results highlights the central role of country-specific institutions, labor market frictions, and pre-existing gender norms in mediating how trade liberalization translates into gendered wage outcomes.

Table 2.1. Empirical Evidence on Trade Liberalization and Gender Outcomes

Country	Study	Trade / Policy Shock	Main Results	Mechanism
Developed Countries				
Denmark	Keller and Utar (2022)	China import competition (MFA quota removal)	Deepen long-run gender inequality: women's earnings fall; men unaffected	Displacement triggers fertility among women near end of biological clock; high reentry costs
U.S.	Black and Brainerd (2004)	Import competition (1980s)	Faster declines in GWGs in import-competing industries	Competition raises cost of discrimination
U.S.	Brussevich (2018)	China import competition	Men harmed more; women gain relatively more in welfare	Gendered intersectoral mobility costs drive differential transitions
U.S.	Sauré and Zoabi (2014)	NAFTA expansion	Wider GWG, lower FLFP in some settings	Capital–labor dilution lowers female productivity in expanding sectors
Italy	Chan (2018)	China import competition	Wider GWG	Female-intensive firms shrink/exit; women face higher adjustment costs
Norway	Bøler et al. (2015)	Paternal leave reform	Exporters' higher GWG declines post-reform	Exporters facing global competition value commitment more; reform reduces perceived female commitment gap
Transitioning Economies				
Taiwan & Korea	Seguino (2000)	Export-led liberalization (1981–1992)	GWG narrowed in Korea, widened in Taiwan	Taiwan: capital mobility weakens women's bargaining power; Korea: restricted capital mobility promoted upgrading
Cross-Country Studies				
80 countries	Oostendorp (2009)	1983-1999	Trade openness narrows GWG, particularly for low-skilled workers globally and for high-skilled workers in richer economies only	
18 OECD countries	Wolszczak-Derlacz (2013)	1970–2005	Market concentration widens gender gaps; trade narrows high-skill gaps in concentrated industries but widens low-skill gaps	
19 countries	Vijaya and Kaltani (2007)	1987-2001	FDI lowers manufacturing wages, with stronger effects on women	

Table 2.1 (continued)

Country	Study	Trade / Policy Shock	Main Results	Mechanism
Developing Countries				
Brazil	Gaddis and Pieters (2017)	1980s–1990s liberalization	Employment and LFPR fall for both; no relative female gain	Reallocation out of male tradables; norms and segregation limit women's gains
Cambodia	Erten and Keskin (2024)	WTO accession	Male paid work falls; women's labor force entry rises	Due to reduced men's earnings, women entered labor force via added-worker effect
Chile	Banerjee et al. (2025)	Chile–Mexico FTA	New exporters increased female white-collar employment by 10%	Exporting induces tech upgrading favoring cognitive tasks
China	Chen et al. (2013)	Trade/FDI exposure	More female employment; smaller wage gaps in exporters	Competition reduces discrimination; wage gaps reflect productivity
China	Braunstein and Brenner (2007)	1990s FDI/export boom	Early FDI favors women; later gains shift to men	Expansion interacted with gender segregation to favor male-intensive sectors
Colombia	Ederington et al. (2009)	1984 tariff cuts	Exporters hired more women, bigger tariff cuts boosted female employment more	Competition raises cost of discrimination, prompting firms to hire more women
India	Menon and Van der Meulen Rodgers (2009)	1991 liberalization	GWG widens in concentrated industries	Firms maintain male premiums where market power is high (persistent discrimination)
Mexico	Hazarika and Otero (2004)	GATT/NAFTA	Gender earnings gap narrows, especially outside maquiladoras	Competition reduces taste-based discrimination
Mexico	Juhn et al. (2014); Aguayo-Tellez et al. (2014)	NAFTA tariff cuts	Tech upgrading from trade boosted women's blue-collar wages and employment, not white-collar	Tech upgrading reduces physical-task intensity
Peru	Mansour et al. (2022)	China import competition	Women experience persistent participation losses, unlike men	Limited non-tradable job options hinder recovery

Collectively, these theoretical and empirical insights underscore that the relationship between trade liberalization and the gender wage gap is shaped by underlying structural composition and institutional context. Whether increased trade exposure disciplines discriminatory practices or reinforces existing inequalities remains an open empirical question - one that this paper investigates in the case of Vietnam.

2.3 Institutional and Economic Context

2.3.1 Overview of the 2001 U.S. - Vietnam Bilateral Trade Agreement

The U.S.-Vietnam Bilateral Trade Agreement (BTA) was a comprehensive agreement covering trade in goods and services along with broader measures to facilitate business activity. It was signed on July 13, 2000 and entered into force in December 2001. Prior to the agreement, Vietnamese exports to the United States were subject to Column 2 tariff rates under the Harmonized Tariff Schedule. These rates, which originated in the U.S. Tariff Act of 1930 (Smoot-Hawley), apply to countries without Normal Trade Relations (NTR) status and are substantially higher than the tariffs levied on Most Favored Nation (MFN) partners. With the implementation of the BTA, the United States immediately extended NTR/MFN tariff treatment to Vietnam, granting its exports the same market access as offered to other MFN trading partners[5]. As a result, the applicable tariff rates on Vietnamese exports fell sigificantly from an average of 23.4% to 2.5%.

The BTA provides a unique opportunity to study the effects of a trade shock in the form of increased access to a major export market for three key reasons. First, the transition from pre-existing tariff schedules meant that neither the U.S. nor Vietnam could negotiate sector-specific tariffs. This reduces concerns about endogeneity from political lobbying. Second, Vietnam had already granted the U.S. Most Favored Nation (MFN) status prior to the BTA's enactment (McCaig, 2011), making the agreement largely one-sided: expanding export opportunities without significantly increasing import competition. Third, while Vietnam committed to broader trade

[5]This MFN tariff schedule was negotiated among World Trade Organization members in 1995 (McCaig et al., 2022)

and investment reforms under the BTA to level the playing field for U.S. firms (STAR-Vietnam, 2003), these commitments were phased in gradually over ten years, whereas the U.S. granted immediate market access to Vietnamese exporters.

Figure 2.1 illustrates that U.S. imports to Vietnam were already on a steady upward trend before the BTA, while Vietnam's exports to the U.S. responded sharply after its implementation. In the first year, exports rose by 127.4%, and more than quadrupled within two years. Figure A.1 further disaggregates Vietnam's export growth by destination (U.S. vs. rest of world). The considerable surge in exports to the United States is not merely a diversion effect - i.e., reallocating shipments away from other markets — but instead reflects a substantial expansion in overall export capacity. Overall, the BTA induced a sudden and substantial trade shock to Vietnam's economy, and the arguably exogenous tariff reductions can be leveraged as a natural experiment to study its effect on labor market outcomes.

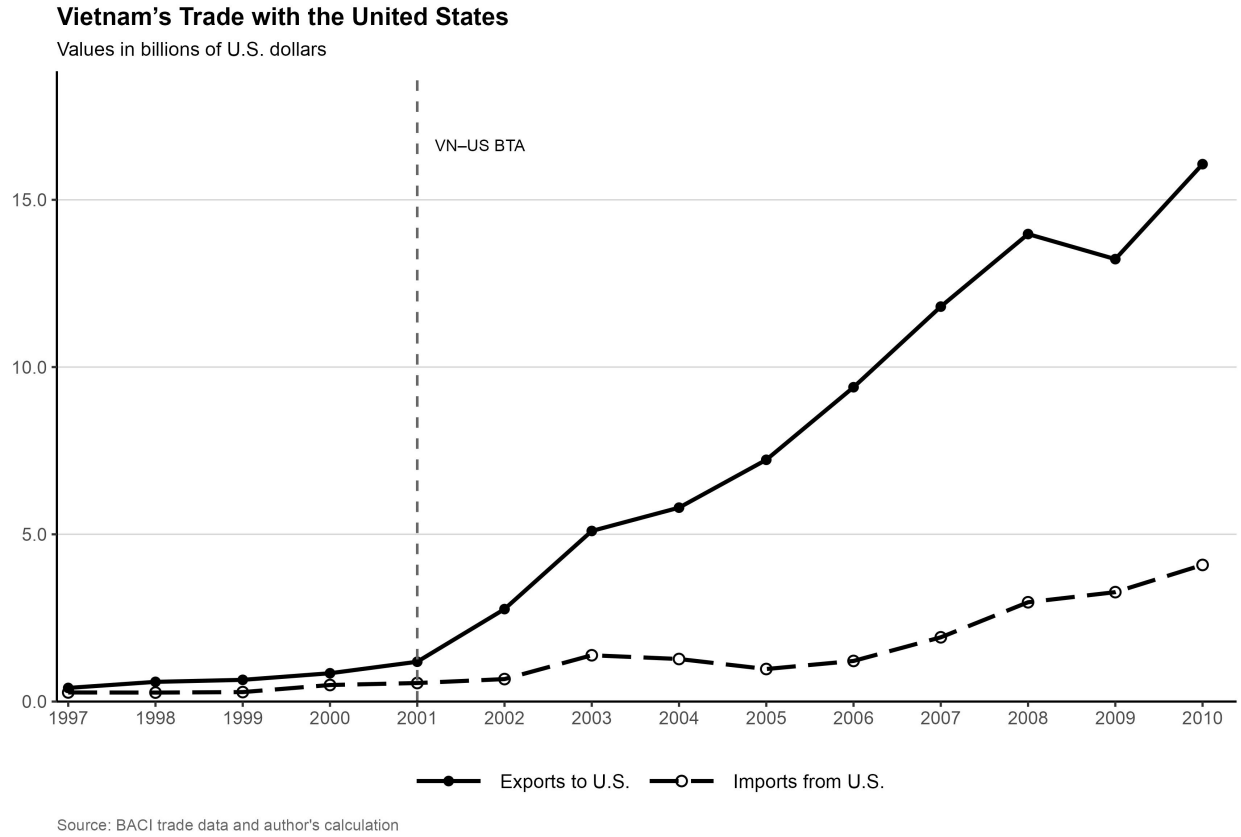


Figure 2.1. Vietnam's Trade with the United States

Beyond the aggregate expansion in exports, an important feature of the BTA is the substantial heterogeneity in the size of tariff reductions across industries. Among the 31 affected 2-digit ISIC industries, the largest reductions occurred in manufacturing, where tariffs dropped from an average of 33.8% to 3.6%. The agriculture sector also saw meaningful reductions (average cut of 7.2%), while the changes in fishing and mining were relatively modest [6]. Industries with the largest tariff reductions include wearing apparel (VSIC 18), tobacco products (VSIC 16), textiles (VSIC 17), and furniture manufacturing (VSIC 36).

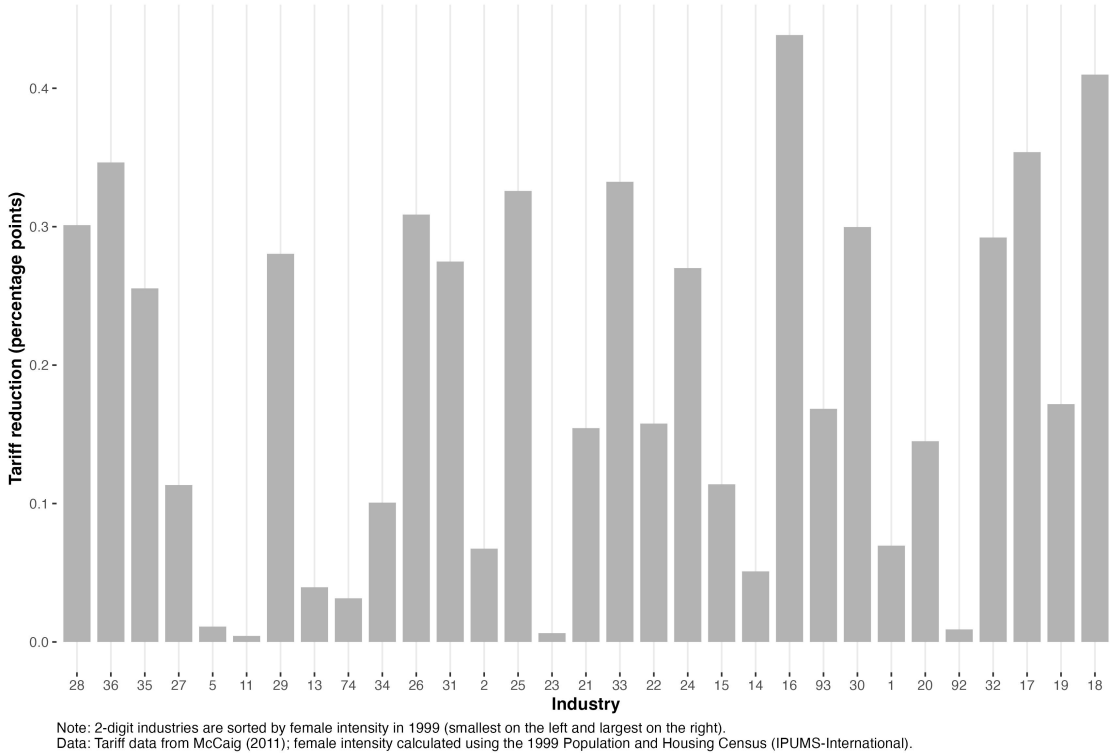
Since one channel through which trade liberalization may influence the gender wage gap is the reallocation of labor across expanding and contracting industries, it is natural to examine whether tariff reductions are correlated with gendered employment patterns. Empirically, the magnitude of tariff cuts does not align systematically with the pre-BTA gender composition of employment. For example, while ISIC 18 and 17 are highly female-intensive (with female employment shares of 81% and 72.3%, respectively), ISIC 36 is predominantly male-intensive (female share of 14.2%), and ISIC 16 has an equal gender split. As shown in Figure 2.2, which plots industry-level tariff reductions against female employment shares in 1999[7], there is no clear relationship between the two. This absence of correlation suggests that any gender-differentiated effects of trade liberalization are unlikely to be driven by pre-existing gender sorting across industries with differing trade exposure.

2.3.2 Gender Disparities in Vietnam's Labor Market

Qualitative evidence from focus group discussions and expert interviews highlights the extent to which gender norms shape labor market behavior in Vietnam (Buchhave et al., 2020). Societal expectations continue to cast women primarily as wives and mothers, with paid work viewed as secondary, while men are expected first and foremost to be breadwinners. These norms influence women's educational and occupational choices from an early age: young women often steer to-

[6] Average tariff reductions were 1.1% and 1.9% for fishing and mining, respectively

[7] Calculated using the 1999 Population and Housing Census, made available through IMPUS International



Note: 2-digit industries are sorted by female intensity in 1999 (smallest on the left and largest on the right).
Data: Tariff data from McCaig (2011); female intensity calculated using the 1999 Population and Housing Census (IPUMS-International).

Figure 2.2. Industry-Level Tariff Cuts Sorted by Female Employment Share (Pre-BTA)

ward fields and jobs perceived as compatible with caregiving responsibilities, typically public sector positions with maternity leave, limited overtime, and fewer physical or mobility demands. By contrast, men face strong pressure to pursue higher-paying, career-advancing paths, and concerns about their work–family balance are rarely raised. Such gendered expectations contribute to persistent disparities in labor market opportunities. Using data from the Vietnam Household Living Standards Surveys[8], I document four stylized facts regarding gender differences in labor market outcomes.

Fact 1: Female labor force participation in Vietnam is consistently high

As can be seen from Figure 2.3, across all five VHLSS waves (2002–2010), female employment rates remain persistently high, ranging from 78% to 82%, only slightly lower than those of mens. This pattern is consistent with Vietnam's historical and institutional context, where women's labor force attachment has been strongly supported both socially and economically

[8]Details on data construction are provided in the next section.

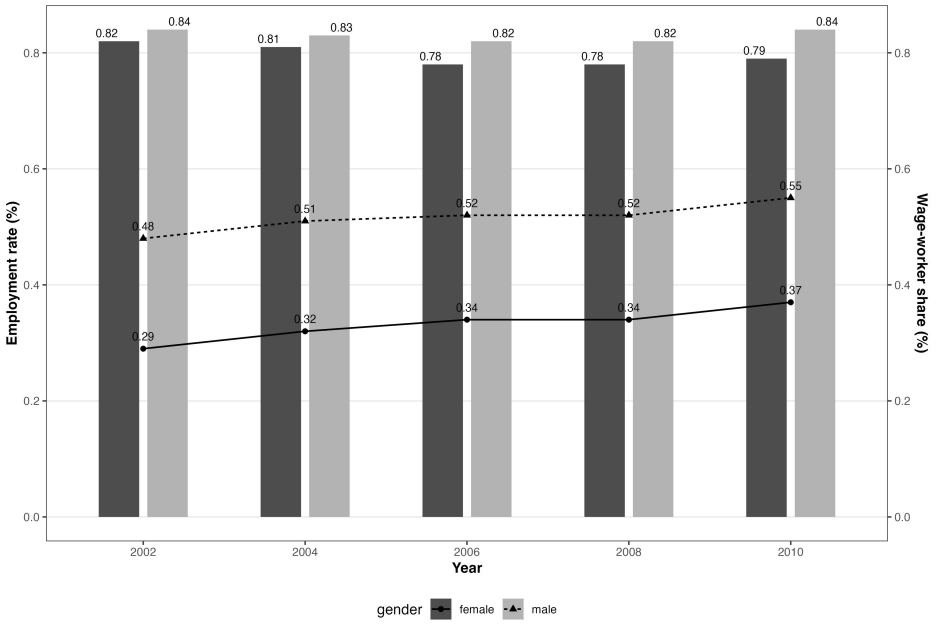


Figure 2.3. Employment Rate and Percentage of Wage Workers by Gender

Note: The percentage of employed individuals is defined as the percentage of people aged 15 to 64 who reported working in the past 12 months. Among these workers, the percentage of wage workers is calculated based on those who reported that their most time-consuming job in the past 12 months was a wage-earning position.

(Banerji et al., 2018). High female participation suggests that gender disparities in labor outcomes are less likely to stem from selection out of employment, and more likely to reflect within-labor-market mechanisms such as occupational segregation or unequal returns.

It is worth noting, however, that even though employment rates are comparable between men and women, a consistently higher percentage of employed men reported working for wages compared to women. As a result, there is a potential concern about selection into wage employment. Since my analysis of the gender wage gap focuses exclusively on wage workers, it only offers partial insights into the gender inequality that exists in the labor market.

Fact 2: Male wages are consistently higher than female wages

As shown in Table 2.2, average wages are higher for men than for women across all survey waves between 2002 and 2010, both in annual and hourly terms. In 2002, for example, men earned an average of 6.0 million VND annually compared to 5.2 million VND for women. This

Table 2.2. Average Wages and Log Wage Gaps by Gender and Year

Year	Annual Wage			Hourly Wage		
	Male	Female	Log Wage Gap	Male	Female	Log Wage Gap
2002	5,966	5,127	0.15	4.0	3.6	0.10
2004	8,075	6,875	0.16	4.8	3.8	0.23
2006	9,349	7,978	0.16	4.9	4.3	0.13
2008	10,312	8,866	0.15	5.3	4.7	0.12
2010	13,205	11,081	0.18	7.0	6.3	0.11

Unit: 1000 VND

Note: All averages are computed using individual-level survey weights. Nominal hourly wages are calculated by dividing total earnings over the past 12 months by estimated annual hours worked. Annual hours are imputed as the product of months worked, days worked per month, and hours worked per day. Both annual and hourly nominal wages are adjusted to national average prices using the regional price index provided in the dataset. To account for inflation, I also apply the Consumer Price Index for 2004–2010, using 2002 as the base year, as reported by the General Statistics Office (GSO). Finally, Observations in the top and bottom 1% of the real wage distribution are removed by gender and year to reduce the influence of outliers.

pattern persists over time; by 2010, men earned 13.2 million VND a year on average, while women earned 11 million VND - a gap of 2 million VND annually.

The wage gap in log terms, a commonly used measure of relative wage differences, ranges from 0.15 to 0.18 for annual wages and from 0.10 to 0.23 for hourly wages over the period. The log hourly wage gap exhibits greater fluctuation, peaking in 2004. This greater variability may partly reflect measurement noise in reported total hours worked, which are used to compute hourly wages. In contrast, the annual wage gap remains relatively stable from 2002 to 2008 before increasing in 2010[9].

Figure A.2 further illustrates these differences: across all five survey years, the distribution of log real wages for men is consistently shifted to the right of the distribution for women, indicating a greater share of male workers earning higher wages. While there is substantial overlap in the middle of the distribution, the divergence at the upper tail, especially in later years, suggests that men are overrepresented among top earners. In contrast, the female distribution appears

[9]This pattern is consistent with evidence from the International Labour Organization's 2012/2013 Global Wage Report, which documents that although gender wage gaps declined in most countries between 2008–2011 relative to 1999–2007, Vietnam is among the few countries where the gap continued to widen.

slightly more concentrated in the lower-to-middle portion of the wage distribution in all survey years. This suggests that women are more likely to be clustered in lower-paying jobs, contributing to the persistent average wage gaps.

Fact 3: Male and female workers are broadly similar in observed characteristics

As can be seen from Table A.1, despite statistically significant differences in several characteristics, male and female workers are broadly comparable across observed traits. Men tend to be slightly older and are consistently more likely to live in urban areas. However, women are more likely to have attained vocational or college education in most years with the exceptions of vocational training in 2002 and 2010, where male rates are marginally higher. While gender differences in means are often statistically significant due to the large sample size, the magnitudes are generally modest, suggesting limited disparity in key demographic and educational variables.

2.4 Data

I use three major data sources for my analysis: labor market outcomes derived from five waves of the Vietnam Household Living Standard Surveys (VHLSSs) from 2002 to 2010, ad valorem equivalent tariff data from McCaig (2011), and provincial employment data derived from the 1999 Population and Housing Census. Detailed descriptions of each data source are provided below.

Household Surveys: The Vietnam Household Living Standards Surveys (VHLSSs) are conducted every two years by the General Statistics Office (GSO) of Vietnam to provide information on household expenditures, employment and various other household and individual characteristics. The VHLSS waves conducted during 2002–2008 used the sampling framework from the 1999 Population and Housing Census, while waves from 2010 onward adopted the framework based on the 2009 Census. In both periods, the sampling designs were representative not only at the national level, but also at the regional and provincial levels (Pacillo et al., 2020).

I use five waves of the VHLSS, spanning from 2002 to 2010. Given that the recall period for expenditure and employment modules covers the twelve months preceding each survey, the 2002 wave reflects individual outcomes in the pre-BTA period, while the waves from 2004 onward capture outcomes in the post-BTA period[10]. From each wave of the VHLSS, I extract data on employment (wages, hours worked, occupation and industry affiliation), educational and demographic characteristics (gender, age, highest level of completed education, ethnicity and marital status), geographic location (province of residence, urban status) and sampling weights. During the study period, several changes in provincial boundaries took place, which required adjustments to maintain consistency across survey years. Specifically, for the 2004 wave onward, I merge Dien Bien, Dak Nong, and Hau Giang back into Lai Chau, Dak Lak, and Can Tho, respectively. For the 2002–2008 waves, I also merge Ha Tay into Ha Noi. These adjustments result in a boundary-consistent set of 60 provinces.

Industrial classifications also vary across waves and require harmonization. The 2002–2006 VHLSS waves use the Vietnamese Standard Industrial Classification (VSIC) 1993, whereas waves from 2008 onward adopt VSIC 2007.[11] I harmonize industry codes across waves using the VSIC 1993–VSIC 2007 concordances in McCaig et al. (2022), with one modification: I map VSIC 2007 code 09 ("Service activities incidental to ores extraction") to VSIC 1993 code 13 ("Mining of metal ores"), rather than to code 10 ("Mining of coal and lignite; extraction of peat") as in McCaig et al. (2022).[12] Additionally, under McCaig et al. (2022)'s concordance, four VSIC 1993 codes - 12, 30, 33, and 37 - do not have matches in the VHLSS 2008 and 2010 waves.[13] To maintain a consistent industry space across waves, I omit observations with these codes from the

[10]The BTA came into effect on December 10, 2001, while the 2002 VHLSS surveyed households throughout the calendar year. As a result, a small fraction of respondents interviewed late in 2002 may have recall periods that partially overlap with the post-BTA implementation phase. Consequently, my estimates may slightly understate the full impact of the BTA.

[11]VSIC 1993 was developed by the General Statistics Office (GSO) based on ISIC Rev. 3. Following the shift to ISIC Rev. 4, VSIC 2007 substantially overhauled the 1993 system—reorganizing and renumbering categories and introducing a fifth classification level.

[12]This reassignment better aligns the underlying activity definitions across revisions.

[13]The four VSIC 1993 codes are: 12 ("Mining of uranium and thorium ores"), 30 ("Manufacture of office, accounting and computing machinery"), 33 ("Manufacture of medical, precision and optical instruments, watches and clocks"), and 37 ("Recycling").

2002–2006 datasets.[14]

Employment type codes also differ across VHLSS waves and were recoded to ensure consistency over time. Each wave categorizes employment status using distinct numeric codes that are not directly comparable. I therefore created a harmonized variable which collapses the original codes into four broad, mutually exclusive categories: (1) unofficial sector, (2) state-owned enterprises (SOEs), (3) private domestic firms and (4) foreign-owned firms. This harmonization produces a consistent measure of employer type that enables meaningful comparisons of employment patterns across survey waves.

Tariff Data: Data of tariff schedules before and after the BTA was taken from (McCaig, 2011). The author extracted column 2 as well as MFN U.S. tariffs from the U.S. International Trade Commission's online Tariff Information Center and computed ad valorem equivalents by estimating the unit value of imports within each 8-digit HTW tariff line using total annual imports from all countries. They then matched the tariff lines to industries by the concordance provided by the World Bank via the World Integrated Trade Solution database to construct industry-level tariffs. I obtained tariff data at the 2-digit industry level based on International Standard Industrial Classification revision 3.

Employment Data: To construct the employment share by industry within each province, I use the 3 percent sample of the 1999 Population and Housing Census made available through Integrated Public Use Microdata Series – International's website. In the 1999 Census, industry affiliation for employed individuals is reported at the 3-digit ISIC level. Since McCaig (2011) pointed out minor inconsistency between the Vietnamese census codes and the corresponding 3-digit ISIC codes, I follow the same procedure in making changes as documented in Appendix B.1.[15].

[14]This results in the loss of 48 observations in 2002, 47 in 2004, and 69 in 2006.

[15]McCaig (2011) made changes to 20 lines of industry codes. After double checking with the IPUMS variable description as well as the ISIC Revision 3 document, I made changes similar to McCaig (2011) except for industry codes 712, 731, 734 and 911. Industries 712, 731 and 734 are not listed in the IPUMS variable description. Industry 911 is labeled as "party activities" in the IPUMS variable description and will be kept to match with the corresponding ISIC code labeled as "activities of business, employers and professional organizations".

Moreover, there are three provinces for which the 1999 Census codes differ from the codes used in VHLSSs as documented by the ILO: Vinh Phuc, Bac Ninh and Lam Dong. I revise these codes to ensure consistency across data sources by changing province codes 219, 223, and 703 to 104, 106, and 607, respectively.

Besides the employment structure at the provincial level, I also utilize the 1999 Census data to construct the female intensity at the industry level. Female intensity is defined as female's employment share within each industry.

2.5 Empirical Methodology

2.5.1 Wage Regression with Trade Shock at the Industry Level

As reported in section 4.2, the BTA generated non-zero tariff reductions across 31 two-digit ISIC industries, with significant variation across industries. The tariff cuts likely induced expansion in some industries and contraction in others, potentially affecting men and women differently given gender segregation across sectors and occupations. As established in section 3.1, the reductions in industry tariffs arise from movements between pre-existing tariff schedules and can therefore be thought of as plausibly exogenous shocks.

To examine the causal effect of trade liberalization on the gender wage gap, I employ a triple difference-in-differences framework using tariff reductions at the industry level as trade shocks and estimate the following specification:

$$\ln w_{jit} = \alpha + \beta_1 \mathbb{I}(t > 2002) \cdot TariffCut_i + \beta_2 \mathbb{I}(t > 2002) \cdot TariffCut_i \cdot Female_{jt} + \delta X_{jt} + \theta_t + \varepsilon_{jit} \tag{2.1}$$

where $\ln w_{jit}$ denotes the natural logarithm of annual wages of individual j working in industry i in year t and X_{jt} is a vector of control variables of workers characteristics including gender, age and its square, educational attainment, marital status, dummy for urban/rural and ethnicity. $TariffCut_i$ measures the proportional reduction in tariffs for industry i following the BTA, de-

fined as the log difference between Column 2 and MFN rates, i.e., $\ln(1+\tau_i^{Col2}) - \ln(1+\tau_i^{MFN})$. $\mathbb{I}(t > 2002)$ is a post-BTA indicator which takes value 1 for the year 2004 onwards and 0 otherwise, and $Female_{jt}$ is an indicator equal 1 if individual j is female. θ_t denotes fixed effects for each year t that capture the effect of aggregate trends on all workers. The error term ε_{jit} is assumed to be mean zero, and standard errors are clustered at the industry level to allow for correlation within any given industry. The analysis uses repeated cross-sections from the five waves of the VHLSS, and the corresponding results are reported in Table 2.3.

Table 2.3. Estimated Effects of Industry-Level Trade Shocks on Individual Wages

Dependent variable: Log of real wages	(1)	(2)
BTA × tariff_reduc	0.0022 (0.0012)	0.0012 (0.0015)
BTA × tariff_reduc × female	-0.0005 (0.0016)	-0.0002 (0.0017)
Year FE	Yes	Yes
Province FE	No	Yes
Observations	158,354	158,354
R^2	0.396	0.440
Within R^2	0.314	0.287
Standard errors	Clustered by industry	

Notes: Observations in the top and bottom 1% of the real wage distribution are removed by gender and year to reduce the influence of outliers. Robust standard errors clustered by industry are reported in parentheses.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

As shown in Table 2, the coefficients on *TariffCut* × *Post_BTA* × *Female* are statistically indistinguishable from zero, indicating that BTA-induced tariff reductions at the industry level did not produce significant differential effects on wages between men and women. However, using industry-level tariff cuts is only one way to measure trade liberalization, and for this to be a salient channel, it requires the strong assumption of perfect labor mobility across regions in response to industry-specific demand shocks. As noted by Dix-Carneiro and Kovak (2023), while earlier empirical work focused on heterogeneity in globalization's effects across industries or worker education levels, more recent research has shifted toward examining differential impacts

across geographic regions, following the framework pioneered by Topalova (2007). This regional approach has also been adopted in seminal studies such as (Autor et al., 2016; Kovak, 2013; McCaig, 2011). I next turn to this approach of measuring trade exposure in my next analyses.

2.5.2 Measuring Trade Exposure at the Provincial Level

Following Topalova (2007), I exploit variation across regions in exposure to tariff cuts that stems from differences in regions' pre-liberalization industry composition of employment. This framework allows me to examine whether provinces that were more exposed to international trade experienced smaller or bigger changes in gender wage inequality than less-exposed provinces. To the extent that exposure to international trade differs across provinces and are exogenous, the differential exposure allows for separating the effects of trade from the effects of concurrent policy changes (Goldberg and Pavcnik, 2007). I construct provincial measures of trade exposure as follows:

$$TradeExp_p = \sum_i \omega_{i,p}^{1999} \Delta\tau_i$$

where p indexes provinces, $\omega_{i,p}^{1999}$ is the share of workers working in industry i in province p in 1999 and $\Delta\tau_i$ is the proportional reduction in tariff of industry i due to the BTA, defined as $ln(1+\tau_i^{Col2}) - ln(1+\tau_i^{MFN})$, with larger values implying more trade liberalization. I also followed the approach of Kovak (2013) by omitting the non-traded sector prior to calculating local employment shares, rather than assuming that non-traded sectors were unaffected by trade liberalization.

Two key assumptions are required to establish the exogeneity of this measure of trade exposure: (1) tariff cuts were not driven by any particular provinces, and (2) labor was not perfectly mobile across provinces. Several factors support the plausibility of the first assumption. First, as discussed in section 3.1, the variation in tariff reductions is plausibly exogenous because it stems from shifts between two sets of pre-existing tariff schedules. Under the BTA, tariff rates applied to Vietnam's exports to the United States moved from Column 2 rates — established un-

der the Tariff Act of 1930 (Pregelj, 2001; McCaig et al., 2022) and reamained largely unchanged before and after the BTA[16] — to the MFN tariff schedule negotiated among WTO members in 1995. Second, Vietnam had 61 provinces at the time of the BTA signing[17], and this relatively large number of independent units further supports the plausibility of the exogeneity assumption.

The second assumption is also plausible in the short and medium run. Vietnam's household registration system, known as *ho khau*, was instituted in the 1950s as a central planning tool used by the government for identification, eligibility for state employment, and law enforcement, as well as to restrict migration into urban areas (Hardy, 2001). Under this system, individuals are tied to their *ho khau* registration location and face limited access to employment opportunities, social services (such as education and healthcare), land rights, and other entitlements outside that location. Although the system has become less rigid over time, concerns persist that *ho khau* continues to constrain internal migration and limit access to public services for those without permanent registration in their place of residence (Demombynes and Vu, 2022). Consistent with this institutional constraint, McCaig (2011) report, using data from the 1999 Population and Housing Census, that only about 3 percent of individuals moved across provinces between 1994 and 1999, providing evidence of limited interprovincial labor mobility prior to the BTA. Finally, as noted by Goldberg and Pavcnik (2007), even if some degree of mobility exists, violation of this assumption would not invalidate the empirical framework but would instead attenuate the estimated differential effects of trade policy as wages equalize across regions.

An important caveat that needs to be addressed with this approach is that it only allows me to capture the provincial deviations from the national trend and does not capture the average impact of trade exposure on the outcome of interest. The total impact, which is the sum of these two components, cannot be determined by this specification. However, to the extent that the identification assumptions are plausible, a causal relationship between trade exposure and the gender wage gap can be inferred.

[16]McCaig et al. (2022) report a correlation of 0.978 between Column 2 rates in 1997 and 2005.
[17]This includes 56 provinces and 5 centrally controlled municipalities.

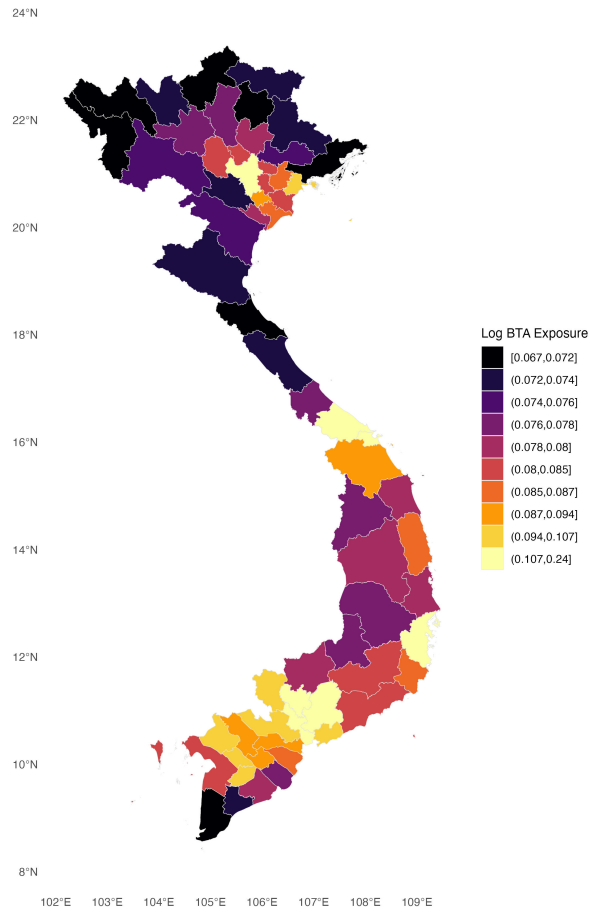


Figure 2.4. Provincial BTA Tariff Exposure

As can be seen from Figure 2.4, there is significant variation in trade exposure at the provincial level. The measure of trade exposure ranges from 0.067 (corresponding to Ca Mau province) to 0.239 (corresponding to Ho Chi Minh City) with a standard deviation of 0.039. To address potential concerns that this measure of trade exposure at the provincial level merely reflects provinces' other characteristics such as population or GDP, I conduct some correlation analyses. The correlation between trade exposure and provincial population in 1999[18] is 0.54, indicating a moderate positive relationship. Trade exposure is more strongly correlated with total industrial gross output in 1997[19], with a correlation coefficient of 0.80. However, as illustrated in Figure A.4, this strong correlation is largely driven by Ho Chi Minh City, which is an outlier in terms of both industrial

[18]Data of provincials' population is from the 1999 Population & Housing Census conducted by the GSO, made available through https://www.citypopulation.de/en/vietnam/prov/admin/

[19]Data of provincials' total industrial gross output in 1997 is from the 1997 Statistical Yearbook of Vietnam, made available through the GSO's website

output and exposure. After excluding Ho Chi Minh City from the sample, the correlation drops to 0.47, suggesting a more moderate positive relationship.

2.5.3 Provincial-level Analysis

Since trade exposure is measured at the provincial level, I begin by focusing on specifications that can be implemented at the same level of aggregation. This allows me to directly examine how provinces that were more exposed to tariff reductions experienced changes in gender-specific wage premiums relative to less exposed provinces.

I assess the effect of trade exposure on female wage premiums using a two-step estimation framework commonly employed in the labor and urban economics literature to study industry- or city-level wage premiums[20]. In the first stage, I regress the logarithm of real annual wages for individual i living in province p in year t on a vector of individual characteristics (X_{it}) and a full set of province–gender dummy variables:

$$\ln w_{it} = \alpha X_{it} + \sum_{p=1}^{P} \sum_{g \in \{0,1\}} \mathbf{1}\{\text{Province}_i = p\}\, \mathbf{1}\{\text{Gender}_i = g\}\, \theta_{pg,t} + \varepsilon_{it}, \tag{2.2}$$

The vector X_{it} includes age and its square, educational attainment, marital status, an urban-rural indicator, ethnicity, and the logarithm of total time worked. The coefficients on the province–gender dummies ($\theta_{pg,t}$) capture the component of wage variation that is not explained by individual characteristics but is systematically associated with the worker's gender and province of residence - that is, the provincial wage premium by gender. Following Krueger and Summers (1988), I normalize the sum of the employment-weighted provincial wage premiums to zero within each gender group and express all estimated premiums as deviations from their respective gender-specific weighted means. The first-stage regressions are estimated separately for each year, as the subscript t in equation (2) indicates.

[20] See, among others, Card and Lemieux (2001), who examine the rise in college wage premiums in the U.S.; Roca and Puga (2017), who study city-level wage premiums in Spain; and **?**, who employ a similar two-step approach to examine how industry wage premiums responded to trade liberalization in Colombia

In the second stage, I separately analyze the evolution of provincial wage premiums for female and male workers. To assess the effect of trade liberalization on the provincial wage premiums for female workers, I estimate the following difference-in-differences specification with a continuous treatment intensity:

$$\hat{\theta}_{pf,t} = \beta \mathbb{I}(t > 2002) \times TradeExp_p + \mu_p + \tau_t + u_{pt}, \tag{2.3}$$

where $\hat{\theta}_{pf,t}$ denotes the estimated provincial wage premium for female workers in province p and year t obtained from the first stage, $TradeExp_p$ is the time-invariant measure of provincial trade exposure constructed in section 5.2, and $\mathbb{I}(t > 2002)$ is an indicator that equals one for post-BTA years (2004 onward) and zero otherwise. Province fixed effects (μ_p) absorb time-invariant provincial characteristics, while year fixed effects (τ_t) control for aggregate shocks common to all provinces. Standard errors are clustered at the province level. Following standard two-step practice, I estimate equation (3) with weighted least squares, with weights equal to the inverse of the estimated variance of $\hat{\theta}_{pf,t}$ from the first stage.

An implicit assumption underlying the first-stage estimation is that gender-specific wage setting varies systematically across province (i.e. $\theta_{pg,t}$ exists and exhibits meaningful variation). Differences in local labor demand conditions, occupational segregation, discrimination, or industrial composition across provinces may all contribute to such systematic variation in wage setting. To assess this, I conduct standard F-tests comparing models with and without province–gender dummies separately for each survey year. The resulting p-values are small at conventional significance levels, indicating that province–gender effects are jointly significant in every year. Moreover, to examine the stability of these effects over time, I compute the correlation between $\hat{\theta}_{pg,t}$ and $\hat{\theta}_{pg,t-1}$. Specifically, I calculate one overall correlation between current and lagged provincial wage premiums for all valid adjacent pairs of years in the panel. The estimated correlation coefficient of 0.7 indicates that provincial gender-specific wage premiums are highly persistent, suggesting that they capture systematic and enduring patterns in local wage-setting rather than

sampling noise.

While the provincial-level analysis aligns the treatment variable (i.e. trade exposure) with the level of aggregation at which it is measured, a limitation of the two-step approach is that, after estimating provincial wage premiums by gender in the first stage, I must subset the data by gender in the second stage. As a result, I cannot directly compare the estimated coefficients across the male and female regressions to infer gender differences in the wage response to trade exposure. Motivated by this limitation, I next turn to an individual-level framework that allows me to estimate these gender-differentiated effects directly within a unified specification.

2.5.4 Individual-level Analysis

In this section, I examine the relationship between trade exposure and individual wages using a triple difference-in-differences framework. This specification exploits variation across provinces, time, and gender to directly test whether trade liberalization had differential effects on men's and women's wages. I estimate the following equation:

$$\begin{aligned} \ln w_{ipt} = \alpha &+ \beta_1 \mathbb{I}(t > 2002) \times TradeExp_p \\ &+ \beta_2 \mathbb{I}(t > 2002) \times TradeExp_p \times Female_{it} \\ &+ \delta X_{it} + \theta_t + \xi_p + \varepsilon_{ipt}. \end{aligned} \tag{2.4}$$

where $\ln w_{ipt}$ denotes the logarithm of annual wages of individual i residing in province p in year t, and X_{it} is a vector of worker characteristics including gender, age and its square, educational attainment, marital status, and indicators for urban residence and ethnicity. $TradeExp_p$ and $\mathbb{I}(t > 2002)$ are defined as in equation (3), and $Female_{it}$ equals one if individual i is female. θ_t represents year fixed effects capturing aggregate shocks that affect all workers, while province fixed effects ξ_p control for time-invariant underlying differences across provinces. The error term ε_{ipt} is assumed to have mean zero, and standard errors are clustered at the province level to allow for correlation within provinces over time.[21]

[21] Following Solon et al. (2015), I do not apply sampling weights in the main specification because the objective is to estimate causal effects rather than descriptive population averages. Given that the VHLSS sampling design is

In this specification, β_2 measures the differential effect of trade exposure on female workers relative to male workers in the same province and time period. A negative coefficient indicates that female workers benefited less from the BTA than their male counterparts. One important caveat is that equation (4) is estimated using repeated cross-sectional data from successive waves of the VHLSS rather than panel data. Without individuals fixed effects, unobserved worker characteristics correlated with both gender and trade exposure may not be fully accounted for. Nonetheless, because the specification compares workers across provinces with different exposure to trade both before and after the BTA, unobservable but time-invariant provincial factors that affect wages are differenced out. The inclusion of time and province fixed effects further controls for aggregate economic trends and persistent provincial differences. I also estimate a specification that includes province-by-year fixed effects, which flexibly absorb all time-varying shocks common to individuals within the same province-year cell. In that case, the interaction term $\mathbb{I}(t > 2002) \cdot TradeExp_p$ is absorbed by these fixed effects. However, the main coefficient of interest, β_2 in equation (3), remains identified and still captures the differential impact of trade exposure on female relative to male workers.

2.6 Empirical Results

2.6.1 Provincial-level Analysis

Regression results of equation (3) are presented in Table 2.4. In both specifications, the interaction between provincial trade exposure and the post-BTA period is negative and statistically significant at the 1% level, indicating that provinces more exposed to tariff reductions experienced a relative decline in female wage premiums after the BTA. The estimated coefficient decreases in magnitude from –2.04 to –1.78 when industry fixed effects are added in the first stage of the two-step approach, as shown in column (2), but remains robustly significant. Between the two, specification (1) is preferred, as the inclusion of industry fixed effects may absorb part of the

stratified but exogenous conditional on controls and fixed effects, unweighted estimation yields consistent and more efficient estimates. Weighted regressions are nevertheless reported as a robustness check to verify that results are not sensitive to the sampling design

Table 2.4. Estimated Effects of Provincial Trade Shocks on Provincial Wage Premiums for Female Workers

Dependent variable:	*female provincial wage premiums*	
	(1)	(2)
Trade_exposure × post_BTA	-2.039*** (0.4098)	-1.782*** (0.4337)
Year Fixed Effects (2nd stage)	Yes	Yes
Province Fixed Effects (2nd stage)	Yes	Yes
Industry Fixed Effects (1st stage)	No	Yes
Standard Errors	Clustered by province	
Observations	300	300
R^2	0.804	0.781
Within R^2	0.069	0.026

Notes: Observations in the top and bottom 1% of the real wage distribution are removed by gender and year before being used in the first-stage wage regressions. Robust standard errors clustered by province are reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

mechanism through which trade liberalization influences wages, particularly if trade reform itself reshaped provincial industrial structures. Quantitatively, the coefficient of –2.04 implies that a one–standard-deviation increase in provincial trade exposure (0.039) is associated with a roughly 8% decline in the female provincial wage premium. As noted in Section 5.2, this specification captures deviations from national trends rather than the aggregate impact of trade liberalization. A negative coefficient means that provinces with higher exposure to tariff cuts saw either smaller increases or larger decreases in female wage premiums relative to the national average. One possible explanation for this pattern is mean reversion, whereby provinces with initially higher female wage premiums experienced slower subsequent growth than those starting from lower levels. Alternatively, larger declines in female wage premiums in highly exposed provinces may reflect adverse adjustment pressures, as trade liberalization disproportionately affected female-intensive sectors or limited women's mobility into expanding industries.

Turning to the results for male provincial wage premiums, reported in Table 2.5, a similar pattern emerges. Across both specifications, the interaction between provincial trade exposure

Table 2.5. Estimated Effects of Provincial Trade Shocks on Provincial Wage Premiums for Male Workers

Dependent variable:	*male provincial wage premiums*	
	(1)	(2)
Trade_exposure × post_BTA	-1.704*** (0.3525)	-1.554*** (0.4413)
Year Fixed Effects (2nd stage)	Yes	Yes
Province Fixed Effects (2nd stage)	Yes	Yes
Industry Fixed Effects (1st stage)	No	Yes
Standard Errors	Clustered by province	
Observations	300	300
R^2	0.809	0.750
Within R^2	0.055	0.032

Notes: Observations in the top and bottom 1% of the real wage distribution are removed by gender and year before being used in the first-stage wage regressions. Robust standard errors clustered by province are reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

and the post-BTA period is negative and statistically significant at the 1% level. The estimated coefficients range from –1.55 to –1.7 between specifications with and without industry fixed effects in the first stage. This indicates that provinces more exposed to tariff reductions experienced relative declines in male wage premiums as well, suggesting that trade liberalization exerted downward pressure on local wage premiums for both genders. However, because provincial wage premiums are normalized separately within each gender group, the coefficients from the male and female regressions are not directly comparable in magnitude. Each captures the relationship between trade exposure and within-gender deviations in provincial wages. Taken together, the results point to a general compression of provincial wage premiums following the BTA, suggesting partial convergence in regional wage differentials.

2.6.2 Individual Analysis

Table 2.6 presents individual-level estimates of the effect of trade exposure on log real wages, focusing on gender differences. In column (1), the coefficient on *Trade_exposure* · *post_BTA* is large, negative, and statistically significant, indicating that trade liberalization is associated with

Table 2.6. Estimated Effects of Trade Exposure on Individual Log Real Wages

Dependent variable:	*log real wage*	
	(1)	(2)
Trade_exposure × post_BTA	-2.466^{***} (0.4147)	
Trade_exposure × post_BTA × Female	-0.2142^{*} (0.0947)	-0.2376^{*} (0.0981)
Province Fixed Effects	Yes	No
Year Fixed Effects	Yes	No
Province × Year Fixed Effects	No	Yes
Standard Errors	Clustered by province	
Observations	158,302	158,302
R^2	0.612	0.622
Within R^2	0.506	0.510

Notes: Observations in the top and bottom 1% of the real wage distribution are removed by gender and year to reduce the influence of outliers. All specifications control for gender, age and its square, education, college and vocational dummies, marital status, urban residence, ethnicity, and log of total time worked. Robust standard errors clustered by province are reported in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

lower wages overall in provinces more exposed to tariff cuts. This is consistent with results from section 6.1 which found that provinces more exposed to trade liberalization experienced relative declines in wage premiums for both gender. Quantitatively, a one-standard-deviation increase in exposure is associated with a 9.6% decline in wages in column (1).

The triple interaction term *Trade_exposure* · *post_BTA* · *female* is negative and statistically significant at the 5% level in both specifications, suggesting that women experience a larger wage penalty in more trade-exposed provinces. In column (1), a one-standard-deviation increase in trade exposure is associated with an additional 0.83% decline in women's wages relative to men. In column (2), where province-by-year fixed effects are included to flexibly account for time-varying provincial shocks, the interaction term remains negative and robustly significant, with a slightly larger magnitude of −0.2376. These findings indicate that the gendered effects of trade exposure persist even after controlling for province-specific trends over time. While the specification with province-by-year fixed effects provides a stricter comparison by exploiting only within-province-year variation, it also absorbs the overall trade exposure effect captured by the two-way

interaction term. As such, column (2) should be interpreted as capturing only the within-province estimate of the gender differential, while column (1) captures the broader, province-level impact of trade liberalization on gender wage inequality.

Taken together, the provincial- and individual-level analyses reveal a consistent pattern: trade liberalization following the BTA was associated with a general compression of wage differentials across provinces and a widening of gender disparities within them. At the provincial level, both male and female wage premiums declined in more trade-exposed regions, suggesting downward pressure on local wage premiums and partial convergence in regional wages. The individual-level evidence further shows that women faced an additional wage penalty in provinces more exposed to tariff cuts, even after accounting for time-varying provincial shocks.

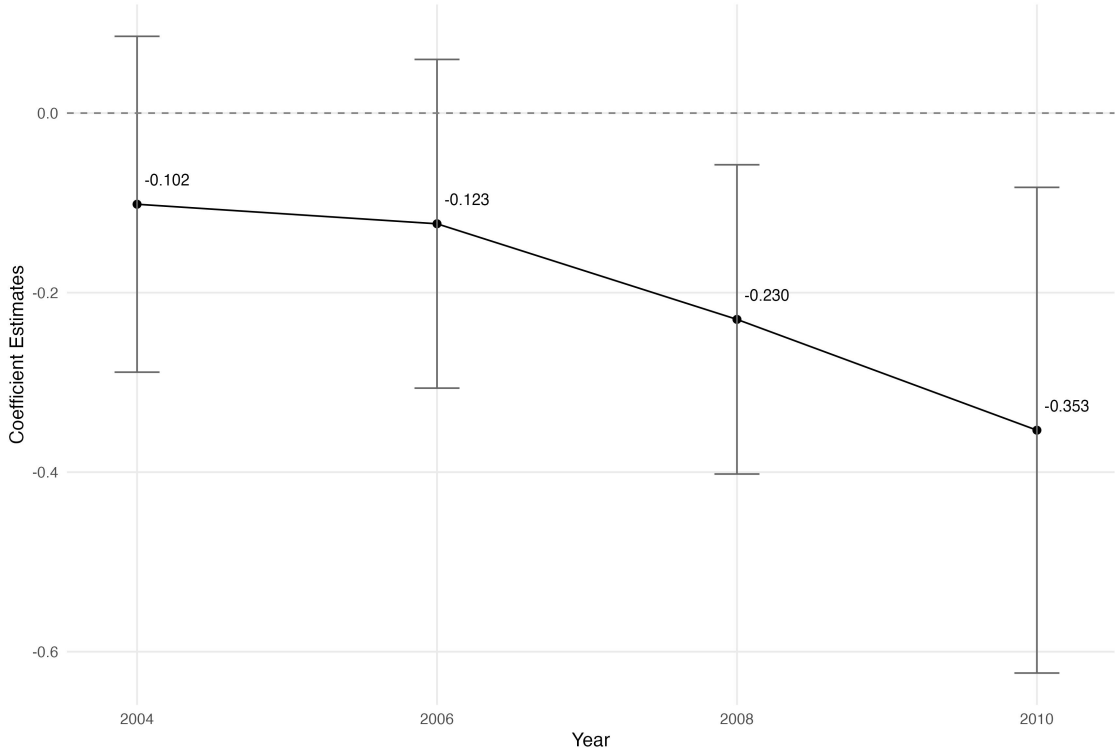


Figure 2.5. Estimated Effects of Trade Exposure on Women's Wages (Relative to Men)

Event Study: To further explore the gendered effects of trade liberalization over time, I estimate separate interaction terms between trade exposure, post-BTA years, and gender, and plot the resulting coefficients in an event study format. Figure 2.5 shows how the additional impact of trade

exposure on women's wages (relative to men's) evolves across the 2004–2010 period. While the gender-specific effects are relatively small and imprecise in the early years, they become more negative and statistically significant in 2008 and 2010. This pattern suggests that the gender gap in wages widened more substantially in provinces more exposed to trade over time, particularly beyond the immediate short term.

2.7 Robustness of Results

In this section, I evaluate the robustness of the main findings to a variety of alternative estimation choices and data treatments. Specifically, I conduct five exercises designed to ensure that the estimated effects are not driven by sampling variation, measurement choices, or model specification. First, I re-estimate the main specification using the VHLSS sampling weights to verify that the results are not sensitive to the survey's sampling desig. Second, I perform a jackknife robustness exercise that sequentially omits one province at a time to confirm that no single province drives the main results. Third, I include industry-by-year fixed effects to control for sector-specific time-varying shocks that could differentially affect male and female workers. Fourth, I replace the continuous measure of trade exposure with a binary treatment indicator to test sensitivity to functionalvform assumptions. Finally, I use hourly earnings instead of annual wages as the dependent variable to verify that the results are robust to alternative definitions of the wage variable. Together, these exercises demonstrate that the estimated gender-differentiated effects of trade liberalization remain stable across alternative sampling schemes, data constructions, and empirical specifications.

Accounting for Survey Weights: The Vietnam Household Living Standards Survey (VHLSS) employs a two-stage stratified sampling design in which enumeration areas (EAs) are first stratified by province and by urban or rural status and then selected with probability proportional to size. Within each selected EA, a fixed number of households are randomly drawn for interview. As a result, sampling probabilities differ across provinces and between urban and rural areas. To ensure that the results are not influenced by these differences in selection probabilities, I re-

estimate Equation (4) using the sampling weights provided in the VHLSS. Following standard practice, I normalize the survey weights within each year so that their average equals one, preserving the relative weight of each observation within the year while maintaining comparability across survey waves.[22]

Table A.2 compares the baseline estimates with those obtained when incorporating the VHLSS sampling weights. The results show that weighting the regressions leaves the main findings largely unchanged. The coefficient on *Trade_exposure* × *Post_BTA* × *Female* remains negative and statistically significant across all specifications, with magnitudes ranging from −0.17 to −0.18. Although slightly smaller than the unweighted baseline estimate of −0.21, these coefficients convey a consistent pattern: women in more-exposed provinces experienced smaller wage gains relative to men after the BTA. Similarly, the coefficient on *Trade_exposure* × *Post_BTA* remains large and negative. The stability of the coefficients across weighted and unweighted specifications confirms that the main results are not artifacts of the sampling design or of unequal selection probabilities across provinces.

Jackknife Robustness: To ensure that the results are not driven by a small number of influential provinces, I perform a jackknife robustness exercise in which the main specification is repeatedly estimated while omitting one province at a time. This leave-one-province-out procedure assesses the influence of individual provinces on the estimated gender-differentiated effects of trade liberalization and helps verify that the results are not dominated by outliers with unusually high or low trade exposure.

For each iteration, I record the estimated coefficient on *Trade_exposure* × *Post_BTA* × *Female*, and the distribution of these estimates is summarized in Figure A.5. The coefficients cluster around the baseline estimate, ranging from approximately −0.19 to −0.28, indicating strong overall stability. The only notable deviation arises when *Ho Chi Minh City* is excluded, which yields a more

[22]Formally, for each survey year t, the normalized weight is computed as $w_{it} = w_{it}^{\text{raw}} / \bar{w}_t^{\text{raw}}$, where $\bar{w}_t^{\text{raw}}$ denotes the mean of the raw weights in that year.

negative but substantially less precise estimate, making the coefficient marginally significant. This pattern reflects the city's large share of observations and its role as a key source of identifying variation, rather than a reversal of the underlying relationship.

Industry-year Fixed Effects: To ensure that the estimated effects are not confounded by nationwide shocks that affect certain industries differently over time, I include industry-by-year fixed effects in the specification. These fixed effects absorb all variation that is common to workers employed in the same industry nationwide in a given year. In doing so, they control for sector-specific time-varying shocks such as global price movements affecting a particular industry, sector-wide technological change, or nationwide policy reforms that disproportionately impact certain sectors. With both province and industry-by-year fixed effects included, identification comes from differences in trade exposure across provinces within the same industry-year, comparing workers in the same industry and year but located in provinces with different levels of trade exposure.

Results from Table A.3 show that including industry-by-year fixed effects slightly reduces the magnitude of the estimated coefficients but leaves their sign and statistical significance unchanged. The coefficient on $Trade_exposure \cdot post_BTA \cdot female$ remains negative and significant at the 5 percent level, indicating that the post-BTA decline in women's wages relative to men's persists even after accounting for nationwide, sector-specific shocks. The smaller magnitude of the triple-interaction term, about -0.17 compared with -0.20 in the baseline, suggests that part of the raw effect captured in the baseline specification may reflect industry-level shocks common across provinces, but a substantial and robust gender-differentiated component remains. However, the inclusion of additional fixed effects may also absorb part of the mechanism through which trade affects gender wage inequality, particularly if trade liberalization itself reshaped occupational or industrial structures along gendered lines. Overall, these results reinforce that the observed gender gap in the wage response to trade liberalization reflects genuine heterogeneity across local labor markets rather than artifacts of sectoral composition or national shocks.

Binary Treatment Indicator. Following common practice in the applied DiD literature,[23] I estimate alternative specifications using binary versions of the exposure measures at the provincial level. Specifically, I define $\mathbf{1}(More_exposed)_p$ as an indicator equal to one for provinces with above-median trade exposure and replace the continuous measure of $TradeExp_p$ in Equation (4) with its binary counterpart, $\mathbf{1}(More_exposed)_p$. This dichotomization provides a transparent treated–control comparison and mitigates concerns about functional-form dependence by relaxing the assumption of linearity in the relationship between trade exposure and wages.

Table A.4 reports regression results from Equation (4) using a binary treatment indicator instead of a continuous measure of trade exposure. Column (1) includes province and year fixed effects, Column (2) replaces year fixed effects with industry-by-year fixed effects, and Column (3) controls for province-by-year fixed effects. Across all specifications, the interaction term $BTA \times More_exposed \times Female$ is negative and statistically significant, confirming that the gender-differentiated effects of trade liberalization are robust to using a binary treatment definition. The coefficient on $Post_BTA \times More_exposed \times Female$ ranges from -0.035 to -0.049, implying that women in more-exposed provinces experienced wages approximately 3.5–4.9% lower than comparable men after the BTA. The magnitude of these effects is larger than in the baseline specification using the continuous exposure measure, likely because dichotomizing the treatment accentuates the difference between high- and low-exposure provinces by discarding within-group variation. The persistence of statistical significance across specifications indicates that the main findings are not sensitive to the treatment definition and that the relationship between the BTA and the widening gender wage gap is robust.

Alternative Measurement of Wages: In the main specification, I use annual wages as the dependent variable and control for total time worked on the right hand side. As a robustness check,

[23] See Kelishomi and Nisticò (2024), who examines the effects of international sanctions on labor market outcomes and defines treated workers as those employed in industries with above-median levels of trade exposure. See also Lu and Yu (2015), who study China's WTO accession to analyze the effect of trade liberalization on markup dispersion and assign industries to treatment and control groups based on whether they faced tariffs above or below the sample median.

I re-estimate Equation (4) using the hourly wage rate, measured as annual earnings divided by total time worked, as the dependent variable. This alternative measure normalizes earnings by hours worked, ensuring that the estimated effects are not driven by differences in labor supply between men and women or across provinces. Results from Table A.5 show that the interaction term $BTA \times More_exposed \times Female$ remains negative and statistically significant across specifications, although its magnitude is slightly smaller than in the baseline using annual wages. In Column (2), the coefficient is marginally significant at the 10% level, but the direction and overall pattern are consistent, indicating that the main findings are robust to alternative measurements of wages.

2.8 Mechanisms and Channels of Adjustment

2.8.1 Heterogeneity in Trade Effects

By demographic and skill characteristics: To understand how trade liberalization interacts with existing labor market frictions, I estimate equation (4) separately for subgroups defined by skill level, age, marital status, and urban–rural status. Skill is measured by distinguishing between low-skilled workers, those with nine or fewer years of education and no vocational or college training, and high-skilled workers. Additional subsamples compare younger workers (age 25 and under) with older workers, urban with rural residents, and married with single individuals. These dimensions capture meaningful differences in mobility, job-search constraints, and household responsibilities, allowing a richer characterization of whether the adverse effects of trade liberalization on women's relative wages are more pronounced among specific demographic or labor market groups.

Table 2.7. Heterogeneity in the Effect of Trade Exposure on Women's Wages

	Low Skill	**High Skill**	**Age ≤ 25**	**Age > 25**	**Urban**	**Rural**	**Married**	**Single**
Trade_Exp×Post-BTA	−2.791***	−1.389***	−2.234***	−2.677***	−2.768**	−1.304***	−2.719***	−1.760***
	(0.6485)	(0.2723)	(0.5993)	(0.3545)	(0.9801)	(0.2579)	(0.4207)	(0.4008)
Trade_Exp×Post-BTA×Female	−0.1301	−0.6391**	0.0788	−0.2406*	−0.1462	−0.2416*	−0.3174**	0.1201
	(0.1330)	(0.1940)	(0.1303)	(0.0921)	(0.1682)	(0.1142)	(0.1165)	(0.1748)
Worker characteristics as controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Province Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	87,320	70,982	46,084	112,218	100,709	57,593	102,345	50,027
R^2	0.57183	0.52574	0.60986	0.59931	0.57651	0.56574	0.60265	0.62056
Within R^2	0.44268	0.44391	0.51594	0.48256	0.46633	0.48702	0.48323	0.52259

Notes: Each column reports estimates from equation (4) on subsamples defined by worker characteristics. Low-skilled workers are those with nine or fewer years of education and no vocational or college training; high-skilled workers include those with vocational or college education. Additional subsamples compare younger (age 25 and under) versus older (over 25) workers, urban versus rural residents, and married versus single individuals. Robust standard errors clustered by province are reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2.7 reports the estimated effects of trade exposure on women's wages across these subgroups. The coefficient on the double interaction term (*Trade Exposure* × *Post-BTA*) is negative and statistically significant in every case, consistent with results from my main specification. The estimates of the triple interaction term (*Trade_exposure* × *Post_BTA* × *Female*) reveal important heterogeneity in the gendered impact of trade. High-skilled, older, rural, and married women experience the largest additional wage penalties relative to their male counterparts. These patterns suggest that trade liberalization increased gender inequality along dimensions where women face structural disadvantages. Women with characteristics that make labor market participation more constrained (married, older, rural) experienced the largest losses. By contrast, younger and single women show no detectable gender differentiated effects, consistent with greater flexibility in job search and occupational adjustment. Potential mechanisms underlying these heterogeneous effects include gendered sorting across industries and occupational, differences in mobility or flexibility that generate compensating differentials, and discrimination or biased promotion within expanding export-oriented sectors.

By employment type: I also examine whether the effects of trade exposure vary across segments of the labor market defined by employer type. Using the VHLSS employment classification, I estimate equation (4) separately for workers in the informal sector (including the self-employed and those working for other households), state-owned enterprises (SOEs), private domestic firms, and foreign-invested firms. Table 2.8 reveals substantial heterogeneity in how trade shocks translate into wage changes. The double interaction term (*Trade_exposure* × *Post-BTA*) is large, negative, and statistically significant only in the informal sector, indicating that workers in informal or household-based employment experience the strongest overall wage compression following the BTA. One plausible explanation for this pattern is that the BTA triggered a major reallocation of workers out of microenterprises and into expanding formal firms, and this reallocation is likely correlated with unobserved productivity differences across workers. McCaig and Pavcnik (2018) document that the BTA induced sizable movements of labor from informal household businesses toward formal employers, particularly in industries experiencing the largest

declines in export costs and in more internationally integrated provinces. As more productive and mobile workers transition into formal employment, the remaining informal workforce becomes increasingly negatively selected. This compositional downgrading, combined with highly flexible wage-setting and thinner profit margins in household enterprises, results in disproportionately large declines in average informal-sector earnings.

Table 2.8. Heterogeneity in the Effect of Trade Exposure on Women's Wages by Employer Type

	Unofficial	**SOE**	**Private**	**Foreign**
Trade_Exp×*Post-BTA*	−3.110***	−0.2556	−0.5912	−0.4831
	(0.7006)	(0.4270)	(0.3415)	(0.4755)
Trade_Exp×*Post-BTA*×*Female*	−0.1990	−0.7376***	0.1344	−0.0598
	(0.1687)	(0.1307)	(0.0893)	(0.2500)
Worker characteristics as controls	Yes	Yes	Yes	Yes
Province Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
Observations	80,399	48,523	21,985	7,395
R^2	0.59041	0.45923	0.61718	0.63088
Within R^2	0.46820	0.37404	0.52527	0.55535

Notes: Each column reports estimates from equation (4) for a different employer type. "Unofficial" includes self-employed workers and employees in informal household enterprises; "SOE" refers to state-owned enterprises; "Private" refers to domestic private firms; "Foreign" refers to foreign-invested enterprises. Robust standard errors clustered by province are shown in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Turning to gender differences, the estimates of the triple interaction term (*Trade_exposure* × *Post-BTA* × *Female*) indicate that the gendered effects of trade liberalization differ markedly by employer type. The coefficient is large, negative, and statistically significant only for SOEs, suggesting that women in the state sector experienced substantially greater trade-induced wage losses relative to men. By contrast, the corresponding estimates for private, foreign, and unofficial firms are small and statistically insignificant, indicating no detectable gender-differentiated effects in those segments. One possible explanation is that SOEs have greater scope for non-market incentives and entrenched gender norms to influence wage setting, which can amplify gender disparities when the sector adjusts to trade opportunities. A complementary channel is selection: women may disproportionately choose SOE jobs for their stronger job security and compliance

with statutory benefits (e.g., maternity leave and reinstatement protections)[24].

Overall, these results highlight that the gender consequences of trade liberalization are shaped not only by worker characteristics but also by the institutional features of the firms and sectors in which workers are employed.

2.8.2 Decomposition of the Provincial Gender Wage Gap

To better understand how trade liberalization translated into gendered wage outcomes, I examine the extent to which differences in wages between men and women at the provincial level reflect disparities in observable worker characteristics versus differential returns to those characteristics. This decomposition serves as a useful diagnostic for whether trade-induced adjustments operate primarily through shifts in worker composition, such as changes in women's educational attainment or sectoral sorting, or through changes in how the labor market rewards similar characteristics across genders.

The Oaxaca–Blinder decomposition, widely used in studies of wage inequality (Oaxaca, 1973; Blinder, 1973; Blau and Kahn, 2017), separates the overall wage differential into portions explained by differences in observable characteristics and those arising from differences in returns to those characteristics. For each province p and year t, I estimate separate Ordinary Least Squares wage equations for men and women:

$$w_i = X_i B_i + \varepsilon_i \quad \text{for } i \in \{m, f\}$$

where w_i is the natural log of wages, X is a vector of worker characteristics, B the corresponding coefficients, and ε_i an error term.[25] The covariate set X includes educational attainment, experi-

[24] According to Vietnam's Labour Code of June 1994 (Vietnamese National Assembly, 1994), statutory maternity leave was approximately four months and applied to insured workers across sectors through the compulsory social insurance system. SOEs, however, were more likely to enroll workers in social insurance and to comply fully with statutory leave and job-protection rules, making them relatively more attractive to women seeking stable employment and benefit coverage.

[25] The p and t subscripts are omitted for brevity.

ence and its square, marital status, urban residence and ethnicity.

Let β_m and β_f denote the estimated coefficients and bars denote group means. Because OLS with a constant term yields mean-zero residuals, the average log wage differential can be expressed as:

$$\bar{w}_m - \bar{w}_f = \underbrace{(\bar{X}_m - \bar{X}_f)\beta_m}_{\text{Explained component}} + \underbrace{\bar{X}_f(\beta_m - \beta_f)}_{\text{Unexplained component}} . \tag{2.5}$$

The left-hand side represents the raw gender wage gap; the first term on the right captures the portion explained by differences in observable characteristics, while the second term - the unexplained or residual gap - reflects differences in returns to those characteristics, potentially arising from discrimination, compensating wage differentials associated with flexible hours, or unobserved heterogeneity.

To study how trade liberalization affected each component of the provincial gender wage gap, I regress the *raw gap*, the *explained* component, and the *unexplained* component on provincial trade exposure using the following specification:

$$\hat{Y}_{p,t} = \beta \mathbb{I}(t > 2002) \times TradeExp_p + \mu_p + \tau_t + u_{pt}, \tag{2.6}$$

where $\hat{Y}_{p,t}$ denotes the estimated overall gender wage gap, the explained component, or the unexplained component for province p in year t. All regressions include province fixed effects μ_p and year fixed effects τ_t. Each regression is estimated using weighted least squares, with weights equal to the inverse of the estimated variance of the corresponding Oaxaca-Blinder component, ensuring that more precisely estimated province-year cells receive greater weight. The coefficient β captures whether provinces more exposed to the BTA experienced differential changes in the raw gap or in its explained and unexplained components. A positive coefficient indicates that trade exposure increased the relevant component of the gender wage gap, while a negative coeffi-

cient implies a reduction.

Table 2.9. Estimated Effects of Provincial Trade Shocks on Components of the Gender Wage Gap

Dependent variable:	*Wage gap*	*Unexplained*	*Explained*
Trade_exposure × Post-BTA	0.3322 (0.2120)	0.3815* (0.1681)	0.0298 (0.1132)
Year Fixed Effects	Yes	Yes	Yes
Province Fixed Effects	Yes	Yes	Yes
Observations	300	300	300
R^2	0.81759	0.72791	0.75590
Within R^2	0.00537	0.00855	0.00023

Notes: Each column reports estimates from regressions of provincial wage gap components (total, unexplained, and explained) on trade exposure interacted with a post-BTA indicator. Dependent variables are constructed from the Oaxaca–Blinder decomposition in equation (5). All regressions are estimated using weighted least squares, with weights equal to the inverse of the estimated variance of each dependent variable. Standard errors, reported in parentheses, are clustered by province.
Significance levels: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$.

Table 2.9 examines how provincial trade exposure affected the *overall* gender wage gap and its *explained* and *unexplained* components. The estimates reveal three key patterns. First, the coefficient on the raw gender wage gap is positive (0.332) but statistically insignificant, indicating that provinces more exposed to the BTA did not experience systematically larger or smaller *overall* gender wage gaps. Second, trade exposure is positively and significantly associated with the *unexplained* component of the gender wage gap (0.382, significant at the 5% level). This implies that, following the BTA, provinces with higher trade exposure experienced a widening in the portion of the gender wage gap attributable to differences in returns to characteristics rather than differences in characteristics themselves. In other words, trade liberalization appears to have amplified structural or discriminatory forces underlying gender wage inequality. Third, the effect on the *explained* component is small (0.030) and statistically insignificant, suggesting that trade exposure did not meaningfully alter gender differences in observable worker characteristics such as education, experience, or demographic composition. Instead, the widening of the gender wage gap in more exposed provinces is driven almost entirely by growing disparities in how the labor market rewards men and women with the same observable traits. Taken together, these results

indicate that the BTA increased the structural or residual component of gender wage inequality without altering the observable composition of male and female workers.

2.9 Conclusion

This paper provides new evidence on how trade liberalization shapes gender wage inequality by studying the 2001 U.S.–Vietnam Bilateral Trade Agreement (BTA), a major and exogenous reduction in U.S. tariffs on Vietnamese exports. Exploiting cross-provincial differences in exposure to the BTA, constructed from pre-BTA industry composition and industry-specific tariff cuts, I combine province-level and individual-level analyses to document how this trade shock altered both overall wages and the relative earnings of men and women. At the provincial level, I employ a two-stage estimation procedure to first recover provincial wage premiums by gender and then relate these to province-level trade exposure. At the individual level, I implement a triple difference-in-differences framework to identify gender-differentiated effects of trade liberalization while controlling for province and year fixed effects and a rich set of worker characteristics.

Across provinces, the BTA generated substantial dispersion in wage outcomes: provinces more exposed to tariff reductions experienced significantly smaller wage increases, and both men and women saw compression in provincial wage premiums. Results from the individual-level analysis corroborate this pattern of regional convergence. However, heterogeneity analysis by employment type reveals an important nuance; this convergence is driven entirely by the informal sector. No comparable evidence of wage compression is found in the state-owned, private, or foreign-invested sectors. A plausible explanation is that the BTA accelerated transitions from informal to formal employment, as documented by McCaig and Pavcnik (2018), and that this reallocation is correlated with unobserved worker productivity, leaving behind a negatively selected pool of workers in the informal sector and contributing to lower average wages there.

Turning to gender differences, I find that women experienced an additional wage penalty relative to men within the same province. A one–standard-deviation increase in provincial trade

exposure led to an additional decline in women's wages of roughly 0.8 percent. Results from the Oaxaca–Blinder decomposition indicate that trade exposure widened gender wage inequality primarily through an increase in the unexplained component of the wage gap, consistent with larger differences in returns to characteristics rather than differences in endowments. Heterogeneity analyses further show that the gender penalty is concentrated among women employed in SOEs and among high-skilled, married, older, and rural women, who face greater constraints on mobility and stronger household responsibilities. In contrast, no gender differentials are detected in the private or foreign sectors, suggesting that firms more exposed to market competition and international standards operate under more gender-neutral wage-setting environments. These results imply that institutional and social environments, particularly persistent gender norms and occupational sorting, play a central role in shaping how trade shocks translate into unequal outcomes by gender in Vietnam

This chapter has several limitations. First, since the analysis relies on household survey data rather than firm-level data, the results cannot fully disentangle labor demand responses from labor supply adjustments. Second, while the findings document important changes in gendered labor market outcomes following trade liberalization, the data do not allow direct observation of intra-household bargaining, occupational transitions, or workplace-level mechanisms that may mediate these effects. For example, the analysis focuses primarily on wages and therefore may not capture non-wage dimensions of welfare, such as work time flexibility, working conditions, or other amenities that could partially offset observed wage penalties for women. Future work combining household and firm-level data could provide a more complete understanding of the channels through which trade liberalization affects gender inequality.

Despite these limitations, the findings demonstrate that the distributional consequences of trade liberalization depend not only on exposure to international markets, but also on the institutional and social environments through which adjustment occurs. This finding is particularly concerning in light of emerging evidence on other gendered responses to the BTA. Huynh and

Ku (2025) show that positive income shocks from the BTA reduced women's labor market participation — especially among married, affluent Southerners — highlighting the role of gender role attitudes in shaping behavioral responses to rising household wealth. Similarly, Huynh and Nguyen (2025) find that women in more exposed industries were more likely to have male children and work longer hours, with effects driven solely by maternal rather than paternal exposure. These patterns suggest that trade liberalization is not inherently harmful to women; rather, it interacts with preexisting structural factors — such as gender norms, occupational segregation, and institutional rigidity — that amplify unequal outcomes as economies adjust to globalization. The results highlight the need for Vietnam to confront persistent gender norms that constrain women's economic choices. When women's primary role continues to be viewed as caregiving and paid work is treated as secondary, the new opportunities created by trade liberalization may inadvertently reinforce, rather than reduce, existing gender inequalities.

Chapter III: From Trade War to Job Growth: Evidence from Vietnam

Anca Cristea

University of Oregon

Hoang Pham

Oregon State University

Ngan Tran

University of Oregon

Abstract

We study how the 2018–2019 U.S.–China trade war affected employment in Vietnamese manufacturing firms, exploiting cross-industry variation in U.S. tariff increases on Chinese goods as a plausibly exogenous demand shock to Vietnamese producers. Using eight waves of the Vietnam Enterprise Survey (2015–2022) matched to industry-level tariff exposure measures, we estimate difference-in-differences specifications with firm and year fixed effects. We document three main findings. First, firms in more exposed industries experienced significant employment growth: moving from zero to the mean tariff exposure is associated with a 4.7% increase in employment for the average manufacturing firm, rising to 12.6% among firms with pre-trade-war international trading experience. Second, employment gains are highly heterogeneous across ownership types: foreign-owned firms respond most strongly, followed by state-owned enterprises, while private domestic firms benefit only when they had prior export market experience. Third, employment expansion was not compositionally neutral: exposed firms disproportionately hired female workers, increasing both the level and share of female employment, while average wages declined among the most exposed pre-trade-war traders, consistent with rapid scaling through lower-wage hires. These results highlight the importance of firm heterogeneity and pre-existing trade linkages in shaping the labor market consequences of global trade shocks.

3.1 Introduction

The 2018–2019 U.S.–China trade war was one of the largest disruptions to global trade in recent decades. As the United States imposed successive tariff increases on Chinese goods, raising average tariffs from about 3% to over 19%, firms were forced to reorganize global supply chains and reconsider sourcing decisions (Freund et al., 2024). While a large literature has examined the direct effects of these policies on the U.S. and China, less is known about their consequences for third-country bystanders. This paper focuses on one such country, Vietnam, and asks whether the demand shock resulting from the trade diversion effect of the U.S.-China trade war translated into meaningful employment gains for Vietnamese manufacturing firms.

Vietnam is a compelling setting for this question. Following the 2001 U.S.–Vietnam Bilateral Trade Agreement, the U.S. became a key export destination, accounting for roughly one-fifth of Vietnam's total exports prior to the trade war. At the same time, Vietnamese manufacturers were deeply embedded in regional supply chains and closely linked to Chinese production networks, in part due to geographical proximity as well as substantial Chinese foreign direct investment. These features meant that Vietnamese firms were well positioned to absorb demand redirected away from Chinese exporters. Consistent with this, Alfaro and Chor (2023) identify Vietnam as a major beneficiary of the resulting shift in U.S. import sourcing — what they term the "great reallocation." Importantly, the surge in exports appears to reflect genuine production expansion rather than simple transshipment: Schulze and Xin (2025) document rising domestic value added in key sectors, while Iyoha et al. (2024) show that true rerouting accounted for only 1.8% of Vietnamese exports to the U.S. in 2021.

A key question is how these trade-driven demand shocks translate into labor market outcomes. Standard trade models would predict that a labor-abundant country like Vietnam should experience broad employment gains when export demand rises. However, theories of heterogeneous firms suggest a more nuanced response: increased demand may lead to reallocation toward more productive firms rather than uniform expansion across all firms. In addition, firms can ad-

just along multiple margins, including capital investment, technology upgrading, and workforce composition, so employment growth is not guaranteed. Whether employment expands, and for whom, is therefore an empirical question.

We address this question using firm-level data from eight waves of the Vietnam Enterprise Survey (VES) covering the period 2015–2022, matched to industry-level measures of trade war exposure constructed from product-level U.S. tariff data compiled by Bown (2021). Our econometric identification strategy exploits cross-industry variation in the magnitude of U.S. tariff increases on Chinese goods. Since Vietnam played no role in the trade dispute, this variation is plausibly exogenous to Vietnamese firm outcomes. We estimate difference-in-differences specifications with firm and year fixed effects, controlling for time-varying industry-specific trends that might confound the trade diversion channel.

We document three main findings. First, firms in more exposed industries experienced significant employment growth. Moving from zero to the mean level of tariff exposure increases employment by 4.7% for the average manufacturing firm, with substantially larger effects, around 12.6%, among firms with trade experience prior to the trade war. Second, the employment response is highly heterogeneous across ownership types. Foreign-owned firms exhibit the largest gains, followed by state-owned firms, while private domestic firms benefit only when already integrated into export markets prior to the trade war. Third, the composition of employment changes alongside overall growth. The expansion disproportionately increased female hiring, with exposed firms raising both the level and share of female workers. Average wages, however, declined among the most exposed firms, suggesting that firms hired new workers at the lower end of the wage distribution as they scaled up production to meet new export demands.

Our paper contributes to two main strands of literature. First, it extends a rapidly growing literature documenting the impact of the U.S.–China trade war on bystander countries. Early evidence from Mexico shows significant trade diversion effects. Utar et al. (2025) find that Mexican exporters gained market share in U.S. product categories newly subject to tariffs on Chinese

goods, while Chen et al. (2025) document positive employment responses concentrated in manufacturing industries with high exposure to trade diversion. A parallel literature examines third-country effects in Asia, documenting export gains in economies such as Taiwan, South Korea, and Vietnam (Fajgelbaum et al., 2024; Nicita, 2019). Within this literature, several papers focus on employment effects in Vietnam (Mayr-Dorn et al., 2026; Nguyen and Lim, 2023; Rotunno et al., 2024; Wu, 2024).

We make three main contributions to this strand. First, we confirm that ownership heterogeneity is central: employment effects are strongest among foreign-owned firms. However, we show that domestic private firms also respond to trade diversion, but only conditional on prior export experience—a margin of heterogeneity not previously emphasized in the literature that helps explain why domestic firms appear unresponsive in aggregate. Second, we document systematic differences in the composition of new hires across ownership types, with employment expansion skewed toward female and lower-wage workers, revealing distributional effects that aggregate employment measures obscure. Third, we show that employment effects are persistent and continue to grow over our extended sample through 2022, rather than dissipating after the initial shock. The implied magnitudes—4.7 percent at mean exposure and 12.6 percent among prior traders—are substantially larger than those documented in existing work.

Our paper also contributes to a second strand of literature on the causal impact of trade policy on labor markets. A large body of work, including Dix-Carneiro and Kovak (2017), shows that trade shocks can have long-lasting effects on employment and local labor market outcomes. While our setting does not allow us to study effects over similarly long horizons, we provide evidence that trade war–induced demand shocks generate employment responses that are not only persistent but continue to strengthen over time, consistent with gradual adjustment and reallocation dynamics at the firm level.

The remainder of the paper is organized as follows. Section 2 provides institutional background on the U.S.–China trade war, describes Vietnam's trade patterns before and after 2018,

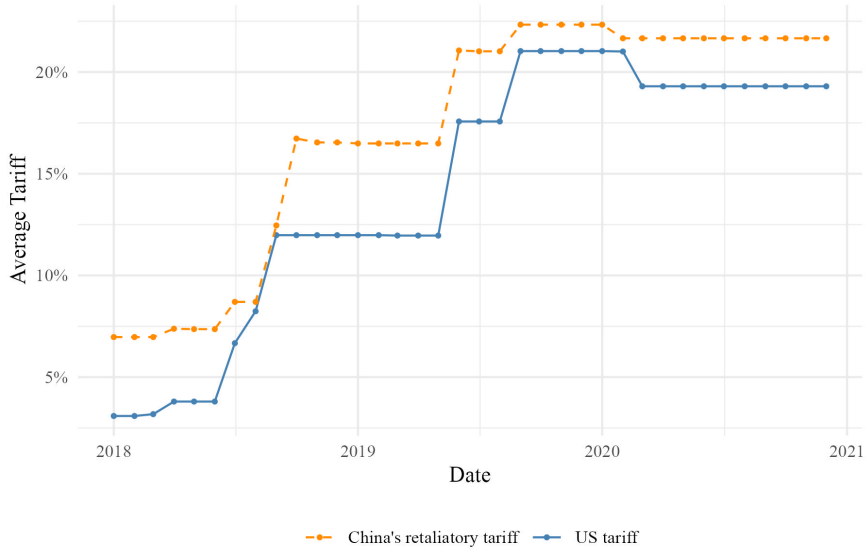


Figure 3.1. Progression of U.S. Import Tariffs on China during 2018-2019 Trade War

Note: Tariff data from Bown (2021).

and outlines the theoretical framework motivating our empirical analysis. Section 3 describes the data sources. Section 4 presents the empirical methodology and measurement of trade war exposure. Section 5 reports the main estimation results and describes the robustness checks performed to solidify our findings. Section 6 concludes.

3.2 Institutional Background

3.2.1 Overview of the U.S.-China Trade War

Prior to 2017, U.S.–China trade relations were largely governed by multilateral rules under the World Trade Organization (WTO), with disputes addressed through conventional mechanisms such as anti-dumping and countervailing duties alongside formal dispute settlement processes. The Trump administration's initiation of five major trade investigations in 2017, targeting steel, aluminum, solar panels, washing machines, and broader unfair trade practices, marked a turning point that set the stage for executive-driven trade policy (Bown, 2021). The trade war began in February 2018 with the imposition of safeguard tariffs on solar panels and washing machines, followed shortly by steel and aluminum tariffs applied under national security provisions. It esca-

lated sharply in July 2018 as the U.S. introduced successive rounds of tariffs targeting hundreds of billions of dollars in Chinese goods, which were met with retaliatory tariffs from China. The conflict continued to intensify throughout 2019 before slowing down with the announcement of the Phase One agreement in December.[26] Signed in January 2020 and implemented in February, the agreement paused further escalation and included limited commitments on trade and structural issues. By that time, average U.S. tariffs on Chinese goods had risen to 19.3%, while China's retaliatory tariffs averaged 21.7% (see Figure 3.1).

Figure 3.2 presents the distribution of cumulative tariff changes at the HS6 level between 2017 and 2020. Approximately 4,873 out of roughly 5,751 HS6 product lines experienced tariff increases. Since tariffs were applied in successive rounds targeting predefined lists of goods, the distribution is discrete with mass concentrated at a few distinct values, most prominently at around 8 and 22 percentage points. The spike at 22 percentage points is particularly pronounced, reflecting the fact that roughly 46% of all HS6 lines received cumulative increases in the narrow range of 21.2 to 22.3 percentage points. When aggregated to the ISIC Rev. 4 industry level using China's pre-trade-war import weights, the distribution becomes substantially smoother, with greater dispersion across industries (Figure C.1[27]). This aggregation generates more continuous variation in exposure and provides the identifying variation for our empirical analysis.

3.2.2 Vietnam's Trade Patterns and Stylized Facts

As the U.S. and China exchanged escalating rounds of tariffs, global supply chains faced significant disruption, pushing firms to reassess where they sourced their inputs and goods (Freund et al., 2024). Alfaro and Chor (2023) characterizes the resulting shift in U.S. import sourcing as the opening stage of a "great reallocation," with Vietnam emerging as one of the principal beneficiaries of diverted trade flows. Vietnam was well positioned to take advantage of the U.S.-China trade war for several reasons. Vietnam has had a strong and growing trading relation with the

[26] See Table C.1 for a detailed timeline of U.S. and China's tariff actions from February 2018 to early 2021

[27] More details about the construction of this industry-level exposure is provided in Section 3.5.1

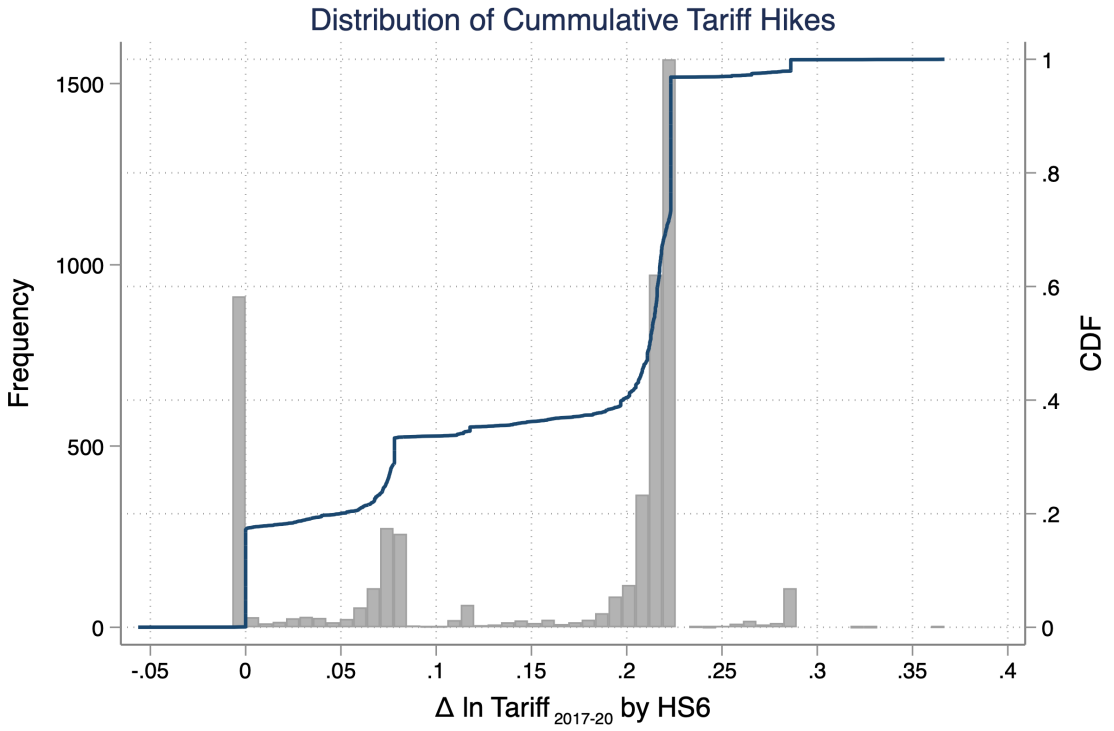


Figure 3.2. Distribution of Cumulative Tariff Increases at HS6 Level

U.S. ever since the U.S.-Vietnam Bilateral Trade Agreement (BTA) signed in 2001[28]. The U.S. has been an important destination for Vietnam's exports, with a steadily growing share, accounting for roughly one-fifth of total exports in the years leading up to the trade war (Figure C.3). By having already established trade networks, Vietnamese firms could take advantage more easily and quickly of the market opportunities created by the withdrawal of Chinese suppliers from the U.S. markets. Moreover, Vietnam has been increasingly immersed in South East Asian global supply chains, with Chinese firms investing heavily in Vietnam years before the 2018-2019 trade war[29]. Familiarity with Chinese products and production processes implied that Vietnamese firms, including foreign owned enterprises, were better positioned to supply the U.S. market with

[28]Previous research finds that the BTA lowered poverty rates, increased long-term employment, and triggered structural transformation from agriculture to manufacturing in Vietnam (McCaig, 2011; McCaig and Pavcnik, 2018; McCaig et al., 2025)

[29]According to the General Statistics Office of Vietnam, total registered capital from China (excluding Taiwan) in Vietnam increased from US$700 million in 2011 to US$2.4 billion in 2018 - more than tripling over the period, with an average annual growth rate of about 18% (Lam, 2019).

close input substitutes.

For context, in what follows, we describe a few facts about Vietnam's aggregate trade patterns before and after the U.S.-China trade war. We use BACI-CEPII data for the period 2014-2022 to get annual import and export values for Vietnam at the HS 6-digit product code by the foreign country. U.S. tariff data at the HS 6-digit level (to match the CEPII data) is available from Bown (2021). More on the data sources and data construction is relegated to the data section.

Fact 1: Vietnam's trade grew rapidly following the U.S.-China trade war

Figure 3.3 reports aggregate trends in Vietnam's trade over the period 2014-2022. The left panel reports Vietnam's exports to the U.S. and to all foreign destinations other than the U.S., combined. Export values are normalized by the respective 2017 values to facilitate comparisons across trends. Consistent with findings from the existing literature, Vietnam has benefited from the U.S.-China trade war by capitalizing on the market opportunities created by the exit of Chinese suppliers to the U.S.. As panel (a) of Figure 3.3 shows, exports to the U.S. grew more rapidly relative to other countries after 2018. During the pre-trade war period, however, the two trend lines progress at a similar rate. Interestingly, Vietnam's export growth in the U.S. market did not come at the expense of exports to other destinations, which seem to point to an increase in production capacity in Vietnam to accommodate the positive demand shock rather than a pure reallocation of exports across destination markets.

Panel (b) of Figure 3.3 plots Vietnam's imports from China and from all other foreign sources other than China, normalized by their respective 2017 values. Here again we notice that the two trend lines evolve at the same rate until 2018, after which there is a non-trivial departure in trajectories with imports from China outpacing those from other foreign sources after 2018. This pattern is consistent with the possibility of trade rerouting during the U.S.–China trade war, whereby Chinese goods are shipped through Vietnam to circumvent tariffs. However, Iyoha et al. (2024) find evidence that not all of the observed increase in Vietnamese imports from China reflects tariff evasion through rerouting. Using firm-level transaction data, they find that true rerouting, de-

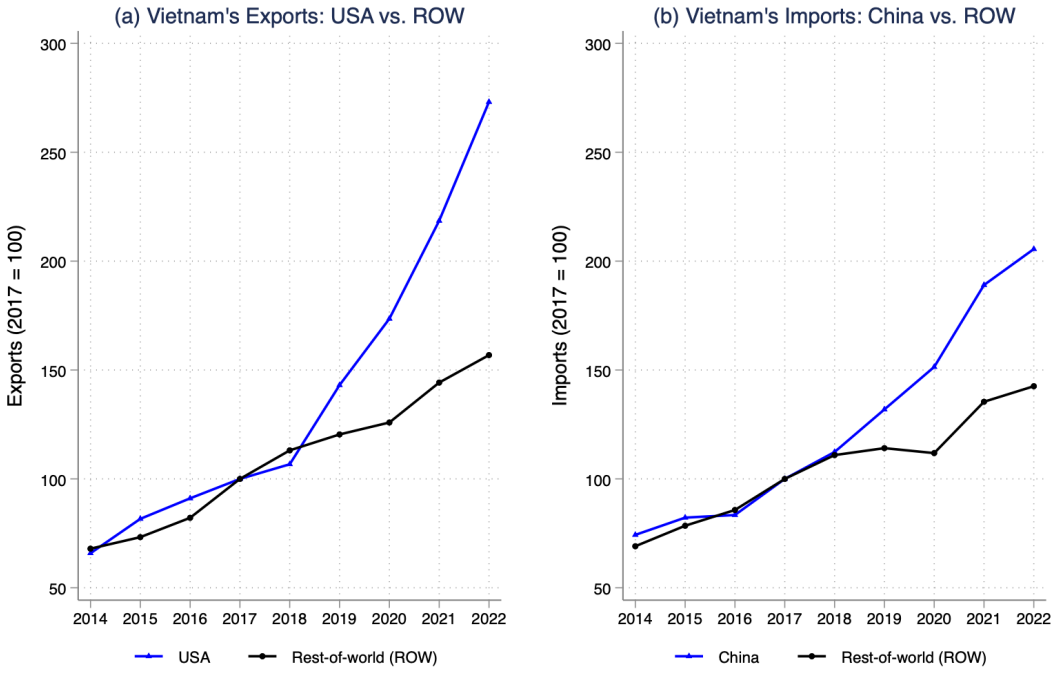


Figure 3.3. Vietnam's Aggregate Trade

fined as the import and re-export of identical products within the same firm, accounted for only 1.8% of Vietnamese exports to the United States in 2021, suggesting that a substantial share of the import surge reflects substantial changes in input sourcing and supply chain reorganization.

Fact 2: The trade war was pervasive, affecting most of Vietnam's trade

The U.S.-China tariff escalation extended rapidly across traded products. Out of approximately 5350 HS 6-digit products imported by the U.S. between 2017-2019, 84% of them involved a tariff hike during that period. A relevant question for us is whether Vietnam has traded these products targeted by the trade war. Figure 3.4 illustrates these patterns using Vietnam's global exports and imports. Targeted products are defined as the HS 6-digit categories that experienced a tariff hike at any point during the 2018-2020 period. It is clear from the reported trends that Vietnam traded extensively in the HS 6-digit products targeted by tariff hikes even before the start of the U.S.-China trade tensions. Not only that, but the share of trade in targeted categories is high, accounting on average for 82% of Vietnam's exports, respectively 90% of Vietnam's imports over

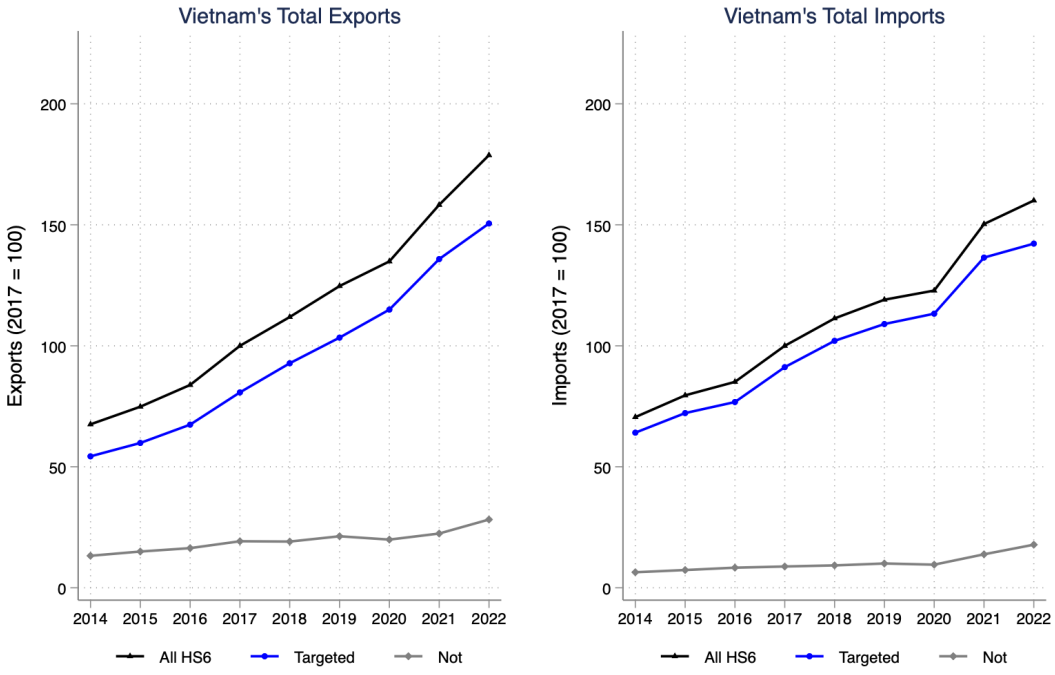


Figure 3.4. Vietnam's Aggregate Trade

the 2014-2022 period.

Fact 3: Vietnam's trade growth is driven by goods affected by the trade war

Figure 3.5 is informative about aggregate trends, however it does not indicate whether the growth in Vietnam's trade with the U.S., respectively with China, is driven by the sectors that experienced the highest surge in tariffs during the 2018-2019 trade war. To investigate this, Figure 3.5 reports a binscatter plot that correlates the total change in U.S. tariffs on China at HS 6-digit product level against average annual changes in Vietnam's porduct level exports over three time intervals: i) the period prior to the trade war (2015-2017), ii) the period during the trade war (2017-2019), and iii) the period after the trade war (2019-2022).[30] The left panel of Figure 3.5 reports growth trends for Vietnam's exports to all destinations other than the U.S., while the right panel reports Vietnam's exports to the U.S. Focusing on aggregate exports first (left panel), it is re-assuring to see that prior to 2018 Vietnam did not export more in HS 6-digit products that

[30]Our binscatter plot is motivated by Fajgelbaum et al. (2024), who implement a similar graph in Figure 1 of their paper.

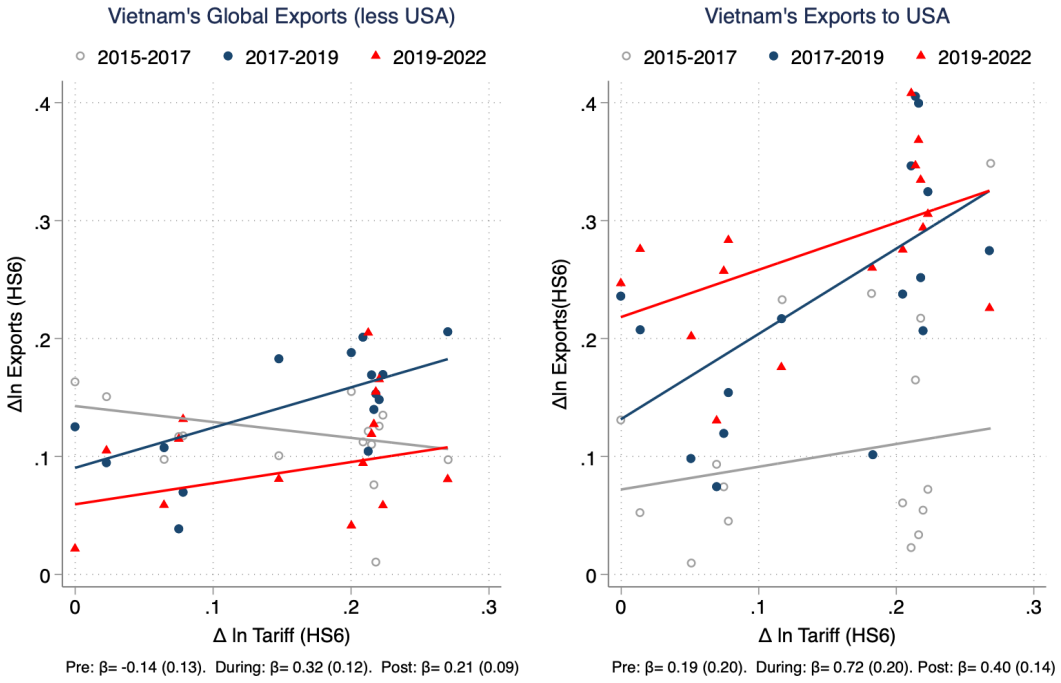


Figure 3.5. Changes in Vietnam's Exports by Sector

would subsequently be targeted by Trump tariffs. If anything, the fitted line points to a slight negative yet statistically insignificant correlation (see the slope coefficients at the bottom of the plot area). When focusing on trade with the U.S. (right panel), there is some weak evidence that exports grew faster in HS 6-digit products that would subsequently be targeted by the Trump tariffs, although again the correlation is statistically insignificant. Importantly for our analysis, during the trade war period of 2018-2019, Vietnam saw its exports grow more rapidly in the HS 6-digit product categories affected by the tariff hikes. This accelerated growth trend continued even after the Phase One trade deal was signed in February 2020 between the U.S. and China, which paused the escalation of tariffs.

Figure 3.6 reports a similar binscatter plot using Vietnam's imports, both global imports from all sources other than China, and separately for imports from China alone. Interestingly, while Vietnam's overall imports increased both before and after the trade war (as shown in Figure 3.3), the growth in imports has not been concentrated in those HS 6-digit sectors hit the hardest by the

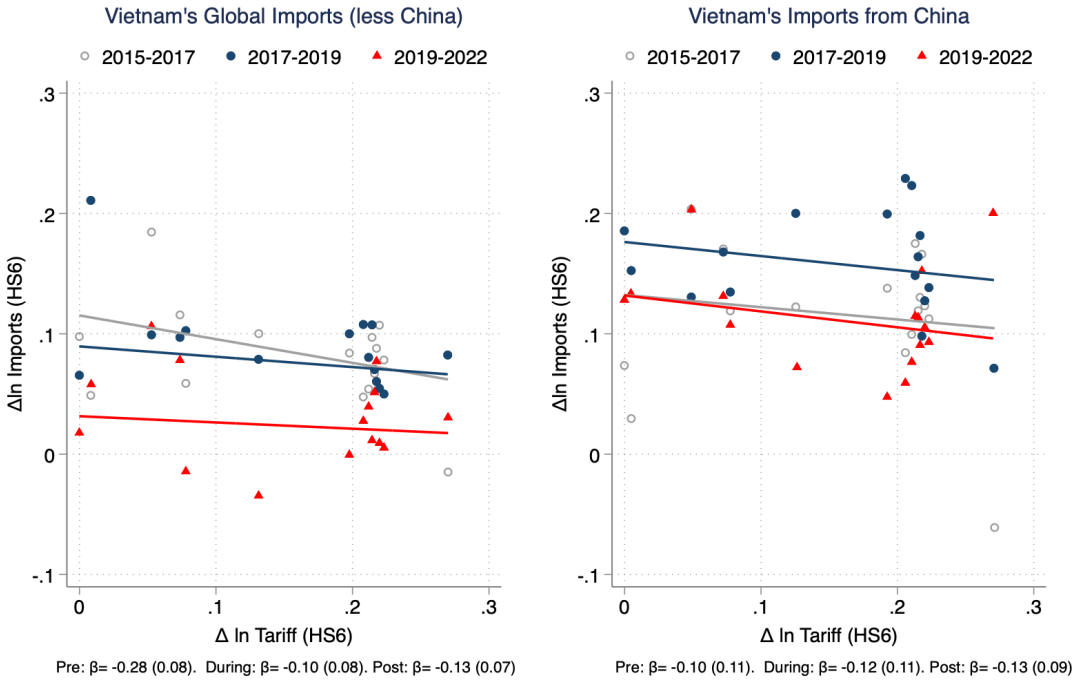


Figure 3.6. Changes in Vietnam's Imports by Sector

Trump tariffs. On the contrary, judging from the negative slope of the fitted lines corresponding to the post trade war period (2019-2022), Vietnam seems to have experienced weaker growth in the imports of targeted products, particularly when it comes to imports from China.

3.3 Conceptual Framework and Empirical Evidence

3.3.1 Employment Effect: Theoretical Framework

A central question in this paper is whether the positive demand shocks induced by US's trade diversion translate into employment expansion, or whether firms primarily adjust along other margins. Standard Heckscher-Ohlin and Stolper-Samuelson modes predict that a country like Vietnam should gain broadly when export demand rises since labor is its abundant factor. However, in canonical models of heterogeneous firms (Melitz, 2003), trade-induced demand shocks lead to a reallocation of resources toward more productive firms[31]. This mechanism does not imply uni-

[31]Both standard neoclassical trade models and canonical heterogeneous-firm models typically assume full employment or frictionless labor reallocation. As a result, while these frameworks are informative about resource allocation

form employment growth: while more productive firms may expand, less productive firms may contract or exit, resulting in ambiguous net employment effects.

Moreover, firms facing increased export opportunities may respond along multiple adjustment margins beyond employment. One important channel is capital deepening and technology adoption. For instance, Bustos (2011) shows that Argentine firms exposed to trade liberalization upgrade technology and increase capital intensity. To the extent that such technological upgrading is labor-substituting, it may dampen employment growth. Similarly, firms may engage in quality upgrading rather than workforce expansion. Verhoogen (2008) finds that Mexican firms entering export markets upgrade product quality and pay higher wages to attract more skilled workers, suggesting a shift toward higher-value production rather than larger employment.

Another important margin of adjustment is worker composition. Brambilla et al. (2012) document that Brazilian firms exporting to high-income destinations employ more skilled workers and pay higher wages compared to firms serving domestic or lower-income markets. This implies that firms may respond to export-driven demand by upgrading the skill composition of their workforce rather than expanding total employment. In these cases, wage adjustments themselves may absorb part of the demand shock, especially if firms compete for higher-quality labor.

In the context of Vietnam, evidence from an earlier trade liberalization episode suggests that trade-induced demand shocks can translate into meaningful employment gains. Studying the 2001 US–Vietnam Bilateral Trade Agreement, McCaig and Pavcnik (2018) document substantial reallocation of labor toward more export-oriented sectors, accompanied by declines in informality. More recent work by McCaig et al. (2025) finds that these employment gains persist and even strengthen over the longer run.

Overall, theory suggests that the employment response to trade-induced demand shocks is not predetermined and depends on the relative importance of reallocation, technology adoption,

and wage effects, they are less well suited for analyzing short- to medium-run employment responses to trade shocks, especially in settings with labor market frictions, adjustment costs, or informal employment.

and changes in workforce composition. This motivates an empirical investigation into how Vietnamese firms adjust employment in response to the U.S.–China trade war

3.3.2 Employment Effect of the Trade War: Empirical Evidence

This paper contributes to a rapidly growing literature on the spillover effects of the U.S.–China trade war, with a particular focus on Vietnam as a key beneficiary of trade and production diversion. Mayr-Dorn et al. (2026) provide the most comprehensive individual-level evidence to date. Exploiting district-level variation in exposure to U.S. tariff hikes, constructed using a Bartik-style shift-share instrument based on pre-war industry employment shares, they show that individuals in more exposed districts were more likely to work in traded industries after 2018, with modest increases in wages and hours. Their most economically meaningful result concerns formalization: exposure reduces informal employment by 0.5 percentage points at the interquartile range, accounting for roughly a quarter of Vietnam's total formalization between 2017 and 2019. The employment-level effect, while statistically robust, is modest in magnitude — a 0.4 percentage point increase in employment probability at the interquartile range — and serves primarily to establish that trade diversion reached local labor markets rather than to quantify the scale of job creation.

Rotunno et al. (2024) is closest to our setting. Using firm-level data from the Vietnamese Enterprise Survey and a binary indicator for firms exposed to tariff increases above 15 percent, they find that treated firms created about 5 percent more jobs, with effects emerging only in 2020 — the final year of their sample and the beginning of our extended window. We confirm this delayed response and show that employment gains continue to grow beyond 2020. Methodologically, their use of recent difference-in-differences estimators that address heterogeneous treatment effects and staggered timing (Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; Dube et al., 2025) complements our continuous-exposure approach.

Nguyen and Lim (2023) examine structural transformation using continuous exposure measures, showing that more exposed provinces reallocate workers from informal agriculture to formal manufacturing, consistent with the formalization results of Mayr-Dorn et al. (2026). Employment effects, however, are confined to foreign firms and appear only in 2019. They also document skill upgrading, finding that only workers in FDI firms experience improvements in IT-related skills. Wu (2024) further sharpens this ownership heterogeneity, showing that employment and export gains are concentrated among foreign-owned manufacturers, particularly Chinese-owned firms. A one-standard-deviation increase in exposure raises employment by about 7 percent for Chinese firms, compared to roughly 1 percent for domestic firms.

Overall, this literature suggests that foreign-owned firms, particularly those embedded in global supply chains, account for most of the employment gains associated with trade diversion. We add to this evidence by showing that domestic private firms also respond, but only when they have prior trading experience with foreign partners. We further show that these employment effects are persistent and continue to grow over time, rather than dissipating after the initial shock. Finally, we document systematic differences in the composition of new hires, with employment expansion skewed toward female workers and workers at lower wages, highlighting distributional effects that are obscured in aggregate employment measures.

3.4 Data

We leverage three major data sources for our analysis: firm-level data from 8 waves of the Vietnam Enterprise Survey 2015-2022, trade data from the BACI international trade dabase provided by CEPII (Gaulier and Zignago, 2010), and tariff data compiled by Bown (2021) which documents U.S. trade protection measures implemented during the 2018–2019 trade war. Detailed description of each dataset is provided below.

3.4.1 Vietnam firm-level data

The Vietnam Enterprise Survey (VES) is a nationally representative survey conducted annually by the Vietnam General Statistics Office (GSO). The survey covers all businesses registered as enterprises under Vietnam's Enterprise Law. All state, collective, and foreign businesses must register as an enterprise to legally operate in Vietnam. Private businesses are legally required to register as enterprises if they employ more than 10 workers or operate in multiple locations. Although registration is not mandatory for small private businesses, some still choose to register and are therefore included in the sample. The VES provides detailed information on firm characteristics, including identification codes, ownership types, industry classification, employment, total labor compensation, assets and liabilities, and business performance indicators (e.g., sales and revenue).

Industry code: Industry classification in the VES is reported at the five-digit Vietnam Standard Industrial Classification (VSIC) level. The VES datasets for 2015–2017 use VSIC 2007,[32] while the VES adopts the updated VSIC 2018 classification beginning in 2018. To ensure consistency in industry definitions over time, we harmonize VSIC 2018 to VSIC 2007 using the concordance developed by McCaig et al. (2025). We then convert five-digit VSIC codes to four-digit ISIC Rev. 4 industry codes. Because VSIC 2007 was constructed on the basis of ISIC Rev. 4, the correspondence between four-digit VSIC 2007 and ISIC codes is generally one-to-one. We manually verified the mapping between four-digit VSIC 2007 and ISIC Rev. 4 classifications.[33]

Ownership classification: Following McCaig et al. (2025), we classify firms into three ownership categories: state-owned, foreign-owned, and private. Foreign-owned firms include wholly foreign-owned enterprises as well as joint ventures between foreign and Vietnamese entities. State-owned firms include enterprises with 100% state ownership and joint-stock companies in

[32] VSIC 2007 was developed by the General Statistics Office (GSO) based on the International Standard Industrial Classification Revision 4 (ISIC Rev. 4) and the draft ASEAN Common Industrial Classification (ACIC), detailed to the three-digit level.

[33] One four-digit VSIC category (0121, Growing of fruits) maps to multiple four-digit ISIC Rev. 4 industries. We therefore use five-digit VSIC detail in this case to assign the correct ISIC classification.

which the state holds a majority ownership share (greater than 50%). Private firms include partnerships, private enterprises, cooperative unions, and joint-stock companies in which domestic private owners hold at least 50% of ownership.

Panel construction: Our primary unit of observation is the firm–province–industry cell. In each wave of the Vietnam Enterprise Survey (VES), firms are identified by a unique 10-digit tax identifier.[34] We construct a consistent identifier across all years by aggregating observations at the firm–province–industry level. Specifically, for each year we define a unit identifier as the combination of (i) the 10-digit firm tax identifier, (ii) the two-digit province code, and (iii) the two-digit ISIC industry code. Observations sharing the same unit identifier within a year are aggregated. Continuous variables such as employment, revenue, value added, and assets are summed within the cell. For categorical variables, including ownership and the four-digit industry classification, we assign the value from the observation with the largest employment.

Balanced panel: A key empirical challenge in the VES is that firm entry and exit are not cleanly observed due to institutional features of enterprise registration in Vietnam. While all state, collective, and foreign firms are required to register, private businesses are only legally required to do so once they exceed 10 employees or operate across multiple locations. As a result, the appearance of a "new" firm in the data may reflect either true entry or a pre-existing informal firm crossing the registration threshold. Similarly, a firm's disappearance from the sample may indicate either genuine exit or temporary non-reporting, particularly among smaller private firms for whom reporting is not required. These features make it difficult to interpret changes in the unbalanced panel as reflecting true firm dynamics. To address this concern, we focus our main analysis on a balanced panel of firms observed continuously over the sample period. This approach ensures that any employment responses we observe reflect within-firm adjustments to trade shocks,

[34] Starting in 2016, the VES also reports a 3-digit extension to the tax identifier. We explored constructing a 13-digit firm identifier by combining the 10-digit tax ID with this extension, but the resulting panel linkage appears inconsistent across years, particularly around 2019. For consistency, we therefore rely on the original 10-digit identifier as well as province and industry codes.

rather than compositional changes arising from endogenous entry, exit, or reporting behavior.[35] Table 3.1 reports summary statistics on firm counts and employment for the balanced panel of manufacturing firms used in our analysis.

Table 3.1. Summary Statistics by Ownership and Year

	2015	2017	2019	2021
Foreign-owned				
Employment	683.5	754.7	765.6	706.5
	(2314)	(2395)	(2396)	(2303)
Number of firms	4,164	4,210	4,198	4,137
Private				
Employment	76.1	76.5	73.9	69.1
	(351)	(353)	(388)	(362)
Number of firms	21,150	21,173	21,211	21,283
State-owned				
Employment	502.9	509.5	431.0	390.9
	(800)	(914)	(573)	(501)
Number of firms	325	256	230	219
Total				
Employment	180.2	192.2	190.4	174.7
	(1015)	(1057)	(1065)	(1011)
Number of firms	25,639	25,639	25,639	25,639

Note: The table reports mean employment by firm type, with standard deviations in parentheses. Data source: Vietnam General Statistics Office & authors' calculations

3.4.2 Trade data

Annual data on Vietnam's imports and exports at HS 6-digit product level for the period 2014-2022 comes from the BACI dataset provided by CEPII (Gaulier and Zignago, 2010). BACI reports bilateral trade flows for 200 countries and over 5000 products (defined at HS 6-digit level based on the 2012 classification revision). The source of BACI data is the United Nations Statistical Division Comtrade database. CEPII applies a series of processing and harmonization steps to

[35] As a robustness check, we also report results using the full unbalanced panel, though these should be interpreted with these considerations in mind.

the original UN Comtrade dataset to improve data quality. For our analysis, we select from BACI all the import and export bilateral trade flows involving Vietnam during the period 2014-2022.

3.4.3 Trade War Tariffs

Data on U.S. import tariffs imposed on China during the 2018-2019 Trade War period come from Bown (2021), who compiled the information from official reports at the HS 10-digit level of product classification. The source data also reports the exact date (i.e., day of the month) when the tariff increase entered into effect. We modify the source data from Bown (2021) to calculate the average annual increase in import tariff. We do so by factoring in the number of months in a calendar year for which a particular tariff increase was in effect. We then aggregate the data from HS 10-digit level to HS 8-digit level using simple averages. At HS 8-digit level, we combine the Trump War tariff increases with data on U.S. MFN import tariffs obtained from the U.S. International Trade Commission (USITC) to calculate the full tariff level imposed by the U.S. on imports from China. Lastly, we aggregate the import tariffs from HS 8-digit level to HS 6-digit level using as weights the share of each HS 8-digit product in U.S.'s imports from China at HS 6-digit level. We use U.S. Census data from Schott (2008) to create the trade weights. We specifically focus on the importance of each HS 8-digit product imported by the U.S. from China within a given HS 6-digit category to capture more precisely the actual impact of the Trade War on Chinese producers, and the expected trade reallocation effect that we should observe in a given HS 6-digit product category.

3.5 Empirical Methodology

3.5.1 Measuring Trade War Exposure at the Industry Level

Since the VES dataset reports firms' primary industry using the ISIC Revision 4 classification, we construct a concordance from HS 6-digit product codes (2012 revision) to ISIC Rev. 4 using UN Statistics Division mappings.[36] Using this concordance, we compute industry-level exposure

[36]We first map HS 6-digit products to the UN CPC Ver. 2.1 classification (5-digit level), and then map CPC to ISIC Rev. 4 industries. Because these mappings are many-to-many, we assign each HS code to a single ISIC industry

to the Trump trade war by aggregating product-level tariffs using China's pre-trade-war shares in U.S. imports as weights:

$$\tau_{jt} = \sum_{p \in j} \frac{\bar{X}_p^{CHN \to US}}{\sum_{p \in j} \bar{X}_p^{CHN \to US}} \cdot \tau_{pt} \tag{3.1}$$

where τ_{jt} denotes the cumulative tariff increase for ISIC Rev. 4 industry j in year t, and τ_{pt} is the cumulative U.S. tariff increase on HS 6-digit product p imported from China in year t (see Section 3.4.3). The weights $\frac{\bar{X}_p^{CHN \to US}}{\sum_{p \in j} \bar{X}_p^{CHN \to US}}$ are constructed using China's exports to the United States averaged over 2015–2017, capturing the pre-trade-war importance of China as a supplier within each industry while smoothing idiosyncratic year-to-year fluctuations in trade flows.[37] This weighting scheme assigns greater exposure to industries where U.S. imports were more reliant on Chinese goods prior to the trade war, and thus where tariff increases are more likely to induce trade diversion toward alternative suppliers such as Vietnam. As a result, the measure generates continuous variation in tariff exposure at the industry level; Figure C.1 presents the corresponding distribution across ISIC Rev. 4 industries.

Because the weights are based on China's pre-trade-war export structure to the United States, this exposure measure is not shaped by Vietnamese firm behavior and instead reflects variation driven by U.S. tariff policy toward China. It therefore provides a plausibly exogenous source of variation in demand shocks faced by Vietnamese firms. More broadly, this construction can be applied to assess trade war spillover effects in other third countries not directly targeted by the policy.

3.5.2 Employment Effects of the U.S.-China Trade War on Vietnamese Firms

To examine the impact of trade war–induced tariff increases on firm-level employment in Vietnam, we estimate the following specification:

based on the most frequent correspondence in the combined mapping, and manually verify consistency with product descriptions.

[37] We follow Fajgelbaum et al. (2020) and Utar (2026)'s approach of using multiple pre-period years to construct trade shares

$$Y_{ijt} = \alpha + \beta \ln(1 + \tau_{jt}) + \gamma C_{jt} + \theta_i + \mu_t + \varepsilon_{ijt} \tag{3.2}$$

where Y_{ijt} denotes the log of end-of-year employment for firm i in industry j at time t. τ_{jt} is the industry-level measure of cumulative tariff increases described in Section 3.5.1. C_{jt} is a vector of industry-specific, time-varying controls, including the average MFN tariff applied to Vietnamese exports by the rest of the world and the share of HS6-level Vietnamese exports to the United States within each ISIC Rev. 4 industry that were not subject to U.S.–China tariff changes, interacted with a full set of year indicators. The average MFN tariff controls for overall market access conditions faced by Vietnamese exporters, absorbing industry-level variation in export competitiveness unrelated to the trade war. The share of untargeted HS 6-digit Vietnamese exports to the U.S. within each industry, interacted with year indicators, flexibly accounts for the possibility that industries with a higher share of exports unaffected by the trade war may have been on differential employment trajectories over time for reasons independent of the tariff escalation. θ_i and μ_t denote firm and year fixed effects, respectively. Firm fixed effects absorb time-invariant unobserved heterogeneity across firms, such as ownership structure, size, or productivity, while year fixed effects capture common macroeconomic shocks and aggregate trends affecting all firms, including those associated with the COVID-19 period. Standard errors are clustered at the industry and province levels.

The coefficient of interest, β, captures the within-firm employment response to changes in industry-level tariff exposure over time. Identification comes from differential exposure across industries and over time to U.S. tariff increases on Chinese goods, driven by the phased escalation of tariffs across product categories. Because tariffs were imposed in successive rounds targeting different sets of goods, industries vary meaningfully in both the timing and magnitude of their exposure. Since τ_{jt} is constructed using pre-trade-war Chinese export shares to the United States, this variation is driven by U.S. tariff policy toward China rather than by contemporaneous conditions in Vietnam. A key identifying assumption is that, absent the trade war, firms in industries with

different levels of exposure would have followed parallel employment trends. The inclusion of industry-level controls and flexible time interactions helps address concerns about differential pre-trends correlated with exposure. At the end of Section 3.6.1, we provide event-study evidence to confirm the validity of this assumption.

3.6 Empirical Results

3.6.1 Employment Effect

Table 3.2. Effect of Trade War Tariffs on Firm Employment

	Full Balanced Panel		Pre-TW Traders Only	
	(1)	(2)	(3)	(4)
Trade War Tariff	0.1716	0.3085*	0.5714*	0.8336***
	(0.1439)	(0.1774)	(0.2804)	(0.3031)
Controls		✓		✓
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	205,123	190,708	69,655	67,496
R^2	0.911	0.913	0.907	0.907
Within R^2	0.00012	0.00051	0.00137	0.00337
SE clustering	isic4 & province			

Notes: The dependent variable is log end-of-year employment. Columns (1) and (3) report specifications without additional controls, including only firm and year fixed effects. Columns (2) and (4) add controls for the average MFN tariff applied to Vietnamese exports and the share of HS6-level Vietnamese exports to the U.S. within each ISIC 4 industry that were not subject to U.S.–China tariff changes, interacted with a full set of year indicators. Columns (1)–(2) use the full balanced panel of manufacturing firms, while columns (3)–(4) restrict the sample to firms with pre-trade war trade experience. Standard errors are clustered at the industry (ISIC4) and province levels. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 3.2 reports estimates of the effect of trade war tariff increases on firm-level employment in Vietnam. Columns (1) and (2) use the full balanced panel of manufacturing firms, while columns (3) and (4) restrict the sample to firms with pre–trade war trading experience.[38] Columns (1) and (3) present baseline specifications with firm and year fixed effects only, while columns (2)

[38]The Vietnam Enterprise Survey (VES) collects information on whether firms engage in trade with foreign partners in the 2015 and 2017 waves, but not in 2016. We define pre–trade war trading experience as reporting such activity in either 2015 or 2017.

and (4) additionally include controls for MFN tariffs applied to Vietnamese exports, as well as the share of untargeted HS6 exports within each industry interacted with year indicators. Across all specifications, the coefficient on the trade war tariff is positive, and becomes statistically significant once controls are included. This pattern suggests that accounting for time-varying, industry-specific trends is important for isolating the trade diversion effect of the trade war. In particular, the inclusion of controls absorbs confounding variation that may otherwise attenuate the estimated impact of trade war exposure.

Focusing on the preferred specifications with controls, the coefficient in column (2) implies that moving from zero to the mean level of tariff exposure (0.15 ln points) is associated with a 4.7% increase in end-of-year employment. The effect is considerably larger among firms with pre–trade war trading experience (in column (4)), where the same change in exposure corresponds to a 12.6% increase in employment. These results suggest that while the trade diversion channel generated modest average employment gains across the broader manufacturing sector, the gains were substantially more pronounced among firms already engaged in international markets prior to the trade war - precisely the firms best positioned to respond to the demand redirected away from Chinese exporters.

By Ownership

To examine whether the employment effects were distributed evenly across ownership types, we estimate our preferred specification with controls separately for subsamples of Vietnamese manufacturing firms. Table 3.3 reports the resulting estimates disaggregated by ownership. A clear pattern emerges: the positive employment response to trade war exposure is concentrated among foreign-invested and state-owned firms. In the full balanced panel, the estimated coefficient for foreign-owned firms is 1.06 and statistically significant at the 1% level, while state-owned firms exhibit a coefficient of 0.99, significant at the 10% level. Evaluated at the mean tariff exposure of 0.15 log points, these estimates imply employment gains of approximately 16.0% and 14.8% for foreign and state-owned firms, respectively. In contrast, private domestic firms show a small and

Table 3.3. Effect of Trade War Tariffs on Firm Employment by Ownership

	Full Balanced Panel			Pre-TW Traders Only		
	Private	State	Foreign	Private	State	Foreign
Trade War Tariff	0.1474	0.9848*	1.0640***	0.5512**	0.7761	1.0980***
	(0.1904)	(0.5566)	(0.3627)	(0.2308)	(0.4949)	(0.3657)
Controls	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Observations	156,142	1,851	32,463	35,014	1,262	31,033
R^2	0.874	0.961	0.932	0.874	0.952	0.932
Within R^2	0.00021	0.0176	0.00622	0.00229	0.01366	0.00666
SE clustering	isic4 & province					

Notes: The dependent variable is log end-of-year employment. Columns (1)–(3) use the full balanced panel of manufacturing firms, while columns (4)–(6) restrict the sample to firms with pre-trade war export experience. Controls include the average MFN tariff applied to Vietnamese exports and the share of HS6-level Vietnamese exports to the U.S. within each ISIC 4 industry that were not subject to U.S.–China tariff changes, interacted with a full set of year indicators. Standard errors are clustered at the industry (ISIC4) and province levels. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

statistically insignificant coefficient of 0.15.

Restricting the sample to firms with pre–trade war trading experience reveals an additional dimension of heterogeneity. In this subsample, private domestic firms exhibit a positive and statistically significant response, with a coefficient of 0.55, while the estimate for state-owned firms becomes statistically insignificant. This suggests that private firms benefit from trade diversion only when they are already integrated into global trade networks prior to the shock. Evaluated at the mean tariff exposure (0.15 log points), this implies an employment gain of approximately 8.3% for private domestic firms with prior trading experience. The employment response among foreign firms remains large and highly significant, with a magnitude comparable to that in the full sample.

Overall, these results are consistent with the broader literature on trade diversion and firm heterogeneity. Foreign-owned firms , typically larger, more productive, and more deeply embedded in global value chains (Helpman et al., 2004; Antràs and Yeaple, 2014), are well posi-

tioned to rapidly absorb redirected demand. In the Vietnamese context, state-owned enterprises (SOEs) have been documented to enjoy preferential access to credit, regulatory advantages, or other forms of political support (Baccini et al., 2019), which might have enabled them to expand production to meet diverted demand from the trade war as well. By contrast, private domestic firms lacking both the global embeddedness of foreign-owned enterprises and the institutional privileges of SOEs, appear able to capitalize on the trade diversion channel only when they had accumulated prior exposure to international markets. Experience with the global market signals pre-existing export relationships, production capacity, or logistics networks that likely enabled these firms to scale up employment in response to new opportunities arising from the U.S.–China trade war.

Robustness of Results: As a robustness check, we replicate our analysis using the full unbalanced panel of firms. While the balanced panel isolates within-firm adjustments among surviving firms, the unbalanced panel additionally captures entry and exit dynamics. Results are reported in Table C.2. Reassuringly, the results from the unbalanced panel are consistent with our baseline findings. In column (1), which uses the full sample of firms, the coefficient on trade war exposure remains positive, although it is not precisely estimated. The effect becomes substantially larger and statistically significant when restricting the sample to firms with pre-trade-war trading experience, as shown in column (2). The estimated coefficient of 1.16 implies a sizable increase in employment among these firms, consistent with the interpretation that firms already integrated into international markets were best positioned to capitalize on trade diversion opportunities created by the U.S.–China trade war. Importantly, the stronger magnitude in the unbalanced panel suggests that extensive-margin adjustments, such as entry of new firms or expansion of previously small firms into the sample, may amplify the overall employment response. At the same time, the qualitative similarity of the results across balanced and unbalanced specifications indicates that our main conclusions are not driven by selection on firm survival, but instead reflect a robust pattern of employment growth in more exposed sectors.

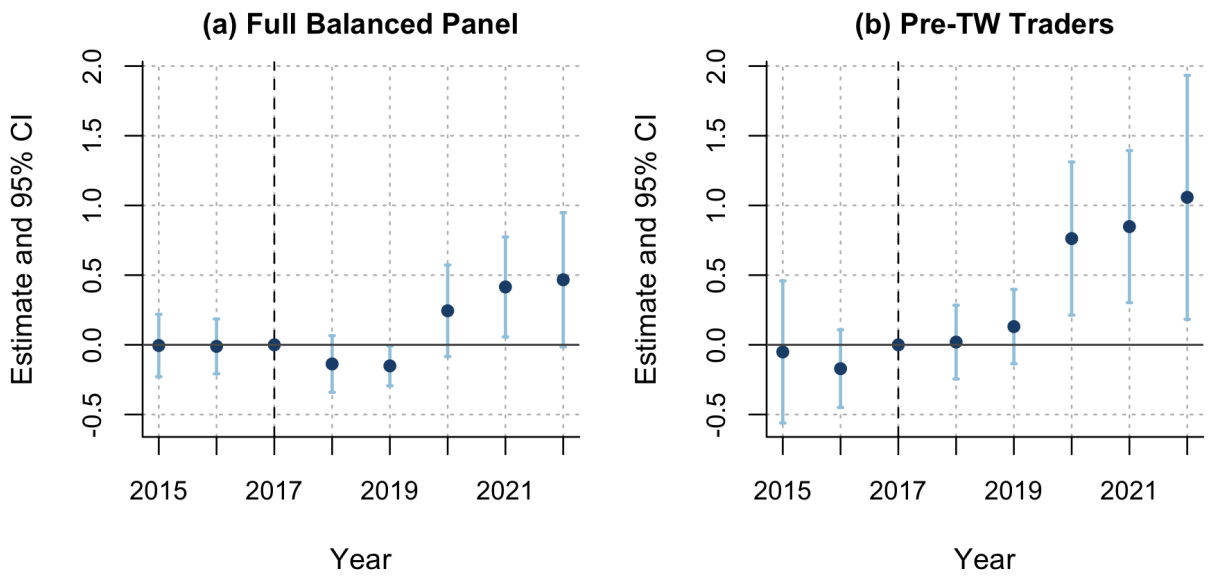


Figure 3.7. Event-study estimates of trade war exposure on firm employment

Event Study:

Figure 3.7 presents event-study estimates of the effect of trade war tariff exposure on firm employment, using 2017 as the base year. Panel (a) shows results for the full balanced panel, while Panel (b) restricts the sample to firms with pre-trade-war trading experience. Before turning to the results, it is useful to clarify a feature of the exposure measure. In the event-study specification, tariff exposure τ_{jt} is time-invariant, capturing the cumulative (maximum) increase in tariffs at the industry level over the course of the trade war, rather than the time-varying exposure used in the main specification. As a result, identification in this setting is primarily cross-sectional, driven by differences across industries in the magnitude of their ultimate exposure. The event-study coefficients therefore trace the dynamic evolution of employment in more versus less exposed industries relative to the base year.

The pre-treatment coefficients in both panels are small, statistically indistinguishable from zero, and exhibit no discernible trend, supporting the parallel trends assumption. In the full sample, the estimated effects for 2018 and 2019 are slightly negative, though imprecisely estimated. This likely reflects that the exposure measure captures the full cumulative tariff increase, whereas realized tariff changes in the early years were smaller, resulting in attenuated short-run effects. In

contrast to muted response in the full sample, firms with pre-trade-war trading experience exhibit a stronger and more immediate response, with economically large and statistically significant increases in employment beginning around 2020 and persisting through the end of the sample period. The delayed but sustained positive effects are consistent with gradual adjustment along the intensive margin, as firms scale up production to meet increased export demand following the reconfiguration of global supply chains.

3.6.2 Worker Composition and Wage Effect

The baseline results establish that Vietnamese manufacturing firms in industries more exposed to trade war tariff increases expanded employment. A natural follow-up question is what kinds of workers were hired in the process. To shed light on this, we re-estimate equation (2) using two alternative outcome variables: the log of female employment and the log of average wages, defined as total compensation divided by total employment[39]. The results are reported in Tables 3.4 and Table 3.5, respectively.

Female Employment

Table 3.4 reports the estimates for female employment. The coefficient on the trade war tariff is positive and statistically significant across all four specifications. Columns (1) and (3) show that firms more exposed to trade war tariffs expanded their female workforce in both the full sample and among pre-trade war traders, with the effect nearly doubling in magnitude for the latter group. Importantly, columns (2) and (4) show that the female employment *share* also increased significantly, indicating that the overall employment expansion documented in the previous subsection was not gender-neutral: firms scaling up in response to increased export demand disproportionately drew on female workers. The result echoes findings from the broader literature linking export-oriented manufacturing expansions in low-income countries to increased female

[39]We are also interested in the formality margin, similar to Nguyen and Lim (2023), who use the total number of workers with social security benefits as the outcome variable. However, from 2020 onwards, the VES no longer reports this measure and instead provides the total amount paid to the social insurance system.

Table 3.4. Effect on Female Employment

	(1)	(2)	(3)	(4)
	Full Sample		Pre-TW Traders Only	
	log(female_emp)	log(female_share)	log(female_emp)	log(female_share)
Trade War Tariffs	0.6767***	0.4574***	1.0380***	0.3400***
	(0.2542)	(0.1417)	(0.3057)	(0.1193)
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	176,547	176,547	66,403	66,403
R^2	0.923	0.726	0.923	0.835
Within R^2	0.00151	0.00181	0.00414	0.00157
SE clustering	isic4 & province			

Notes: Columns (1)–(2) use the full sample; columns (3)–(4) restrict to firms with pre-trade war trading experience. Columns (1) and (3) report effects on female employment while columns (2) and (4) report effects on the female employment share. Female Share is defined as female employment divided by total employment.$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

employment[40]. Further disaggregation by ownership shows that these gains are driven primarily by private and foreign firms (see Table C.3 in the Appendix) while state-owned firms exhibit no comparable gains and, if anything, show negative but imprecisely estimated coefficients.

Wage Effect

Table 3.5 turns to the wage effects of trade war exposure. Average wages are measured as total compensation divided by end-of-year employment. In the full balanced panel, the coefficient on the trade war tariff is statistically indistinguishable from zero; however, restricting the sample to firms with pre-trade war trading experience reveals a negative and statistically significant effect on average wages[41]. We interpret this finding as possibly reflecting a compositional shift in the workforce rather than a deterioration in pay for incumbent workers. Firms rapidly scaling up pro-

[40]Ederington et al. (2009) document that exporters increased their employment of female workers following trade liberalization in Colombia starting in 1985. In the context of Vietnam, Sakakibara (2023) shows that the U.S.–Vietnam Bilateral Trade Agreement (BTA) led to an expansion in female employment in manufacturing, particularly in wearing apparel. Similarly, Heath and Mobarak (2015) find that the growth of the ready-made garments industry significantly increased labor force participation among Bangladeshi women.

[41]Since year fixed effects absorb aggregate inflation common to all firms, deflating total compensation is not required. We nevertheless deflate using the CPI and obtain identical results.

Table 3.5. Effect on Average Wages

	(1)	(2)
Dependent variable: log(total_comp/emp_end)		
Trade War Tariff	0.0014	-0.2514*
	(0.1210)	(0.1168)
Firm FE	✓	✓
Year FE	✓	✓
Controls	✓	✓
Observations	160,712	63,790
R^2	0.987	0.993
Within R^2	0.00082	0.00031
SE clustering	isic4 & province	

Notes: Column (1) uses the full sample; column (2) restricts to firms with pre-trade war exposure. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

duction to meet increased export demand likely hired new workers at the margin who were less experienced or previously employed in lower-productivity informal jobs, mechanically pulling down average wages. Disaggregating by ownership type (Table C.4) provides further support for this interpretation. The negative average wage effect among pre-trade war traders is concentrated among private domestic and state-owned firms, while foreign-owned firms show no statistically significant wage response. This pattern is consistent with foreign firms hiring at wage levels more comparable to their existing workforce or expanding along margins other than low-wage production workers. Private domestic and state-owned firms, by contrast, appear to have absorbed workers at the lower end of the wage distribution as they scaled up production in response to new export demand.

3.7 Conclusion

This paper examines how the 2018–2019 U.S.–China trade war affected Vietnamese manufacturing firms, exploiting cross-industry variation in tariff exposure to identify the employment effect of trade diversion. Vietnam entered this period well positioned to absorb redirected demand: it had established trade relationships with the U.S., deep integration into Southeast Asian supply

chains, and a manufacturing base already familiar with Chinese production processes. Our stylized facts confirm that Vietnam's exports to the U.S. grew sharply after 2018, and that this growth was concentrated precisely in the HS 6-digit product categories most heavily targeted by Trump tariffs.

Our firm-level analysis yields three main findings. First, firms in industries more exposed to trade war tariffs saw meaningful employment gains. Moving from zero to the mean level of tariff exposure is associated with a roughly 5% increase in employment for the average manufacturing firm, rising to about 13% among firms with prior international trading experience. This heterogeneity is intuitive: firms already embedded in global trade networks with established export relationships, logistics capacity, and production knowledge were better positioned to scale up rapidly in response to the demand shock created by Chinese exporters' exit from the U.S. market. Our analysis also sheds light on the dynamics of adjustment. Event-study estimates indicate that employment gains materialize with a lag, becoming economically and statistically significant only after 2020.

Second, the employment gains were not evenly distributed across ownership types. Foreign-owned firms exhibited the largest and most precisely estimated responses; state-owned enterprises also expanded employment in response to trade war exposure. Private domestic firms, by contrast, benefited only when they had accumulated prior trading experience, underscoring that access to international markets is a precondition for capitalizing on trade diversion opportunities.

Third, the composition of new employment matters. Firms expanding in response to tariff exposure disproportionately hired female workers, increasing both the level and the share of female employment. This finding aligns with a broader literature linking export-oriented manufacturing growth in developing countries to increased female labor force participation. At the same time, average wages declined among firms with pre-trade war trading experience, particularly in private and state-owned enterprises. We interpret this as a compositional effect: firms rapidly scaling up production to meet new export demand drew on workers at the lower end of the wage and skill

distribution, pulling down average compensation even as total employment expanded.

Taken together, these results present a nuanced picture of how trade diversion shapes labor markets in a bystander economy. While the U.S.–China trade war generated meaningful job creation in Vietnam's manufacturing sector, the gains were concentrated among firms and workers already connected to global markets, and the new jobs created appear to have been on the lower end of the wage distribution. The results underscore the importance of firm capabilities and pre-existing trade linkages in mediating the effects of global shocks, and suggest that policies aimed at improving firm access to export markets and upgrading production capacity could enhance the ability of domestic firms to capitalize on similar opportunities in the future.

Chapter IV: Assessing Bias in Hypothetical Ability Ratings

Italo Lopez Garcia

University of Southern California and IZA

Nicole Maestas

Harvard University and NBER

Kathleen Mullen

University of Oregon, NBER and IZA

Ngan Tran

University of Oregon

Abstract

We examine whether ability perceptions are biased along the dimensions of gender, race, and ethnicity using data from the 2024 American Work Capacity and Ability Survey. Respondents rate the abilities of hypothetical individuals whose performance and demographic identity are independently randomized, allowing us to isolate the causal effect of perceived identity on ability assessments across four cognitive domains. We find little systematic evidence of bias: ability ratings respond strongly to observed performance, and this relationship is highly similar across demographic groups. Results are robust to tests for intersectional bias, in-group favoritism, and heterogeneity by respondent self-assessment accuracy. These findings suggest that when performance signals are clear and comparable, evaluators rely primarily on observed outcomes rather than demographic cues.

4.1 Introduction

Despite progress in recent decades, wage inequalities persist across gender, race and ethnicity. According to the U.S. Department of Labor, in 2023, women working full-time year-round were paid 84% of what men were paid on average (George and Livingston, 2024). Additionally, among full-time U.S. workers, Black employees earned 76% and Hispanic/Latino employees earned 73% of the average weekly earnings of White workers in 2017-2019 (U.S. Department of Labor, 2024). Occupation and industry differences have been identified as among the most significant contributors to wage disparities (Blau and Kahn, 2017; Grodsky and Pager, 2001)[42].

Occupations group jobs that involve similar tasks performed in exchange for wages. These tasks require specific knowledge, skills, and abilities. If hiring managers or supervisors systematically misperceive the abilities of certain groups, such distortions can affect hiring, promotion, and task assignment decisions, ultimately contributing to occupational segregation and wage inequality. In this way, biased beliefs about ability may serve as a key mechanism linking individual-level perceptions to aggregate labor market disparities.

Using data from the 2024 American Work Capacity and Ability Survey, we examine whether ability ratings assigned to hypothetical individuals vary systematically with gender, race, or ethnicity, holding observed performance constant. Our identification strategy leverages the randomized assignment of identity-signaling names across otherwise identical profiles, allowing us to isolate the causal effect of perceived demographic identity on ability assessments across four cognitive domains: Written Comprehension, Oral Comprehension, Inductive Reasoning, and Number Facility. We have developed an analysis plan that outlines our methodologies in advance of data analysis, serving as a pre-specified record of our intended approach (see embargoed preregistration "Assessment of Hypothetical Ability Ratings," submitted to OSF on March 18, 2025:

[42] Blau and Kahn (2017) found that differences in human capital (i.e. education and experience) explained 8% of the gender wage gap in 2010 while industry and occupation accounted for 51% of the disparity. When it comes to race, Grodsky and Pager (2001) found that the majority of the racial gap in earnings can be explained by individual differences in human capital, region and marital status (55 to 60 percent), whereas industry and occupation accounted for 20%

https://osf.io/6vqfw).[43]

Our main finding is that there is little systematic evidence of bias. Across all domains, ability ratings respond strongly and monotonically to observed performance, and this relationship is highly similar across gender, racial, and ethnic groups. Neither the level of ratings nor the translation of performance into perceived ability differs significantly between male and female profiles, or between White, Black, and Hispanic profiles, once performance is held constant. Joint hypothesis tests fail to reject the null of no bias in any domain, and pooled specifications confirm the robustness of these null results to the inclusion of respondent fixed effects.

We also find little evidence that these null effects mask heterogeneity along other dimensions. Tests for intersectional bias show that gender and race do not combine to produce systematic disparities beyond their individual main effects, with only one localized exception in the oral domain for Black women. Tests for in-group favoritism yield no evidence that respondents evaluate same-gender or same-race profiles more generously. Heterogeneity in respondent self-assessment accuracy—including the pronounced male overconfidence documented in prior work—does not translate into systematic amplification of gender bias in evaluations of others.

We contribute to several strands of the literature. First, we relate to a large body of work documenting discrimination in labor market outcomes, particularly audit studies that vary identity signals to measure differential treatment (Bertrand and Mullainathan, 2004; Kline et al., 2022). While these studies provide compelling evidence of disparities at the point of hire, they measure reduced-form outcomes, such as callback rates, that do not reveal the beliefs and perceptions underlying evaluators' decisions. In particular, they cannot isolate whether disparities arise because evaluators perceive minority or female candidates as less able, less suitable for the role, or less likeable, nor can they distinguish taste-based from statistical discrimination. More broadly, Bohren et al. (2023) demonstrate that standard audit methods cannot distinguish taste-based from

[43]The pre-analysis plan included five cognitive domains, including memorization (immediate and delayed recall). The present analysis focuses on four domains and excludes memorization measures.

inaccurate statistical discrimination, and propose eliciting evaluators' subjective beliefs directly as a solution. Our approach implements this by directly eliciting ability perceptions in a controlled setting where performance is held constant, allowing us to test whether demographic identity shapes the translation of observed performance into subjective assessments of ability.

Second, our paper connects to models of statistical discrimination (Phelps, 1972; Arrow, 1973), in which employers use group identity as a proxy for unobservable productivity when individual signals are noisy or unavailable. A key prediction of these models is that discrimination should attenuate, or disappear entirely, when precise performance information is provided, since the informational friction that motivates the use of group-level stereotypes is resolved. Our experimental design directly tests this prediction by holding performance constant and making it fully observable and comparable across individuals. Our null results are consistent with this mechanism: when performance signals are explicit, quantifiable, and directly comparable across individuals, respondents do not appear to use demographic identity as a supplementary signal of ability. This suggests that biased ability perceptions may be most consequential in real-world settings where performance signals are noisier, less comparable, or subject to interpretation.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data. Section 4 outlines the empirical strategy, following the pre-analysis plan. Section 5 presents the results, and section 6 concludes.

4.2 Related Literature

This study builds on three related strands of research: correspondence studies that document discrimination in labor market outcomes, laboratory and field experiments that examine the mechanisms underlying biased evaluations, and a growing literature on how performance is interpreted and rewarded across groups.

A large body of correspondence studies has provided compelling evidence of discrimination in hiring based on race and gender. Neumark et al. (1996) conducted an audit study in Philadel-

phia and found that women were significantly less likely than men to receive interviews or job offers at high-end restaurants, suggesting both employer and customer-driven discrimination. Bertrand and Mullainathan (2004) demonstrated striking racial bias by sending fictitious resumes to job postings in Boston and Chicago, showing that resumes with White-sounding names received 50% more callbacks than those with African-American names, with less responsiveness to resume quality for African-American applicants. More recently, Kline et al. (2022) extended this approach at scale by submitting over 83,000 fictitious applications to 108 major U.S. employers. They found that distinctively Black names reduced employer contact rates by 2.1 percentage points, with discrimination concentrated among a small set of firms. While average gender gaps in callbacks were minimal, there was substantial firm-level heterogeneity, with some employers favoring men and others women. While these studies provide compelling evidence of disparities in hiring outcomes, they do not directly observe the beliefs that underlie these decisions.

A closely related literature examines how identical performance is evaluated across individuals. Reuben et al. (2014) show that evaluators are less likely to select women for math-related tasks, even when performance is comparable, reflecting biased beliefs about ability. Sarsons (2017) find that women receive less credit for coauthored work in academia, suggesting that contributions are differentially attributed across gender. Anne (2017) documents systematic gender bias in student evaluations of teaching, where male instructors receive higher ratings despite no differences in learning outcomes. These studies highlight that disparities can arise not only from access to opportunities, but also from how performance is interpreted and rewarded.

Complementing this evidence, laboratory and field experiments have explored the mechanisms underlying differences in evaluation and behavior. Gneezy et al. (2003) show that men outperform women in competitive environments, even when performance is similar under non-competitive compensation schemes. Correll et al. (2007) document status-based discrimination, finding that mothers are perceived as less competent and offered lower starting salaries, while fathers are not penalized. Niederle and Vesterlund (2007) demonstrate that men are significantly

more likely than women to enter competitive environments, driven in part by differences in confidence and preferences for competition.

By using randomized vignettes embedded in a survey, our study leverages the causal identification strengths of experimental design to contribute to this literature. In contrast to correspondence studies that focus on binary outcomes, and evaluation studies where performance is often imperfectly observed, our setting allows us to directly examine how evaluators map clearly defined and comparable performance signals into perceptions of ability across gender and racial/ethnic groups.

4.3 Data

Our dataset comes from the American Work Capacity and Ability Survey (AWCAS), a survey that was administered between June 19 and August 18, 2024, to English-speaking participants in the Understanding America Study (UAS) [44]. The UAS is a nationally representative sample of Americans aged 18 and older who speak English or Spanish and who have agreed to participate in regular online social science surveys. The survey consists of three sections: 1) Self-Reported Abilities, 2) Demographic & Labor Force Status, and 3) Cognitive Performance Tests and Follow-Up Questions [45]. The main purpose of the AWACS survey was to construct a measure of health-related work capacity based on interactions between functional abilities and occupational requirements and assess its properties; see embargoed preregistration "Assessment of Health-Related Work Capacity to Improve Independence of Older Adults" submitted to OSF on September 18, 2024 (https://osf.io/2cdjm). Section 3 of the survey was intended to assess respondents' objectivity in reporting their ability levels in Section 1 and, if needed, correct the work capacity measure for self-report bias[46].

[44] A pilot survey was administered to 87 participants between March 11 and April 16, 2024. Their responses will be excluded from the analysis sample

[45] For more information about the UAS, see https://uasdata.usc.edu/index.php. The AWACS was survey number 582.

[46] The AWCAS data were under embargo and only available to Lopez-Garcia, Maestas and Mullen when the pre-analysis plan was developed in Winter 2025 by Tran, who did not have access to the data at the time of writing

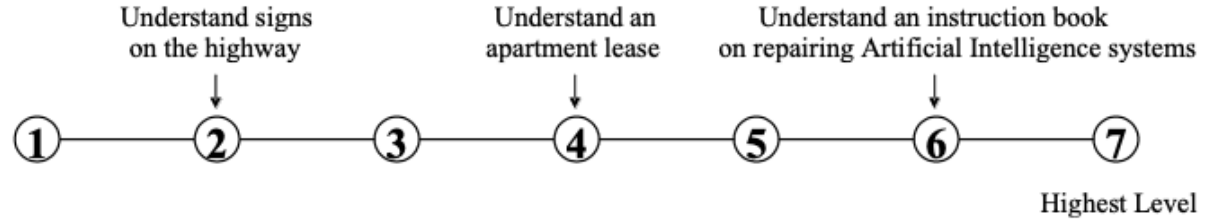
Section 3 of the AWCAS survey (Cognitive Performance Tests and Follow-Up Questions) serves as the primary data source for the proposed analysis. It comprises six subsections testing different cognitive abilities: Written Comprehension, Oral Comprehension, Memorization (Immediate Recall), Number Facility, Memorization (Delayed Recall) and Inductive Reasoning. Each subsection begins with a series of questions designed to assess participants' competency within that cognitive domain. After completing the performance questions, participants were asked to provide an objective estimate of the number of answers they got correctly (where the maximum number possible varied by domain). Next, they were asked to translate the number of correct answers achieved by a hypothetical person to the corresponding level of ability on the same scale (from 0 to 7, where 0 indicates 'not relevant' and 7 represents the highest level of ability) that they used to rate their own abilities in Section 1 (Self-Reported Abilities). Each ability scale also has a corresponding set of anchors that provide examples of activities that correspond to specific ability levels. The scale and level anchors were taken from the ability questionnaire used to describe ability requirements for U.S. occupations collected and posted by the U.S. Bureau of Labor Statistics in its Occupational Information Network (O*NET) database.

A key element of the research design that allows us to detect biases in perceptions of abilities along gender and/or racial/ethnic lines is the randomization of the hypothetical person's name and the number of correct answers for each translational question. The hypothetical identities used in the experiment represent two genders - female and male - and three racial/ethnic groups - Black, Hispanic, and White. The names were selected from Tzioumis (2018) and are those that were both common and predominantly associated with a single racial/ethnic group. See Table E.1 for the list of names by gender and race/ethnicity for the hypothetical person.

An example of the translational question, for Written Comprehension, is as follows:

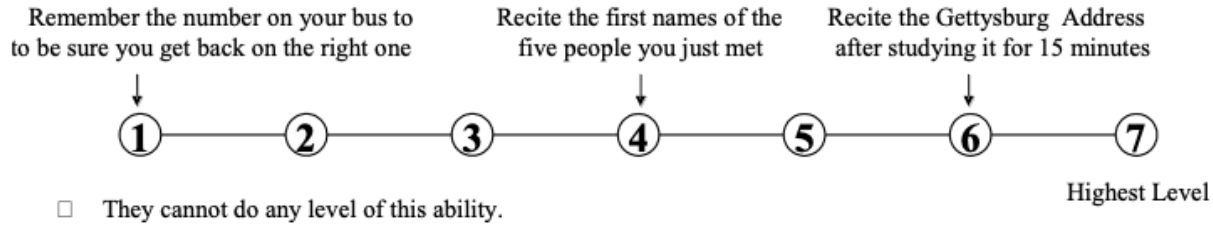
Suppose that [Name], a hypothetical person, got [number] out of 3 questions right in the written comprehension test. Please rate [Name]'s level of WRITTEN COMPREHENSION by clicking the circled number that best matches their level of ability.

The translational questions for Oral Comprehension, Number Facility, and Inductive Rea-



soning have the same structure. The exception is Memorization, which asks respondents to rate hypothetical individuals' abilities given performance information for both immediate and delayed recall, where the number correct for delayed recall must be less than or equal to the number correct for immediate recall:

Suppose that [Name], a hypothetical person, was able to recall [number1] out of 15 words correctly immediately after they were presented and [number2] out of 15 words correctly a little while after they were presented. Please rate [Name]'s level of MEMORIZATION by clicking the circled number that best matches their level of ability.



4.4 Empirical Strategy

We estimate gender/racial bias in ability ratings by comparing the reported ability levels assigned to hypothetical individuals with names indicating different identities, while holding constant their hypothetical performance (i.e., the number of correct answers on the cognitive test). This allows us to isolate the effect of identity on the perceived ability of others. We begin by separately detecting gender-based and race-based biases. We start off with fully flexible specifications allowing for heterogeneous bias based on the number of randomized correct answers, followed by statistical tests to determine whether reducing to more parsimonious specifications is plausible. We

also consider different transformations of the outcome variable and alternative specifications to address potential non-linearity and test for improved model fit. Furthermore, given the variation in the number of questions across categories, normalization will be performed to enable running pooled regressions across categories using respondents' fixed effects. Next, we investigate whether the intersectionality of gender and race affects the presence and extent of biases. Finally, we test whether respondents' characteristics impact biases, either through alignment with the hypothetical persons' demographic traits or through inaccurate estimation of their own abilities. We provide a detailed explanation of our empirical strategy in the following subsections.

4.4.1 Detecting gender-based and race-based biases separately

Gender-based biases

We start with specifications that allow for the relationship between performance test scores and ability ratings to be as flexible as possible, without imposing any linearity or other restrictions. We will estimate the following series of equations:

$$y_{ij} = \sum_{k=0}^{\max_j} \delta_{jk}^{G} 1(\text{CorrectAns}_{ij} = k) + \sum_{k=0}^{\max_j} \beta_{jk}^{F} 1(\text{CorrectAns}_{ij} = k) \cdot Female_{ij} + \varepsilon_{ij} \tag{4.1}$$

where i indexes respondent and j indexes category of ability, y_{ij} is the ability rating provided by individual i for ability j, $1(\text{CorrectAns}_{ij} = k)$ for $k = 0, 1, ..., max_j$ is a series of indicators that take value 1 if the randomized number of correct answers was 0, 1,..., max_j, respectively, and $Female_i j$ is an indicator if the randomized putative gender of the hypothetical person is female. If the coefficients for β_{jk}^{F} for $k = 0, 1, ..., max_j$ are negative, it provides evidence that respondents' ratings are biased against women for ability j.

We will conduct several F-tests to assess whether reducing to more parsimonious specifications is feasible. For instance, if all β^F coefficients in equation (1) are not significantly different from each other, the randomized number of correct answers can be treated as a count variable

rather than a series of indicators. Similarly, we will conduct an F test to determine whether certain restrictions on the δ^G coefficients in equation (1) are warranted; e.g., if linearity is not rejected, the randomized number of correct answers can also be simplified into a count variable.

Race-based biases

We begin with a specification that accounts for heterogeneity in racial biases based on the number of correct answers and estimate the following equation:

$$y_{ij} = \sum_{k=0}^{\max_j} \delta_k^R \, 1(\text{CorrectAns}_{ij} = k) + \sum_{k=0}^{\max_j} \beta_k^B \, 1(\text{CorrectAns}_{ij} = k) \cdot Black_j + \sum_{k=0}^{\max_j} \beta_k^H \, 1(\text{CorrectAns}_{ij} = k) \cdot Hisp_j + \varepsilon_{ij} \quad (4.2)$$

where i indexes respondent and j indexes category of ability, y_{ij} is the ability rating provided by individual i for ability j, $1(\text{CorrectAns}_{ij} = k)$ for $k = 0, 1, ..., max_j$ is a series of indicators that take value 1 if the randomized number of correct answers was 0, 1,..., max_j, respectively. $Black_j$ and $Hisp_j$ are indicators that take value 1 if the hypothetical person in question j is Black or Hispanic, respectively. If the coefficients β_{jk}^B for $k = 0, 1, ..., max_j$ are negative, this indicates evidence of biases against Black individuals, and if the coefficients β_{jk}^H for $k = 0, 1, ..., max_j$, it is evidence of biases against Hispanic individuals. We will proceed by performing F-tests to determine whether reducing to more parsimonious specifications, as discussed in Section 4.1.1, is feasible.

Transformations of outcome variables and alternative specifications

In addition to linear models, we will explore whether various transformations of the outcome variable, including logarithmic and quantile normalizations, as well as alternative regression specifications such as exponential and polynomial functional forms, fit the data well enough to reduce dimensionality for the more complex specifications below. This comprehensive approach enables us to identify the most suitable functional form for the outcome variable while preserving

the integrity of the analysis.

Pooled regression across all abilities

In the Performance Tests and Follow-Up Questions, the total number of questions varies across different categories of ability. For instance, Written Comprehension includes 3 questions, Oral Comprehension has 6 questions, and Number Facility comprises 5 questions. Consequently, in order to conduct a pooled regression across all categories, we will first normalize the randomized number of correct answers by dividing it by the total number of questions within each category. This normalization ensures that the outcome variable is comparable across categories, regardless of the differences in question count. Since ability is already rated on the same scale (0 to 7), we will then perform a pooled regression similar to equation (1) to detect gender biases and similar to equation (2) to detect racial biases. In this analysis, the randomized number of correct answers will be replaced with its normalized version, which ranges between 0 and 1. An important advantage of this approach is that it allows us to include respondents' fixed effects, controlling for individual-specific characteristics and improving the precision of our estimates.

4.4.2 Testing intersectionality of gender and race

A large body of literature in Sociology emphasizes the importance of intersectionality in investigating inequalities, highlighting how social identities such as race, gender, and class intersect to shape individuals' experiences (McCall, 2005; Collins and Bilge, 2020). However, relatively few studies in Economics have found empirical evidence consistent with strong intersectional effects. For example, Bertrand and Mullainathan (2004) reported mixed results regarding whether racial discrimination is more pronounced for men or women. Similarly, Nunley et al. (2015) examined the labor market for recent college graduates and found that the black-white male differential is not statistically different from the black-white female differential. Building on this literature, we assess whether the joint effect of gender and race on perceived ability deviates from the sum of their separate effects.

To do so, we estimate a fully interacted specification that allows both baseline ratings and the relationship between performance and perceived ability to vary flexibly across gender, race/ethnicity, and their interactions:

$$
\begin{aligned}
y_{ij} = {} & \beta_0 + \beta_1 CorrectAns_{ij} + \beta_2 Female_j + \beta_3 Black_j + \beta_4 Hisp_j \\
& + \beta_5 (Female_j \times Black_j) + \beta_6 (Female_j \times Hisp_j) \\
& + \beta_7 (CorrectAns_{ij} \times Female_j) + \beta_8 (CorrectAns_{ij} \times Black_j) + \beta_9 (CorrectAns_{ij} \times Hisp_j) \\
& + \beta_{10} (CorrectAns_{ij} \times Female_j \times Black_j) \\
& + \beta_{11} (CorrectAns_{ij} \times Female_j \times Hisp_j) + \varepsilon_{ij},
\end{aligned}
\tag{4.3}
$$

where i indexes respondents and j indexes the hypothetical individual. All variables are defined as in previous sections.

This specification allows for non-additive (intersectional) effects along two dimensions. First, the coefficients β_5 and β_6 capture whether individuals at the intersection of gender and race/ethnicity receive baseline ratings that differ from what would be predicted by the sum of the separate gender and race effects. Second, the coefficients β_{10} and β_{11} capture whether respondents update their beliefs about ability differently for Black women and Hispanic women as performance varies. Under an additive model with no intersectionality, these interaction terms would be equal to zero. Deviations from zero therefore provide evidence of intersectional bias in either baseline perceptions or the responsiveness of ratings to performance.

4.4.3 Testing if respondents' characteristics affect biases

Because the number of correct answers and the putative identities of the hypothetical person are randomized, we have not included respondents' characteristics in our regression equations thus far. However, empirical evidence on factors influencing biases motivates us to investigate whether respondents' characteristics affect biases - either through alignment with the hypothetical persons' demographic traits or through inaccurate estimation of their own abilities.

Matching of gender, race/ethnicity

Motivated by the work of Parsons et al. (2011), who examined pitch calls in Major League Baseball and found that strikes are called less often when the umpire and pitcher do not share the same race or ethnicity, we investigate whether the alignment of respondents' and hypothetical persons' demographic characteristics influences biases. For instance, male respondents might assign lower ability ratings to female-sounding names compared to male ones than female respondents do, or respondents may rate names associated with other races or ethnicities lower than those matching their own.

To test whether gender matching affects ability ratings, we estimate the following equation:

$$\begin{aligned} y_{ij} = {} & \beta_0 + \beta_1 \,\text{CorrectAns}_{ij} + \beta_2 \,\textit{Female}_j + \beta_3 \,\textit{OwnFemale}_i \\ & + \beta_4 \,(\textit{Female}_j \times \textit{OwnFemale}_i) \\ & + \beta_5 \,(\text{CorrectAns}_{ij} \times \textit{Female}_j) + \beta_6 \,(\text{CorrectAns}_{ij} \times \textit{OwnFemale}_i) \\ & + \beta_7 \,(\text{CorrectAns}_{ij} \times \textit{Female}_j \times \textit{OwnFemale}_i) + \varepsilon_{ij}, \end{aligned} \tag{4.4}$$

where i indexes respondents and j indexes category of ability. The variable $\textit{Female}_j$ is an indicator for whether the hypothetical individual is female, and $\textit{OwnFemale}_i$ indicates whether the respondent is female. This specification allows both baseline ratings and the relationship between performance and perceived ability to vary by respondent gender, gender of hypothetical person, and their interaction. The coefficient β_4 captures whether respondents assign different baseline ratings to individuals of the same gender relative to those of the opposite gender. The coefficient β_7 captures whether the marginal effect of performance on perceived ability differs for same-gender evaluations, providing a test of gender-based in-group bias in how respondents translate performance into ratings.

We conduct a similar analysis to examine whether alignment in race or ethnicity between respondents and hypothetical individuals influences ability ratings.

Heterogeneity due to incorrect self-assessment of abilities

In examining the causes of gender differences across various domains, a significant body of literature has explored the role of noncognitive and behavioral differences. For instance, Niederle and Vesterlund (2007), in a controlled laboratory experiment, demonstrated that men tend to exhibit significantly greater overconfidence in their relative performance compared to women. If overconfidence in one's own abilities is correlated with assessments of others' abilities, particularly those who share similar characteristics, ability ratings may reflect biases influenced by the accuracy of self-assessment. Since we collect information on respondents' estimates of the number of questions they answered correctly and can compare this to their actual performance, we can classify respondents as accurately estimating, overestimating, or underestimating their own abilities. We will then investigate whether biases in ability ratings vary along this dimension.

4.5 Results

4.5.1 Descriptive Results

We begin by presenting descriptive evidence on how respondents map observed performance into perceived ability, and whether these mappings differ by the gender or race/ethnicity of the hypothetical individual. Figures 4.1 and 4.2 plot average ability ratings against the randomized number of correct answers across four ability domains. Across all ability categories, ability ratings increase monotonically with the number of correct answers. This pattern holds consistently for all four cognitive domains - Written Comprehension, Oral Comprehension, Number Facility, and Inductive Reasoning, indicating that respondents meaningfully incorporate performance information when forming ability assessments. The relationship appears approximately linear, suggesting that respondents treat increments in correct answers as roughly equally informative across the range of performance.

Figure 4.1 shows that the mapping from performance to perceived ability is highly similar for male- and female-associated names. The two lines largely overlap across all performance levels and ability domains, with only minor deviations at specific points. There is no clear vi-

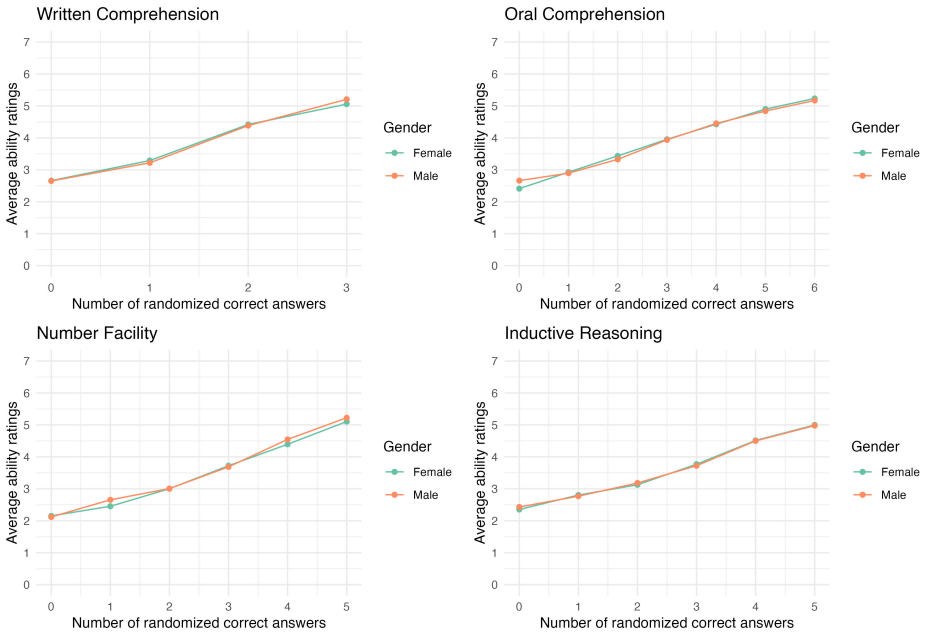


Figure 4.1. Average Ability Ratings by Randomized Performance and Gender

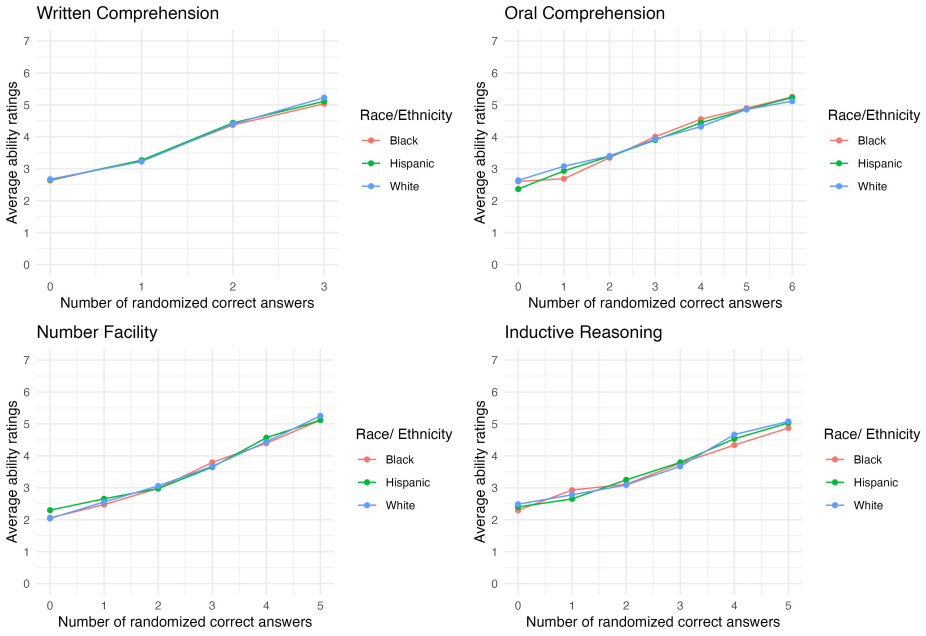


Figure 4.2. Average Ability Ratings by Randomized Performance and Race/ Ethnicity

sual evidence of systematic gender-based differences in average ratings at this descriptive stage. Figure 4.2 presents analogous patterns by race/ethnicity. Ability ratings for Black-, Hispanic-, and White-associated names follow nearly identical trajectories as performance increases. While small differences emerge at certain performance levels, these differences are not consistent across domains and do not suggest a clear ordering across racial or ethnic groups.

Overall, the descriptive evidence indicates that respondents strongly anchor ability ratings to observed performance, with little visual evidence of systematic differences by gender or race/ethnicity. However, these raw patterns may mask more subtle forms of bias that vary across the distribution of performance or across respondents. The descriptive patterns also suggest an approximately linear relationship between performance and perceived ability across all domains. Accordingly, we proceed with linear specifications in the regression analysis that follows, which formally tests for the presence and magnitude of any such biases while exploiting the randomized assignment of identities and performance levels.

4.5.2 Gender-Based Bias

Gender Bias by Performance Level

Table 4.1 reports the estimated β^F_{jk} coefficients from equation (1), capturing the effect of a female-sounding name on ability ratings at each level of the randomized number of correct answers, separately for each ability domain. The results provide little systematic evidence of gender-based bias in ability ratings. Across the four domains, the vast majority of coefficients are small in magnitude and statistically indistinguishable from zero. Where statistically significant effects do appear, they are isolated to specific score levels rather than consistent across the full distribution, which argues against a uniform penalty for female-sounding names.

There are, however, a few isolated instances of statistically significant differences that are worth noting. In Written Comprehension, the coefficient at three correct answers is negative and marginally significant ($\beta^F_3 = -0.150$, $se = 0.078$). In Oral Comprehension, the coefficient at zero

Written Comprehension	**Oral Comprehension**	**Inductive Reasoning**	**Number Facility**
$\beta_0^F = 0.011$ (0.079)	$\beta_0^F = -0.251$ (0.103)**	$\beta_0^F = -0.079$ (0.104)	$\beta_0^F = 0.033$ (0.107)
$\beta_1^F = 0.071$ (0.077)	$\beta_1^F = 0.032$ (0.102)	$\beta_1^F = 0.035$ (0.107)	$\beta_1^F = -0.202$ (0.106)*
$\beta_2^F = 0.036$ (0.078)	$\beta_2^F = 0.108$ (0.102)	$\beta_2^F = -0.056$ (0.106)	$\beta_2^F = -0.004$ (0.110)
$\beta_3^F = -0.150$ (0.078)*	$\beta_3^F = 0.018$ (0.100)	$\beta_3^F = 0.047$ (0.107)	$\beta_3^F = 0.036$ (0.105)
	$\beta_4^F = -0.025$ (0.102)	$\beta_4^F = 0.013$ (0.104)	$\beta_4^F = -0.152$ (0.110)
	$\beta_5^F = 0.061$ (0.103)	$\beta_5^F = 0.018$ (0.101)	$\beta_5^F = -0.120$ (0.105)
	$\beta_6^F = 0.062$ (0.104)		

* significant at 10%, ** at 5%, *** at 1%

Note: Coefficients reported are estimators of β_{jk}^F in equation (1)

Table 4.1. Effects of Female Identity on Perceived Ability

correct answers is negative and significant at the 5 percent level ($\beta_0^F = -0.251$, $se = 0.103$), suggesting that respondents assign meaningfully lower ratings to female-named hypothetical persons who answer no questions correctly relative to their male-named counterparts. In Number Facility, the coefficient at one correct answer is negative and marginally significant ($\beta_1^F = -0.202$, $se = 0.106$). However, none of these effects persist across the full score distribution within their respective domains, and no significant coefficients emerge in Inductive Reasoning at any score level. Overall, the pattern is consistent with the descriptive evidence in Figure 1: male and female rating curves are nearly indistinguishable across most of the score distribution. The sporadic significant coefficients at the low end of the performance distribution are not consistent across domains and do not persist across adjacent score levels, limiting their interpretability as evidence of systematic bias.

To assess whether the perceived ability penalties associated with female identity are jointly significant across performance levels, we test the null hypothesis that all interaction terms are jointly equal to zero, i.e., $H_0 : \beta_0^F = \beta_1^F = \cdots = \beta_{\max_j}^F = 0$. The F-statistics and corresponding p-values are reported below:

- Written Comprehension: $F = 1.19$, $p = 0.31$
- Oral Comprehension: $F = 1.14$, $p = 0.335$

- Inductive Reasoning: $F = 0.20$, $p = 0.98$
- Number Facility: $F = 1.19$, $p = 0.31$

Across all four ability domains, we fail to reject the null hypothesis that the coefficients are jointly equal to zero. This indicates that, conditional on performance, female-associated names are not systematically assigned different ability ratings than male-associated names. These results reinforce the conclusion that there is no robust evidence of gender-based bias in perceived ability ratings.

Pooled Regression

We normalize the number of randomized correct answers by its maximum value within each category to construct a pooled performance measure ranging from 0 to 1, making it comparable across skills. Figure E.1 plots the unconditional distribution of perceived ability ratings by gender. The two distributions closely overlap across the full range of ratings. Both are centered around the middle of the scale, with the highest mass at ratings 4 and 5, and substantially fewer observations at the extremes. While there are minor differences, for instance, female-associated profiles receive slightly more ratings at level 4 and marginally fewer at the upper tail (ratings 6 and 7), these differences are small and do not reflect a systematic shift in the distribution. Overall, the similarity in both the location and dispersion of the distributions suggests that, unconditionally, respondents assign very similar ability ratings to male- and female-associated profiles.

We next estimate the following pooled specification:

$$y_{ij} = \beta_0 + \beta_1 \cdot CorrectAns_{ij} + \beta_2 \cdot Female_{ij} + \beta_3 \cdot (CorrectAns_{ij} \times Female_{ij}) + \alpha_i + \gamma_j + \varepsilon_{ij} \quad (4.5)$$

The results, reported in Table 4.2, are estimated both with and without respondent fixed effects. Across both specifications, the coefficient on the female indicator is small and statistically insignificant, indicating no systematic difference in average ability ratings between female- and male-associated profiles conditional on performance. The interaction between performance and

	(1)	(2)
Dependent variable: Ability Ratings		
Correct	2.950*** (0.0423)	2.963*** (0.0436)
Female	-0.0194 (0.0364)	-0.0080 (0.0369)
Correct × Female	0.0079 (0.0596)	-0.0411 (0.0612)
Ability Fixed Effects	✓	✓
Respondent Fixed Effects		✓
Observations	13,512	13,512
R^2	0.4261	0.6564
Within R^2	0.4209	0.4736

* significant at 10%, ** at 5%, *** at 1%

Table 4.2. Pooled Regression Results with Ability and Respondent Fixed Effects

female identity is also close to zero and statistically insignificant, suggesting that respondents do not differentially reward performance based on gender. That is, the mapping from performance to perceived ability is similar for male- and female-associated profiles. Including respondent fixed effects increases explanatory power but leaves the coefficients of interest essentially unchanged. This indicates that the absence of gender differences is not driven by unobserved heterogeneity across respondents. Overall, the pooled results reinforce the earlier findings: we find no evidence of systematic gender-based bias in perceived ability ratings.

As an additional check, we examine whether gender differences arise at the tails of the rating distribution. We define indicators for receiving a high rating (6–7) and a low rating (0–2), and estimate linear probability models including respondent and ability fixed effects. The results show no statistically significant effect of female identity on the likelihood of receiving either high or low ratings, nor any difference in how performance translates into tail outcomes across genders. Detailed results are reported in Appendix Table E.2.

4.5.3 Race-Based Bias

Race Bias by Performance Level

Table 4.3 reports the estimated coefficients from equation (2), capturing the effect of Black

	Black	Hispanic
Panel A: Written Comprehension		
$k = 0$	−0.040 (0.098)	−0.018 (0.095)
$k = 1$	0.039 (0.095)	0.039 (0.093)
$k = 2$	−0.028 (0.095)	0.037 (0.095)
$k = 3$	−0.191 (0.094)**	−0.116 (0.096)
Panel B: Oral Comprehension		
$k = 0$	−0.034 (0.127)	−0.272 (0.123)**
$k = 1$	−0.390 (0.128)***	−0.146 (0.122)
$k = 2$	−0.057 (0.124)	−0.014 (0.125)
$k = 3$	0.079 (0.121)	−0.024 (0.124)
$k = 4$	0.232 (0.123)*	0.120 (0.125)
$k = 5$	0.041 (0.127)	0.001 (0.124)
$k = 6$	0.134 (0.129)	0.115 (0.125)
Panel C: Inductive Reasoning		
$k = 0$	−0.190 (0.128)	−0.086 (0.126)
$k = 1$	0.149 (0.130)	−0.129 (0.130)
$k = 2$	0.021 (0.131)	0.160 (0.128)
$k = 3$	0.092 (0.130)	0.123 (0.132)
$k = 4$	−0.329 (0.126)***	−0.134 (0.129)
$k = 5$	−0.205 (0.124)*	−0.047 (0.125)
Panel D: Number Facility		
$k = 0$	0.019 (0.131)	0.260 (0.131)**
$k = 1$	−0.082 (0.128)	0.099 (0.130)
$k = 2$	−0.077 (0.134)	−0.083 (0.136)
$k = 3$	0.124 (0.129)	−0.024 (0.128)
$k = 4$	−0.046 (0.132)	0.129 (0.134)
$k = 5$	−0.136 (0.128)	−0.133 (0.130)

Table 4.3. Effects of Race on Perceived Ability

Notes: Coefficients reported are interaction effects relative to White (reference group). Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

and Hispanic identity on ability ratings at each performance level, relative to White candidates. As with the gender results, the large majority of coefficients are small and statistically insignificant, and no domain exhibits a consistent pattern of negative effects across the full score distribution. Among the statistically significant estimates, the most notable pattern emerges in Inductive Reasoning, where candidates with Black-associated names receive lower ratings at $k = 4$ ($\beta_4^B = -0.329$, $se = 0.126$) and $k = 5$ ($\beta_5^B = -0.205$, $se = 0.124$)—two adjacent score levels at the upper end of the performance distribution—suggesting a potential penalty when performance signals are strongest. In Oral Comprehension, candidates with Black-associated names receive lower ratings at $k = 1$ ($\beta_1^B = -0.390$, $se = 0.128$) but higher ratings at $k = 4$ ($\beta_4^B = 0.232$, $se = 0.123$), a sign reversal that argues against a uniform downward bias. For Hispanic candidates, statistically significant effects are limited to Oral Comprehension at $k = 0$ and Number Facility at $k = 0$, the latter being positive, and no domain exhibits a coherent directional pattern. Overall, while the concentrated negative effects for Black-associated names in Inductive Reasoning at higher performance levels are noteworthy, the results suggest that race-based differences in perceived ability are limited in scope, non-monotonic across the performance distribution, and not robust across domains.

Pooled Regression

We begin by examining the unconditional distribution of perceived ability ratings by race/ethnicity. Figure E.2 shows that the distributions for Black-, Hispanic-, and White-associated profiles largely overlap and are centered around the middle of the rating scale. While there are modest differences across groups, these patterns are not pronounced, suggesting broadly similar average ratings across racial identities. We next estimate a pooled specification analogous to equation (5), replacing the gender indicator with indicators for Black- and Hispanic-associated profiles, with White-associated profiles as the omitted reference category:

$$
\begin{aligned}
y_{ij} = \beta_0 &+ \beta_1 \cdot CorrectAns_{ij} + \beta_2 \cdot Black_{ij} + \beta_3 \cdot Hispanic_{ij} \\
&+ \beta_4 \cdot (CorrectAns_{ij} \times Black_{ij}) \\
&+ \beta_5 \cdot (CorrectAns_{ij} \times Hispanic_{ij}) + \alpha_i + \gamma_j + \varepsilon_{ij}
\end{aligned} \tag{4.6}
$$

	(1)	(2)
Dependent variable: Ability Ratings		
Correct	2.936*** (0.0614)	2.910*** (0.0640)
Black	-0.0437 (0.0518)	-0.0650 (0.0502)
Hispanic	-0.0240 (0.0512)	-0.0189 (0.0494)
Correct × Black	0.0170 (0.0822)	0.0759 (0.0839)
Correct × Hispanic	0.0375 (0.0813)	0.0227 (0.0817)
Ability Fixed Effects	✓	✓
Respondent Fixed Effects		✓
Observations	13,512	13,512
R^2	0.4262	0.6564
Within R^2	0.4210	0.4736

* significant at 10%, ** at 5%, *** at 1%

Table 4.4. Pooled Regression Results by Race/Ethnicity

The results, reported in Table 4.4, are estimated both with and without respondent fixed effects. Across both specifications, the coefficients on the Black and Hispanic indicators are small and statistically insignificant, indicating no systematic differences in average ability ratings relative to White-associated profiles, conditional on performance. The interaction terms between performance and racial identity are also close to zero and statistically insignificant. This suggests that respondents do not differentially reward performance across racial groups. In other words, the relationship between performance and perceived ability is similar for Black-, Hispanic-, and White-associated profiles. Including respondent fixed effects substantially increases explanatory power but leaves the coefficients of interest largely unchanged. This indicates that the absence of differential treatment by race is not driven by unobserved heterogeneity in respondents' rating behavior. Overall, the pooled regression results provide no evidence of systematic race-based bias

in perceived ability ratings.

4.5.4 Intersectionality

	(1) Written	(2) Oral	(3) Inductive	(4) Number
Dependent variable: *Ability Ratings*				
Constant	2.574*** (0.082)	2.744*** (0.083)	2.356*** (0.092)	1.868*** (0.094)
Correct	0.904*** (0.043)	0.385*** (0.023)	0.506*** (0.031)	0.663*** (0.031)
Main effects				
Female	–0.038 (0.113)	–0.235** (0.119)	–0.188 (0.132)	0.088 (0.134)
Black	–0.045 (0.115)	–0.320*** (0.123)	–0.067 (0.131)	0.078 (0.133)
Hispanic	–0.062 (0.114)	–0.273** (0.118)	–0.133 (0.131)	0.283** (0.133)
Intersection (levels)				
Female × Black	0.140 (0.162)	0.299* (0.173)	0.178 (0.187)	–0.174 (0.189)
Female × Hispanic	0.185 (0.159)	0.150 (0.168)	0.213 (0.186)	–0.227 (0.190)
Heterogeneity in updating				
Correct × Female	–0.037 (0.060)	0.084** (0.033)	0.083* (0.043)	–0.051 (0.044)
Correct × Black	–0.038 (0.061)	0.111*** (0.034)	0.005 (0.043)	–0.029 (0.044)
Correct × Hispanic	–0.027 (0.062)	0.085*** (0.033)	0.051 (0.043)	–0.094** (0.044)
Intersection (updating)				
Correct × Female × Black	–0.032 (0.086)	–0.105** (0.048)	–0.090 (0.061)	0.038 (0.062)
Correct × Female × Hispanic	–0.009 (0.086)	–0.060 (0.047)	–0.095 (0.061)	0.082 (0.063)
Observations	4,096	3,844	3,987	3,992
R^2	0.370	0.374	0.321	0.370
Adj. R^2	0.368	0.372	0.319	0.369

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4.5. Intersectionality in ability ratings: fully interacted specification

Table 4.5 reports estimates from the fully interacted specification, which allows both baseline ratings and the responsiveness of ratings to performance to vary across gender, race/ethnicity, and their intersections. Across all four ability domains, the coefficients on the interaction terms between gender and race/ethnicity in levels are small in magnitude and statistically insignificant. This indicates no systematic evidence that individuals at the intersection of gender and race receive baseline ratings that differ from what would be predicted by the sum of the separate gender and race effects. Turning to the relationship between performance and perceived ability, we

find limited evidence of heterogeneity in updating across groups. In the written, inductive reasoning, and numerical domains, the triple interaction terms between performance, gender, and race/ethnicity are small and statistically insignificant, suggesting that respondents update their beliefs about ability similarly across intersectional groups.

The oral domain presents a partial exception. Ratings increase more steeply with performance for female and minority individuals, as indicated by positive and statistically significant coefficients on the interactions between number of correct answers and gender and race. However, the triple interaction between number of correct answers, female, and Black is negative and statistically significant at the 5% level, indicating that Black women receive less additional credit for correct answers than would be implied by the sum of the separate gender and race effects. While Black women still receive positive returns to performance, these returns are attenuated relative to an additive benchmark. Overall, the results provide little evidence of systematic intersectional bias in perceived ability ratings. The only detectable deviation from additivity is confined to the oral domain and does not generalize across other measures of ability.

4.5.5 Matching of gender

Table 4.6 reports estimates from the fully interacted specification, which allows both baseline ratings and the relationship between performance and perceived ability to vary with respondent gender, target gender, and their interaction. Across all four domains, the main effect of the hypothetical person's gender (*Female*) is small in magnitude and statistically insignificant, indicating that female profiles are not systematically rated differently from male profiles. The main effect of respondents' gender (*Own_Female*) is negative in all domains and statistically significant in the Number domain at the 5% level and in the Inductive domain at the 10% level. This suggests that female respondents tend to assign slightly lower ratings on average in some domains, reflecting differences in baseline rating levels rather than differential treatment across profiles. Turning to the relationship between performance and perceived ability, the interaction terms between the number of correct answers and gender are generally small and statistically insignificant. The only

	(1) Written	(2) Oral	(3) Inductive	(4) Number
Dependent variable: *Ability Ratings*				
Constant	2.623***	2.597***	2.406***	2.148***
	(0.074)	(0.078)	(0.083)	(0.087)
Correct	0.825***	0.431***	0.530***	0.599***
	(0.039)	(0.022)	(0.028)	(0.029)
Main effects				
Female	0.150 (0.104)	–0.083 (0.111)	–0.063 (0.121)	–0.056 (0.122)
Own Female	–0.145 (0.095)	–0.075 (0.101)	–0.202*	–0.263**
			(0.108)	(0.111)
Heterogeneity in updating				
Correct × Female	–0.063 (0.056)	0.036 (0.031)	–0.002 (0.039)	–0.005 (0.041)
Correct × Own Female	0.100* (0.051)	0.029 (0.028)	–0.005 (0.036)	0.039 (0.037)
Gender matching				
Female × Own Female	–0.129 (0.134)	–0.015 (0.143)	0.025 (0.155)	0.007 (0.158)
Correct × Female × Own Female	0.015 (0.072)	–0.008 (0.040)	0.033 (0.051)	–0.008 (0.052)
Observations	4,096	3,844	3,987	3,992
R^2	0.371	0.372	0.322	0.372
Adj. R^2	0.369	0.371	0.321	0.371

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4.6. Regression results for gender matching effects (fully interacted specification)

exception arises in the Written domain, where there is weak evidence that female respondents place greater weight on performance, as indicated by a marginally significant positive coefficient on the interaction between performance and respondent gender. However, this pattern does not generalize across domains.

Importantly, there is no evidence of gender-based matching effects. The interaction between respondent and target gender (*Female* × *Own Female*) is small and statistically insignificant in all domains, indicating no systematic differences in baseline ratings for same-gender pairs. Similarly, the triple interaction between performance, target gender, and respondent gender is close to zero and statistically insignificant across all domains, suggesting that respondents do not differentially translate performance into ratings when evaluating individuals of the same gender. Overall, these results provide no evidence of gender-based in-group bias in either baseline ratings or in the mapping from performance to perceived ability.

4.5.6 Matching of race/ethnicity

Respondents were asked to report their race/ethnicity by selecting all applicable categories from a list including Black, Hispanic, White, Asian, Native American, and Pacific Islander. As a result, these categories are not mutually exclusive. Appendix Table E.3 reports the distribution of responses across the Black, Hispanic, and White indicators used in the analysis. The majority of respondents (3,220 out of 4,096) select exactly one category, while approximately 14% report multiple categories (526 select two and 25 select all three), and 325 respondents do not select any of these three categories. To account for this, we adopt two complementary approaches in the analysis. First, we restrict the sample to respondents who identify with a single race/ethnicity among Black, Hispanic, or White, allowing for a clean definition of in-group matching that aligns with the single-race profiles used in the experiment. Second, we retain the full sample and define a flexible "same-race" indicator that allows for multiple racial/ethnic identities. This approach preserves sample size while accommodating multiracial respondents. As we show below, the results are qualitatively similar across both specifications.

Single-Race Sample Restriction

	(1) Written	(2) Oral	(3) Inductive	(4) Number
Dependent variable: *Ability Ratings*				
Constant	2.435*** (0.068)	2.614*** (0.071)	2.134*** (0.078)	1.754*** (0.080)
Correct Answers	0.933*** (0.037)	0.436*** (0.020)	0.599*** (0.025)	0.706*** (0.027)
Main effects				
Black	0.030 (0.097)	–0.281*** (0.103)	–0.032 (0.112)	–0.022 (0.111)
Hispanic	0.062 (0.094)	–0.330*** (0.100)	0.003 (0.111)	0.149 (0.113)
Own Black	0.454** (0.191)	0.421** (0.195)	0.631*** (0.222)	0.645*** (0.238)
Own Hispanic	0.197 (0.446)	–0.172 (0.525)	0.698 (0.444)	0.722 (0.462)
Heterogeneity in updating				
Correct × Black	–0.065 (0.052)	0.085*** (0.029)	–0.014 (0.036)	–0.019 (0.037)
Correct × Hispanic	–0.062 (0.051)	0.099*** (0.028)	0.004 (0.036)	–0.058 (0.038)
Correct × Own Black	–0.193* (0.099)	–0.077 (0.055)	–0.233*** (0.076)	–0.291*** (0.078)
Correct × Own Hispanic	–0.009 (0.217)	–0.029 (0.124)	–0.046 (0.146)	–0.204 (0.166)
Race matching (levels)				
Black × Own Black	–0.061 (0.287)	0.109 (0.280)	0.413 (0.313)	0.417 (0.337)
Hispanic × Own Hispanic	0.899 (0.661)	–0.358 (0.721)	0.438 (0.724)	0.072 (0.702)
Race matching (updating)				
Correct × Black × Own Black	0.167 (0.150)	–0.030 (0.079)	–0.136 (0.105)	–0.033 (0.108)
Correct × Hispanic × Own Hispanic	–0.311 (0.357)	0.155 (0.174)	–0.409· (0.227)	0.140 (0.260)
Observations	3,220	3,015	3,132	3,140
R^2	0.390	0.406	0.357	0.408
Adj. R^2	0.387	0.402	0.354	0.405

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4.7. Regression results for race/ethnicity matching effects (single-race sample, fully interacted specification)

Table 4.7 reports estimates from the fully interacted specification for the sample of respondents who identify with a single race/ethnicity. This specification allows both baseline ratings and the relationship between performance and perceived ability to vary with respondent race, race/ethnicity of the hypothetical persons, and their interaction. Across all four domains, there is little evidence of same-race matching effects. The interaction terms between respondent and hypothetical person' race/ethnicity are generally small and statistically insignificant, indicating no systematic dif-

ferences in baseline ratings for same-race pairs. Similarly, the triple interaction terms between performance, respondent race and hypothetical person's race are close to zero and statistically insignificant in most cases, suggesting that respondents do not differentially translate performance into ratings when evaluating individuals of their own race or ethnicity.

Instead, the results reveal heterogeneity driven primarily by respondent characteristics. In particular, Black respondents tend to assign higher ratings on average, as indicated by the positive and statistically significant coefficients on the *Own Black* indicator in all domains. At the same time, the interaction between performance and respondent race is negative for Black respondents and statistically significant in most domains, indicating that they place less weight on performance when forming ability ratings. This pattern suggests that differences across respondents operate primarily through baseline rating levels and the overall responsiveness to performance, rather than through differential treatment of specific groups.

The oral domain presents a distinct pattern. Minority profiles receive lower baseline ratings, as reflected in the negative and statistically significant coefficients on the indicators for Black and Hispanic hypothetical individuals. However, ratings for these profiles increase more steeply with performance, as indicated by positive and statistically significant interactions between performance and race/ ethnicity of the hypothetical person. This implies that while minority individuals start from a lower baseline in the oral domain, their performance is weighted more heavily in the evaluation process. Importantly, these effects are not driven by same-race matching, as the corresponding interaction terms with respondent race are not statistically significant. Overall, the results provide little evidence of race-based in-group bias in either baseline ratings or in the way respondents translate performance into perceived ability. Differences in evaluations are instead driven by general respondent tendencies and domain-specific patterns, particularly in the oral task.

Full Sample with Flexible Matching

Table 4.8 reports results using the full sample and a flexible definition of racial matching.

	(1) Written	(2) Oral	(3) Inductive	(4) Number
Dependent variable: *Ability Ratings*				
Constant	2.580*** (0.096)	2.607*** (0.097)	2.269*** (0.106)	1.895*** (0.109)
Correct Answers	0.840*** (0.050)	0.422*** (0.027)	0.570*** (0.035)	0.638*** (0.036)
Black	0.002 (0.106)	-0.154 (0.108)	0.017 (0.117)	0.008 (0.120)
Hispanic	0.008 (0.103)	-0.182* (0.105)	-0.032 (0.114)	0.186 (0.118)
Same Race	-0.036 (0.096)	0.028 (0.096)	-0.009 (0.103)	0.020 (0.108)
Heterogeneity in updating				
Correct × Black	-0.014 (0.055)	0.062** (0.030)	-0.059 (0.039)	-0.012 (0.040)
Correct × Hispanic	0.006 (0.054)	0.059** (0.029)	-0.015 (0.038)	-0.054 (0.039)
Correct × Same Race	0.059 (0.050)	0.004 (0.027)	-0.026 (0.035)	-0.001 (0.036)
Observations	4,096	3,844	3,987	3,992
R^2	0.369	0.373	0.321	0.369
Adj. R^2	0.368	0.372	0.320	0.368

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$, $^{\cdot}p < 0.1$.

Table 4.8. Regression results for race/ethnicity matching effects (full sample with flexible same-race indicator)

Across all four domains, the coefficient on the same-race indicator is small and statistically insignificant, indicating no systematic differences in baseline ratings for same-race evaluations. The interaction between performance and the same-race indicator is likewise small and statistically insignificant across all domains. The only notable pattern arises in the oral domain, where minority profiles receive lower baseline ratings (marginally significant for Hispanic) but exhibit a steeper relationship between performance and ratings, consistent with the restricted sample. Overall, these results align with those from the restricted sample and provide little evidence of systematic in-group bias based on shared racial or ethnic identity.

4.5.7 Heterogeneity due to incorrect assessment

Table 4.9 reports the distribution of self-assessment accuracy by gender across the four task domains. Overconfidence is prevalent across all domains but is especially pronounced among men, most notably in the numerical task, where over 80% of men overestimate their performance. Women also exhibit substantial overconfidence in some domains, but to a lesser extent. Overall, this indicates that miscalibration, primarily in the form of overconfidence, is widespread and

Table 4.9. Self-Assessment Accuracy by Gender and Task

	Written		Oral		Inductive		Number	
	Female	Male	Female	Male	Female	Male	Female	Male
Accurate	80.1%	81.1%	37.7%	28.8%	63.4%	58.3%	35.5%	16.8%
Overconfident	17.4%	17.4%	60.4%	70.0%	28.2%	36.7%	62.3%	82.4%
Underconfident	2.5%	1.5%	2.0%	1.2%	8.3%	5.0%	2.2%	0.8%
Observations	2471	1625	2316	1528	2402	1585	2403	1589

Notes: Respondents are classified as *overconfident* if their estimated number of correct answers exceeds their actual score by more than one question, *underconfident* if it falls short by more than one question, and *accurate* otherwise.

exhibits clear gender differences that vary by task.

Table 4.10 examines whether gender differences in ability ratings vary with the accuracy of respondents' self-assessments. Across all domains, there is little evidence of systematic gender bias among respondents who accurately assess their own performance, as the coefficient on the female indicator and its interaction with performance are generally small and statistically insignificant. An exception arises in the numerical domain, where accurate respondents exhibit a modest tendency to reward male performance more strongly.

Turning to miscalibration, overconfident respondents tend to give higher ratings on average, particularly in the oral and inductive domains, but place less weight on actual performance, as reflected in the negative and statistically significant interaction between performance and overconfidence. However, overconfidence does not systematically amplify gender bias. If anything, in the oral and numerical domains, the positive and significant triple interaction between performance, the female indicator for hypothetical person, and overconfidence suggests that overconfident respondents are relatively more responsive to female performance, attenuating or even reversing the modest pro-male bias observed among accurate respondents.

The patterns for underconfident respondents are less consistent but reveal notable heterogeneity in the written domain. Underconfident respondents assign significantly higher baseline rat-

	(1) Written	(2) Oral	(3) Inductive	(4) Number
Dependent variable: *Ability Ratings*				
Constant	2.528*** (0.052)	2.389*** (0.083)	2.128*** (0.068)	1.903*** (0.108)
Correct Answers	0.898*** (0.028)	0.496*** (0.024)	0.563*** (0.022)	0.618*** (0.036)
Main effects				
Female	-0.013 (0.073)	0.047 (0.119)	-0.003 (0.096)	0.045 (0.148)
Overconfident	0.140 (0.123)	0.270*** (0.104)	0.487*** (0.116)	0.138 (0.125)
Underconfident	-0.622** (0.301)	-0.352 (0.388)	0.128 (0.224)	-0.578 (0.378)
Heterogeneity in updating				
Correct × Female	-0.032 (0.039)	-0.041 (0.034)	0.015 (0.032)	-0.119** (0.049)
Correct × Overconfident	-0.086 (0.066)	-0.076*** (0.029)	-0.111*** (0.038)	0.004 (0.041)
Correct × Underconfident	0.021 (0.168)	0.087 (0.109)	-0.040 (0.071)	0.151 (0.120)
Miscalibration and gender bias (levels)				
Female × Overconfident	0.281 (0.173)	-0.243* (0.148)	-0.076 (0.167)	-0.146 (0.174)
Female × Underconfident	1.414*** (0.436)	0.849 (0.551)	-0.338 (0.304)	0.545 (0.600)
Miscalibration and gender bias (updating)				
Correct × Female × Overconfident	-0.055 (0.092)	0.118*** (0.041)	0.001 (0.054)	0.159*** (0.057)
Correct × Female × Underconfident	-0.427* (0.236)	-0.291* (0.157)	0.058 (0.098)	-0.322 (0.212)
Observations	4,091	3,840	3,982	3,985
R^2	0.374	0.376	0.326	0.381
Adj. R^2	0.373	0.374	0.324	0.379

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4.10. Miscalibration and gender differences in ability ratings

ings to female profiles, yet the marginally significant negative triple interaction indicates that they place relatively less weight on female performance compared to male performance. This implies a divergence between baseline perceptions and performance-based evaluations: while underconfident respondents appear more favorable toward women at low levels of performance, they reward male performance more strongly as performance increases.

Overall, these results provide little evidence that miscalibration uniformly exacerbates gender bias. Instead, different forms of miscalibration are associated with distinct patterns of evaluation, affecting both the level of ratings and the extent to which respondents rely on observed performance.

4.6 Conclusion

This paper investigates whether gender and race/ethnicity influence how individuals form perceptions of cognitive ability, using a controlled survey experiment in which respondents evaluate hypothetical individuals whose performance and demographic identity are independently randomized. This design allows us to cleanly separate the effect of identity from the effect of performance, addressing a key challenge in the existing literature on discrimination and bias.

The central finding is that respondents overwhelmingly anchor their evaluations to observed performance. Across all four cognitive domains, ability ratings increase monotonically and approximately linearly with the number of correct answers, and this relationship is remarkably stable across gender and racial/ethnic identities. Both descriptive patterns and regression results consistently show that the mapping from performance to perceived ability is similar across groups.

We find little evidence of systematic bias in average ratings. Profiles associated with women are not rated differently from profiles associated with men, either in levels or in the responsiveness of ratings to performance. Similarly, Black- and Hispanic-associated profiles receive ratings that are statistically indistinguishable from those assigned to White-associated profiles once

performance is held constant. While some isolated statistically significant differences emerge at specific points in the performance distribution, these effects are not consistent across domains or score levels and do not indicate a coherent pattern of bias.

The results also provide little support for more subtle forms of bias. We find no evidence of in-group favoritism based on shared gender or race/ethnicity. Intersectional analyses likewise reveal minimal departures from additivity, with the only notable exception arising in the oral domain, where Black women receive slightly attenuated returns to performance. However, this pattern is domain-specific and does not generalize across other measures of ability.

Instead, the analysis highlights the importance of evaluator characteristics, particularly in shaping baseline rating levels and the weight placed on performance. Differences across respondents, such as variation in self-assessment accuracy, play a more meaningful role than the identity of the evaluated individual. Overconfident respondents tend to assign higher ratings but rely less on performance, particularly in the oral and inductive domains. In the written domain, underconfident respondents appear more favorable toward women at low performance levels but reward male performance more strongly as performance increases. These findings suggest that heterogeneity in evaluation arises more from differences in how individuals process information than from systematic discrimination against specific groups.

In summary, the results point to a setting in which evaluators largely rely on observable performance signals when forming assessments of ability, with limited evidence of pervasive bias based on gender or race/ethnicity. While this does not rule out the existence of bias in other contexts or under different informational environments, the findings suggest that clear and interpretable performance information can substantially constrain the scope for discriminatory evaluations.

This chapter has several limitations. First, the analysis is based on evaluations of hypothetical profiles rather than real-world hiring or promotion decisions, and respondents may behave dif-

ferently when evaluations carry tangible economic consequences. Second, respondents observe clear and interpretable performance signals, which likely reduce reliance on demographic stereotypes relative to environments where productivity is more difficult to assess or only imperfectly observed. As a result, the findings may be most applicable to settings in which performance information is transparent and standardized. Future research could explore whether similar patterns emerge in higher-stakes environments or in settings where evaluators face greater uncertainty about underlying ability.

Chapter II Appendix

A.1 Supplementary Tables and Figures

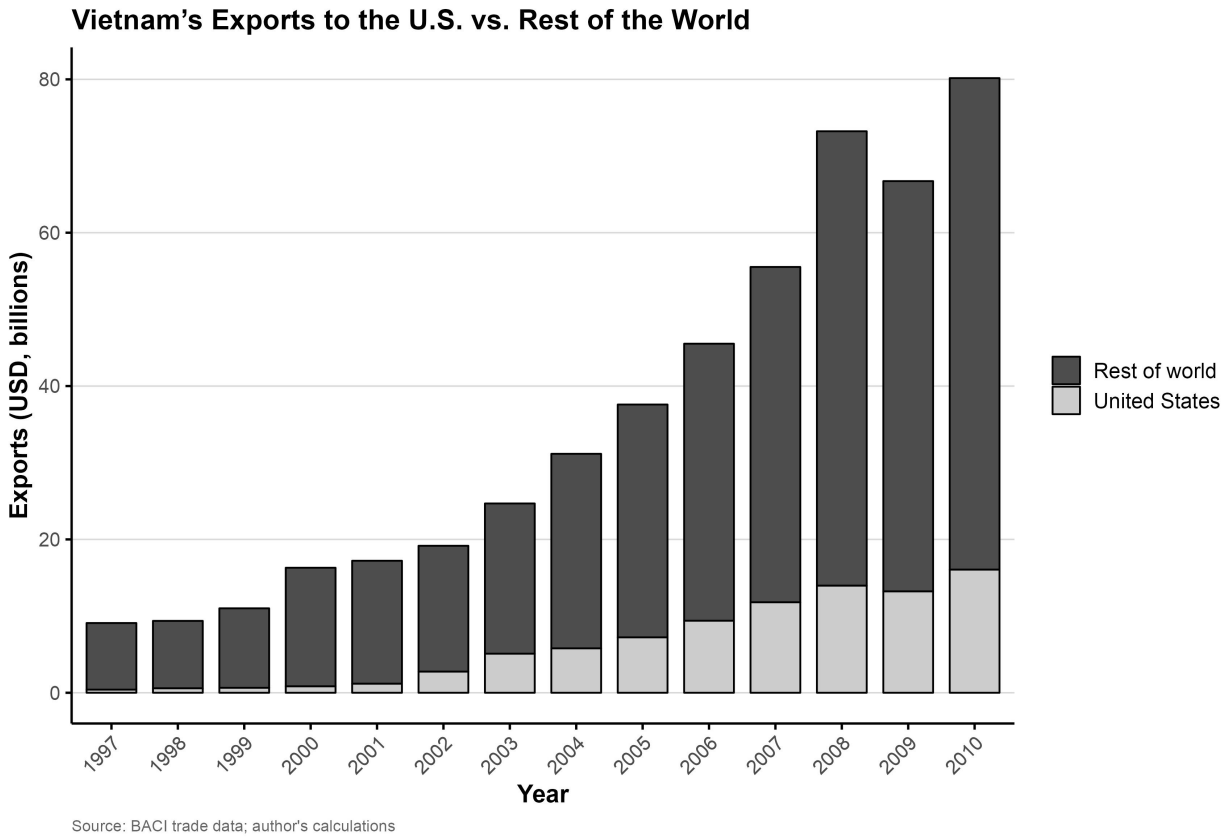


Figure A.1. Vietnam's Export

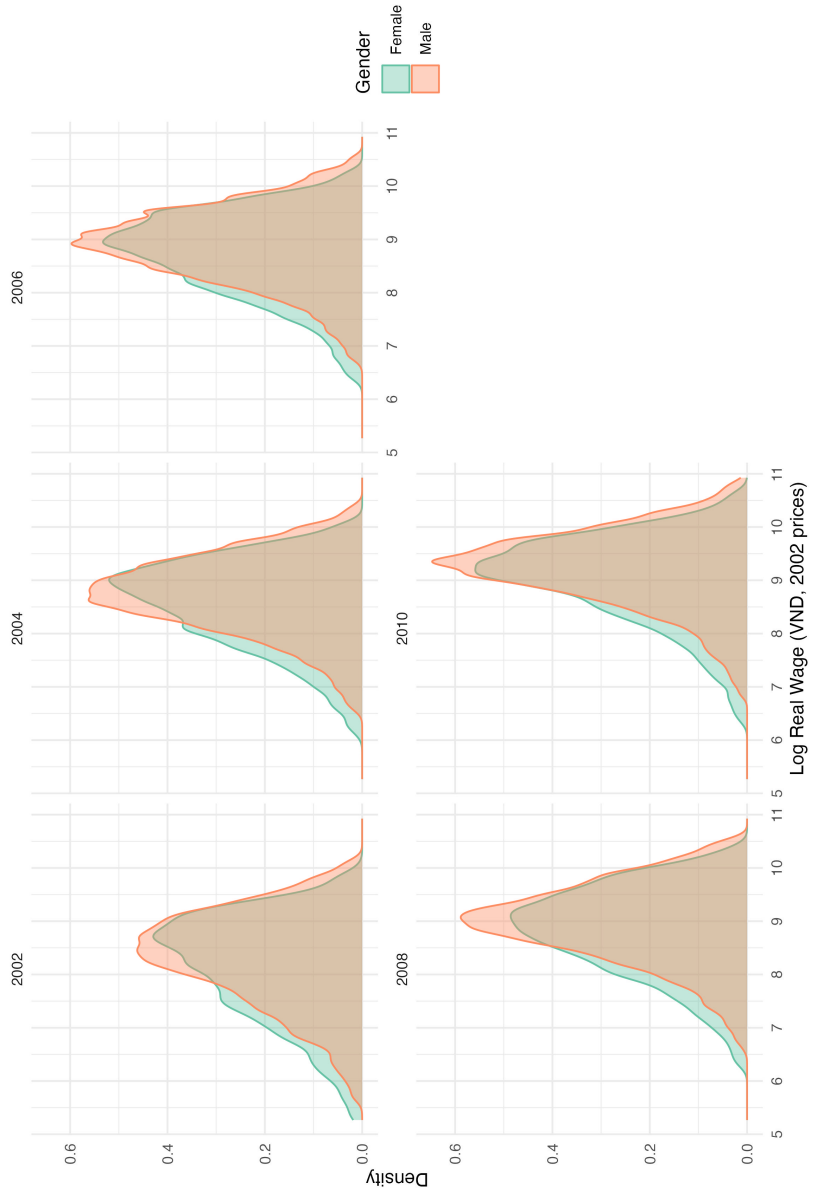


Figure A.2. Wage Distribution by Gender

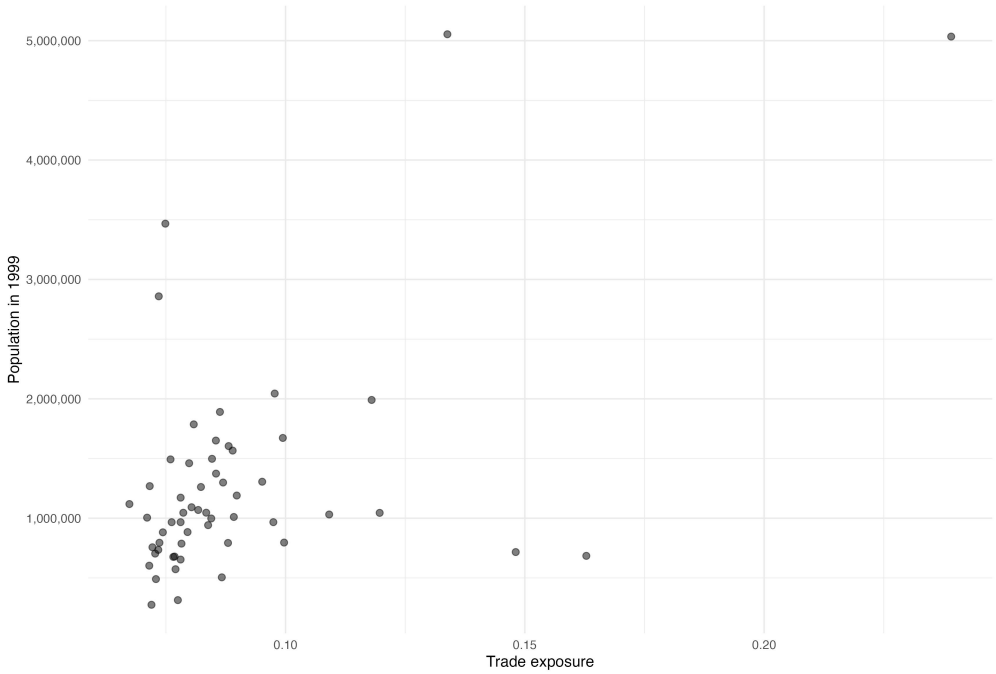


Figure A.3. Provincial Trade Exposure vs. Population (1999)

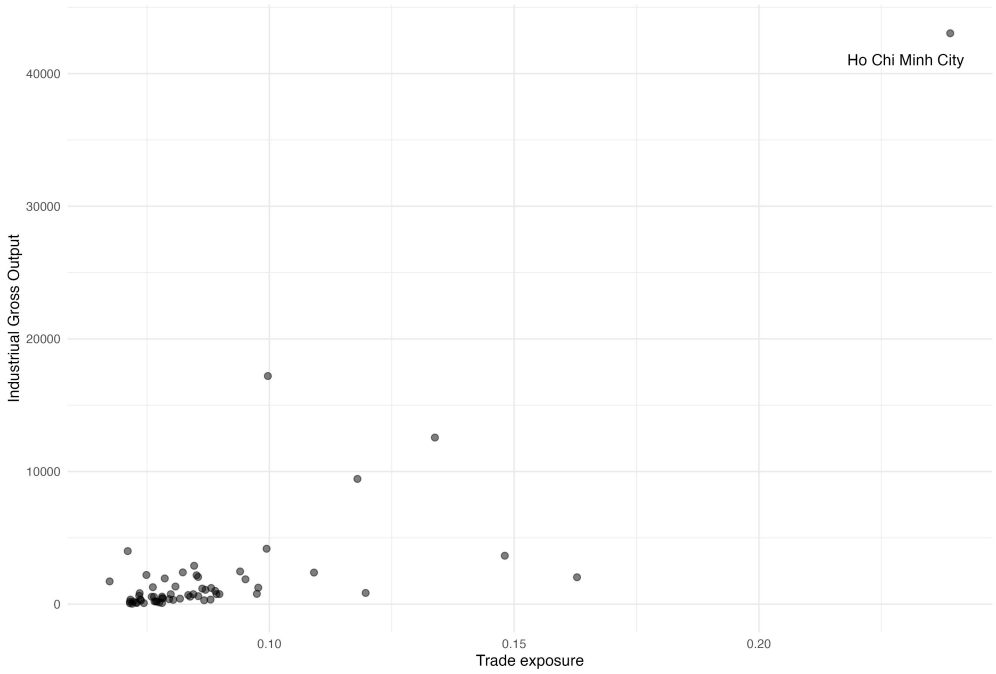


Figure A.4. Provincial Trade Exposure vs. Industrial Gross Output (1997)

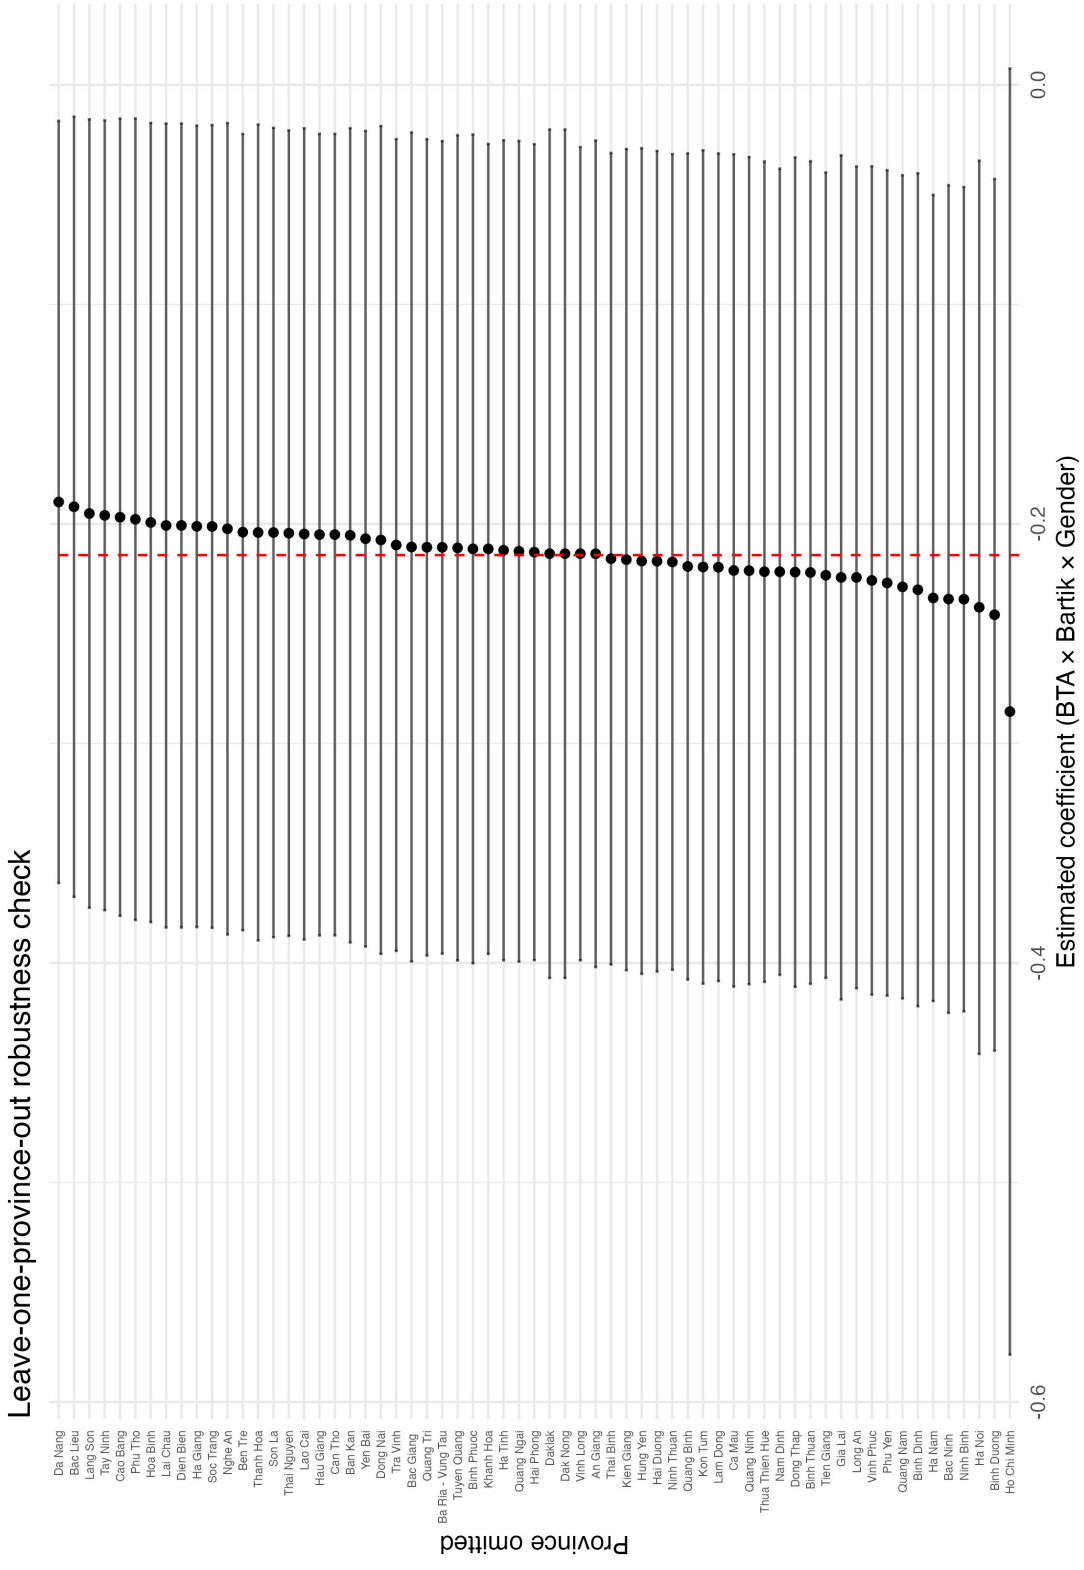


Figure A.5. Distribution of leave-one-province-out estimates of the coefficient on *Trade_exposure* × *Post_BTA* × *Female*.

Table A.1. Summary Statistics by Gender and Year

	Male	Female	Male – Female
2002			
age	33.260	32.167	1.093*** (0.142)
edu	7.869	7.678	0.191*** (0.047)
vocational	0.093	0.145	-0.052*** (0.004)
college	0.065	0.056	0.009*** (0.003)
marital	1.699	1.758	-0.059*** (0.008)
urban	0.717	0.670	0.047*** (0.006)
ethnic	1.607	1.607	0.000 (0.032)
2004			
age	34.078	32.324	1.754*** (0.145)
edu	8.513	8.489	0.024 (0.045)
vocational	0.167	0.191	-0.024*** (0.005)
college	0.110	0.135	-0.025*** (0.004)
marital	1.685	1.735	-0.050*** (0.008)
urban	0.647	0.590	0.058*** (0.006)
ethnic	1.397	1.428	-0.031 (0.030)
2006			
age	34.535	32.854	1.681*** (0.130)
edu	8.677	8.768	-0.091** (0.039)
vocational	0.168	0.180	-0.012*** (0.004)
college	0.119	0.158	-0.039*** (0.004)
marital	1.689	1.749	-0.060*** (0.007)
urban	0.623	0.564	0.059*** (0.005)
ethnic	1.452	1.400	0.052** (0.024)
2008			
age	34.996	33.513	1.483*** (0.128)
edu	8.826	8.981	-0.154*** (0.038)
vocational	0.181	0.189	-0.008** (0.004)
college	0.128	0.174	-0.046*** (0.004)
marital	1.709	1.788	-0.079*** (0.007)
urban	0.628	0.558	0.070*** (0.005)
ethnic	1.417	1.358	0.058** (0.023)
2010			
age	35.473	33.825	1.648*** (0.121)
edu	8.863	9.057	-0.194*** (0.036)
vocational	0.191	0.165	0.026*** (0.004)
college	0.142	0.205	-0.063*** (0.004)
marital	1.764	1.857	-0.093*** (0.007)
urban	0.612	0.553	0.059*** (0.005)
ethnic	1.517	1.440	0.077*** (0.026)

Table A.2. Estimated Effects of Trade Exposure on Individual Log Real Wages: Accounting for Survey Weights

Dependent variable:	*log real wage*		
	(1) Baseline	(2) Weighted	(3) Weighted
Trade_exposure × Post-BTA	−2.466***	−2.210***	
	(0.4147)	(0.3024)	
Trade_exposure × Post-BTA × Female	−0.2142*	−0.1708*	−0.1827*
	(0.0947)	(0.0701)	(0.0738)
Province Fixed Effects	Yes	Yes	No
Year Fixed Effects	Yes	Yes	No
Province × Year Fixed Effects	No	No	Yes
Standard Errors	Clustered by province		
Observations	158,302	158,302	158,302
R^2	0.612	0.621	0.630
Within R^2	0.506	0.502	0.504

Notes: Observations in the top and bottom 1% of the real wage distribution are removed by gender and year to reduce the influence of outliers. All specifications control for gender, age and its square, education, college and vocational dummies, marital status, urban residence, ethnicity, and log of total time worked. Columns (2) and (3) use VHLSS sampling weights normalized to have a mean of one within each survey year. Robust standard errors clustered by province are reported in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table A.3. Estimated Effects of Trade Exposure: Adding Industry-by-year FEs

Dependent variable:	*log real wage*	
	(1)	(2)
Trade exposure × Post-BTA	−2.466*** (0.4147)	−1.095** (0.3940)
Trade exposure × Post-BTA × Female	−0.2142* (0.0947)	−0.1780* (0.0802)
Province Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	No
Industry × Year Fixed Effects	No	Yes
Standard Errors	Clustered by province	
Observations	158,302	158,302
R^2	0.612	0.647
Within R^2	0.506	0.427

Notes: Observations in the top and bottom 1% of the real wage distribution are removed by gender and year to reduce the influence of outliers. All specifications control for gender, age and its square, education, college and vocational dummies, marital status, urban residence, ethnicity, and log of total time worked. Robust standard errors clustered by province are reported in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table A.4. Estimated Effects of Trade Exposure: Binary Treatment Specification

Dependent variable:	*log real wage*		
	(1)	(2)	(3)
More_exposed × Post-BTA	−0.1734***	−0.1220*	
	(0.0464)	(0.0461)	
More_exposed × Post-BTA × Female	−0.0516**	−0.0420**	−0.0514**
	(0.0163)	(0.0128)	(0.0163)
Province Fixed Effects	Yes	Yes	No
Year Fixed Effects	Yes	No	No
Industry × Year Fixed Effects	No	Yes	No
Province × Year Fixed Effects	No	No	Yes
Standard Errors	Clustered by province		
Observations	158,302	158,302	158,302
R^2	0.612	0.647	0.623
Within R^2	0.506	0.428	0.510

Notes: Observations in the top and bottom 1% of the real wage distribution are removed by gender and year to reduce the influence of outliers. All specifications control for gender, age and its square, education, college and vocational dummies, marital status, urban residence, ethnicity, and log of total time worked. Robust standard errors clustered by province are reported in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table A.5. Estimated Effects of Trade Exposure: Alternative Wage Measure (Hourly Wage)

Dependent variable:	*log real hourly wage*		
	(1)	(2)	(3)
More_exposed × Post-BTA	−0.1791***	−0.1433**	
	(0.0476)	(0.0502)	
More_exposed × Post-BTA × Female	−0.0479**	−0.0324*	−0.0475**
	(0.0169)	(0.0131)	(0.0168)
Province Fixed Effects	Yes	Yes	No
Year Fixed Effects	Yes	No	No
Industry × Year Fixed Effects	No	Yes	No
Province × Year Fixed Effects	No	No	Yes
Standard Errors		Clustered by province	
Observations	158,302	158,302	158,302
R^2	0.369	0.417	0.386
Within R^2	0.236	0.148	0.235

Notes: The dependent variable is the log of the hourly real wage, calculated as annual real earnings divided by total time worked. Observations in the top and bottom 1% of the wage distribution are removed by gender and year to reduce the influence of outliers. All specifications control for gender, age and its square, education, college and vocational dummies, marital status, urban residence, ethnicity, and log of total time worked. Robust standard errors clustered by province are reported in parentheses. $^{\cdot}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

A.2 Data Appendix

A.2.1 Changes Made to Industry Codes in 1999 Population and Housing Census (Accessed Through IPUMS–International)

Old industry code	New industry code	Number of observations
701	731	717
702	732	45
711	701	173
721	711	22
722	712	16
723	713	261
732	722	16
733	723	6
735	725	33
739	729	31
901	921	2197
902	922	148
903	923	352
904	924	659
913	911	127
920	900	1550

Table A.6. Industry Code Conversion and Number of Observations
Source: McCaig (2011), IPUMS International Variable Description, ISIC Revision 3 Document.

Chapter III Appendix

Supplementary Tables & Figures

Table C.1. Timeline of U.S. and Chinese Trade Policy Measures

Date	Tariff Description	HS10
Panel A: U.S. Tariffs on China		
02/07/2018	Sec. 201: Solar and washing machines (announced 01/22/2018)	16
03/23/2018	Sec. 232: Steel and aluminum (some countries exempted)	800
07/06/2018	Sec. 301 List 1 ($34B)	1,629
08/23/2018	Sec. 301 List 2 ($16B)	401
09/24/2018	Sec. 301 List 3 ($200B)	8,664
02/07/2019	Sec. 201 tariffs reduced (solar/washers)	16
06/01/2019	Sec. 301 List 3 tariff $\uparrow$ (10% $\rightarrow$ 25%)	8,664
09/01/2019	Sec. 301 List 4A ($112B, 15%)	5,301
02/07/2020	Sec. 201 tariffs reduced again	99
02/08/2020	Sec. 232 extended to derivative products	–
02/14/2020	Sec. 301 List 4A tariff $\downarrow$ (15% $\rightarrow$ 7.5%)	5,301
Panel B: Chinese Retaliation		
04/02/2018	Retaliation to Sec. 232	85
05/01/2018	MFN change: pharmaceuticals	1
07/01/2018	MFN change: consumer goods, autos, ITA	88
07/06/2018	Retaliation to Sec. 301 List 1 ($34B)	451
08/23/2018	Retaliation to Sec. 301 List 2 ($16B)	0
09/24/2018	Retaliation to Sec. 301 List 3 ($60B)	81
11/01/2018	MFN change: industrial goods	16
01/01/2019	Temporary MFN rates (2019)	10
01/01/2019	Suspension of auto retaliation	0
06/01/2019	Retaliation increase (subset of $60B)	88
07/01/2019	MFN cut on ITA products	0
09/01/2019	Retaliation increase (subset of $70B)	563
01/01/2020	Temporary MFN rates (2020)	5
02/14/2020	Retaliation (subset of $75B)	563
07/01/2020	MFN cut on ITA products	0
01/01/2021	Temporary MFN rates (2021)	5

Source: Bown (2021)

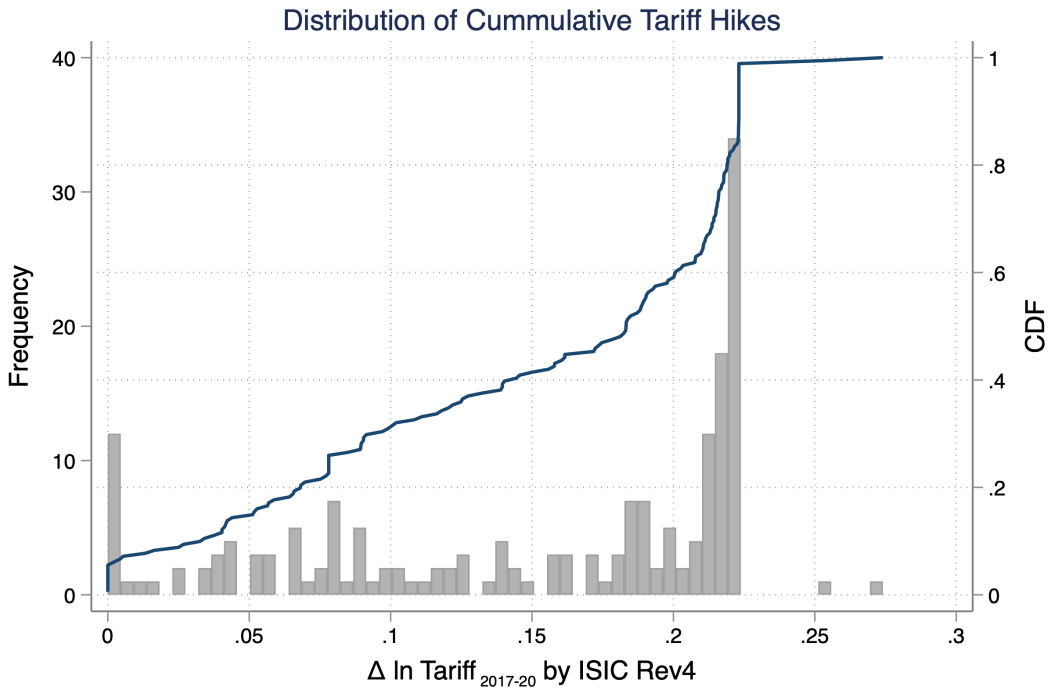


Figure C.1. Distribution of Cumulative Tariff Increases at ISIC4 Level

Notes: Source: Bown (2021) and authors' calculations.

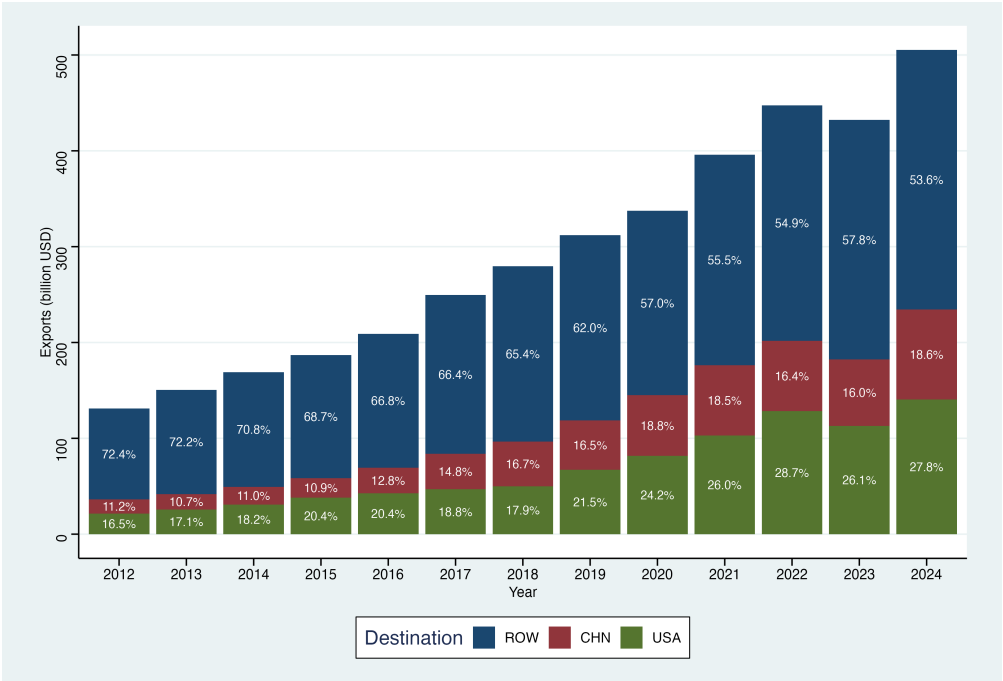


Figure C.2. Vietnam's Exports by Destination

Notes: Source: BACI and authors' calculations.

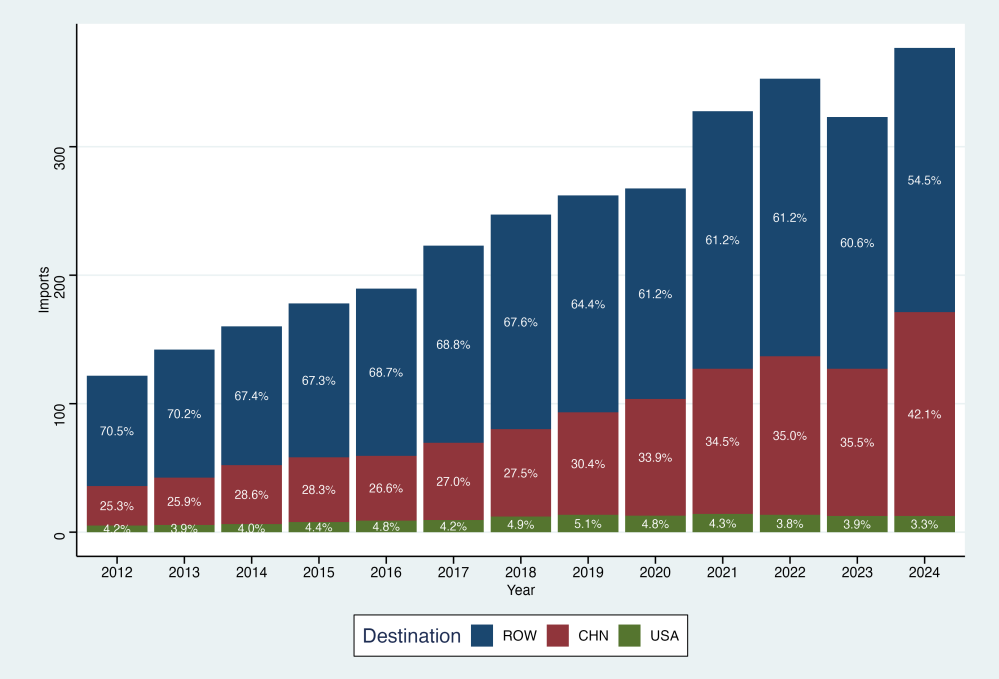


Figure C.3. Vietnam's Imports by Destination

Notes: Source: BACI and authors' calculations.

Table C.2. Effect of Trade War Exposure on Employment - Unbalanced Panel

	(1)	(2)
Dependent variable: log(employment)		
Trade War Tariff	0.2634	1.1550**
	(0.1999)	(0.3764)
Firm FE	✓	✓
Year FE	✓	✓
Controls	✓	✓
Observations	623,339	102,922
R^2	0.881	0.880
Within R^2	0.00032	0.00390
SE clustering	isic4 & province	

Notes: Column (1) uses the full unbalanced sample; column (2) restricts to firms with pre-trade-war trading experience. Controls include the average MFN tariff applied to Vietnamese exports and interactions between the share of non-treated Vietnamese exports and year indicators. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$, $^{\cdot}p < 0.1$.

Table C.3. Effect of Trade War Tariffs on Female Employment Share by Ownership

	Full Sample			Pre-TW Traders Only		
	Private	State	Foreign	Private	State	Foreign
Trade War Tariff	0.4928**	-0.3668	0.3237**	0.3806*	-0.1649	0.3339**
	(0.1433)	(0.2949)	(0.1157)	(0.1700)	(0.2795)	(0.1099)
Controls	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Observations	142,229	1,844	32,226	34,126	1,257	30,836
R^2	0.685	0.930	0.882	0.786	0.945	0.884
Within R^2	0.00194	0.00618	0.00226	0.00164	0.00332	0.00241
SE clustering	isic4 & province					

Notes: The dependent variable is the log of the female employment share, defined as log(emp_end_female/emp_end). Columns (1)–(3) use the full sample of firms, while columns (4)–(6) restrict the sample to firms with pre-trade war trading experience. Controls include the average MFN tariff applied to Vietnamese exports and interactions between the share of non-treated Vietnamese exports and year indicators. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C.4. Effect of Trade War Tariffs on Average Wages by Ownership

	Full Sample			Pre-TW Traders Only		
	Private	State	Foreign	Private	State	Foreign
Trade War Tariff	0.1244	-0.4342	-0.2068	-0.2516**	-0.6175*	-0.1932
	(0.0867)	(0.2803)	(0.1930)	(0.1150)	(0.3635)	(0.1919)
Controls	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Observations	126,381	1,846	32,224	31,503	1,261	30,830
R^2	0.981	0.996	0.995	0.989	0.995	0.995
Within R^2	0.00131	0.00552	0.00034	0.00038	0.00469	0.00035
SE clustering	isic4 & province					

Notes: The dependent variable is log average wages, measured as log(total_comp/emp_end). Columns (1)–(3) use the full sample of firms, while columns (4)–(6) restrict the sample to firms with pre-trade war trading experience. Controls include the average MFN tariff applied to Vietnamese exports and interactions between the share of non-treated Vietnamese exports and year indicators. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.05$, $^{\cdot}p < 0.1$.

Chapter IV Appendix

Appendix

	White	Black	Hispanic
Male	Michael	Willie	Jose
	John	Darnell	Juan
	David	Jermaine	Carlos
Female	Mary	Latoya	Ana
	Susan	Latasha	Rosa
	Jennifer	Tamika	Irma

Table E.1. Names by Gender and Race/Ethnicity

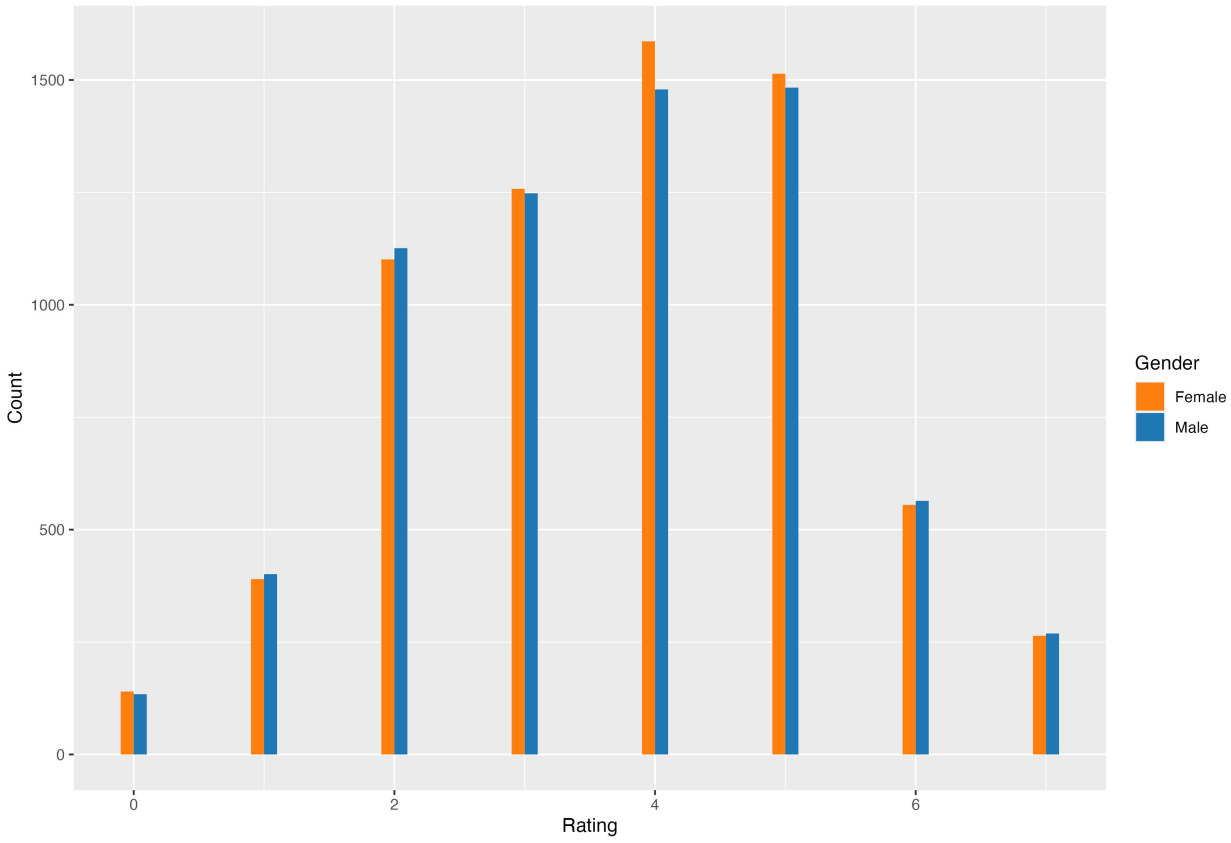


Figure E.1. Distribution of Ratings by Gender (Unconditional)

	(1)	(2)
Dependent variable: High Rating		
Correct	0.3454*** (0.0127)	0.3557*** (0.0139)
Female	-0.0011 (0.0067)	0.0011 (0.0077)
Correct × Female	-0.0117 (0.0178)	-0.0258 (0.0182)
Ability FE	✓	✓
Respondent FE		✓
Observations	13,512	13,512
R^2	0.1317	0.4388
Within R^2	0.1304	0.1504

	(3)	(4)
Dependent variable: Low Rating		
Correct	-0.6320*** (0.0132)	-0.6230*** (0.0151)
Female	-0.0025 (0.0140)	0.0040 (0.0138)
Correct × Female	0.0000 (0.0185)	-0.0019 (0.0195)
Ability FE	✓	✓
Respondent FE		✓
Observations	13,512	13,512
R^2	0.2711	0.5248
Within R^2	0.2648	0.2868

Notes: High rating is defined as 6–7 and low rating as 0–2. Standard errors in parentheses. Robust standard errors in columns (1) and (3); clustered at the respondent level in columns (2) and (4). * p<0.10, ** p<0.05, *** p<0.01.

Table E.2. Tail Analysis: Linear Probability Models for High and Low Ratings

Table E.3. Distribution of Reported Race/Ethnicity Indicators

Number of categories selected	Count
0	325
1	3,220
2	526
3	25
Total	4,096

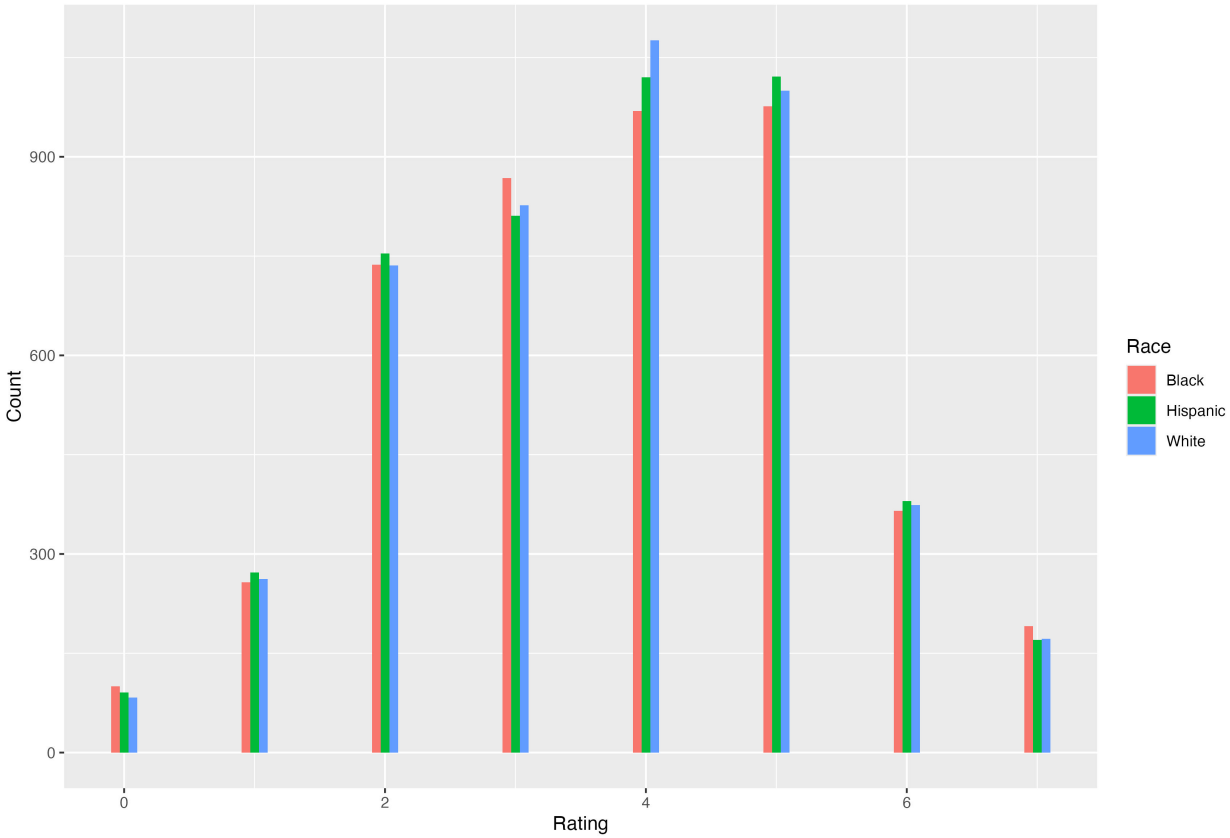


Figure E.2. Distribution of Ratings by Race (Unconditional)

References

Acemoglu, D., Autor, D. H., and Lyle, D. (2004). Women, war, and wages: The effect of female labor supply on the wage structure at midcentury. *Journal of political Economy*, 112(3):497–551.

Aguayo-Tellez, E., Airola, J., Juhn, C., and Villegas-Sanchez, C. (2014). Did trade liberalization help women? the case of mexico in the 1990s. In *New analyses of worker well-being*, pages 1–35. Emerald Group Publishing Limited.

Alfaro, L. and Chor, D. (2023). Global Supply Chains: The Looming "Great Reallocation". Technical report, National Bureau of Economic Research.

Anker, R., Melkas, H., Korten, A., et al. (2003). *Gender-based occupational segregation in the 1990s*, volume 16. Citeseer.

Anne, B. (2017). Gender biases in student evaluations of teachers. *Journal of Public Economics*, 145:27–41.

Antràs, P. and Yeaple, S. R. (2014). Multinational firms and the structure of international trade. *Handbook of international economics*, 4:55–130.

Arrow, K. J. (1973). The theory of discrimination. In Ashenfelter, O. and Rees, A., editors, *Discrimination in Labor Markets*, pages 3–33. Princeton University Press, Princeton, NJ.

Autor, D. H., Dorn, D., and Hanson, G. H. (2016). The china shock: Learning from labor-market adjustment to large changes in trade. *Annual review of economics*, 8(1):205–240.

Baccini, L., Impullitti, G., and Malesky, E. J. (2019). Globalization and state capitalism: Assessing vietnam's accession to the wto. *Journal of International Economics*, 119:75–92.

Banerjee, U., Peñarrieta, L. C., and Chakraborty, P. (2025). Can trade policy change gender equality? evidence from chile. *Journal of International Economics*, page 104143.

Banerji, A., Gjonbalaj, A., Hlatshwayo, S., and Van Le, A. (2018). Asian women at work. *Finance & Development*, 55(003):A004.

Becker, G. S. (1957). *The economics of discrimination*. University of Chicago press.

Berik, G., Rodgers, Y. v. d. M., and Zveglich, J. E. (2004). International trade and gender wage discrimination: Evidence from east asia. *Review of Development Economics*, 8(2):237–254.

Bertrand, M. and Mullainathan, S. (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013.

Black, S. E. and Brainerd, E. (2004). Importing equality? the impact of globalization on gender discrimination. *ILR Review*, 57(4):540–559.

Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of economic literature*, 55(3):789–865.

Blinder, A. S. (1973). Wage discrimination: reduced form and structural estimates. *Journal of Human resources*, pages 436–455.

Bohren, J. A., Haggag, K., Imas, A., and Pope, D. G. (2023). Inaccurate statistical discrimination: An identification problem. *Review of Economics and Statistics*, 2023:1–45.

Bøler, E. A., Javorcik, B. S., and Ulltveit-Moe, K. H. (2015). Globalization: a woman's best friend? exporters and the gender wage gap.

Borrowman, M. and Klasen, S. (2020). Drivers of gendered sectoral and occupational segregation in developing countries. *Feminist Economics*, 26(2):62–94.

Bown, C. P. (2021). The US–China trade war and Phase One agreement. *Journal of Policy Modeling*, 43(4):805–843.

Brambilla, I., Lederman, D., and Porto, G. (2012). Exports, export destinations, and skills. *American Economic Review*, 102(7):3406–3438.

Braunstein, E. and Brenner, M. (2007). Foreign direct investment and gendered wages in urban china. *Feminist Economics*, 13(3-4):213–237.

Brussevich, M. (2018). Does trade liberalization narrow the gender wage gap? the role of sectoral mobility. *European Economic Review*, 109:305–333.

Buchhave, H., Cunningham, W. V., Nguyen, G. T., and Weimann-Sandig, N. (2020). *Perceptions of Gender Disparities in Vietnam's Labor Market*. World Bank Group.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *American economic review*, 101(1):304–340.

Çağatay, N. and Özler, Ş. (1995). Feminization of the labor force: The effects of long-term development and structural adjustment. *World development*, 23(11):1883–1894.

Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica*, 87(3):741–835.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Card, D., Cardoso, A. R., and Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly journal of economics*, 131(2):633–686.

Card, D. and Lemieux, T. (2001). Can falling supply explain the rising return to college for younger men? a cohort-based analysis. *The quarterly journal of economics*, 116(2):705–746.

Chan, J. (2018). Does import competition worsen the gender gap? evidence from matched employer–employee data. *Economics Letters*, 165:13–16.

Chen, N., Novy, D., and Solórzano, D. (2025). Trade diversion and labor market outcomes. Technical report, CESifo Working Paper.

Chen, Z., Ge, Y., Lai, H., and Wan, C. (2013). Globalization and gender wage inequality in china. *World Development*, 44:256–266.

Collins, P. H. and Bilge, S. (2020). *Intersectionality*. John Wiley & Sons.

Correll, S. J., Benard, S., and Paik, I. (2007). Getting a job: Is there a motherhood penalty? *American Journal of Sociology*, 112(5):1297–1338.

Demombynes, G. and Vu, L. H. (2022). Vietnam's household registration system. Technical report, World Bank Group, Washington, D.C. English.

Dix-Carneiro, R. and Kovak, B. K. (2017). Trade liberalization and regional dynamics. *American Economic Review*, 107(10):2908–2946.

Dix-Carneiro, R. and Kovak, B. K. (2023). Globalization and inequality in latin america. Working Paper 31459, National Bureau of Economic Research.

Do, Q.-T., Levchenko, A. A., and Raddatz, C. (2016). Comparative advantage, international trade, and fertility. *Journal of Development Economics*, 119:48–66.

Dollar, D. and Kraay, A. (2004). Trade, growth, and poverty. *The economic journal*, 114(493):F22–F49.

Dube, A., Girardi, D., Jordà, Ò., and Taylor, A. M. (2025). A local projections approach to difference-in-differences. *Journal of Applied Econometrics*, 40(7):741–758.

Ederington, J., Minier, J., and Troske, K. R. (2009). Where the girls are: Trade and labor market segregation in colombia. Technical report, IZA Discussion Papers.

Erten, B. and Keskin, P. (2024). Trade-offs? the impact of wto accession on intimate partner violence in cambodia. *Review of Economics and Statistics*, 106(2):322–333.

Fajgelbaum, P., Goldberg, P., Kennedy, P., Khandelwal, A., and Taglioni, D. (2024). The us-china trade war and global reallocations. *American Economic Review: Insights*, 6(2):295–312.

Fajgelbaum, P. D., Goldberg, P. K., Kennedy, P. J., and Khandelwal, A. K. (2020). The return to protectionism. *The quarterly journal of economics*, 135(1):1–55.

Freund, C., Mattoo, A., Mulabdic, A., and Ruta, M. (2024). Is US trade policy reshaping global supply chains? *Journal of International Economics*, 152:104011.

Gaddis, I. and Pieters, J. (2017). The gendered labor market impacts of trade liberalization: Evidence from brazil. *Journal of Human Resources*, 52(2):457–490.

Galor, O. and Weil, D. N. (1996). The gender gap, fertility, and growth. *The American Economic Review*, 86(3):374–387.

Gaulier, G. and Zignago, S. (2010). BACI: International Trade Database at the Product-Level (The 1994–2007 Version).

George, E. and Livingston, G. (2024). What you need to know about the gender wage gap. U.S. Department of Labor Blog. https://blog.dol.gov/2024/03/12/what-you-need-to-know-about-the-gender-wage-gap.

Gneezy, U., Niederle, M., and Rustichini, A. (2003). Performance in competitive environments: Gender differences. *The quarterly journal of economics*, 118(3):1049–1074.

Goldberg, P. K. and Pavcnik, N. (2007). Distributional effects of globalization in developing countries. *Journal of Economic Literature*, 45(1):39–82.

Grodsky, E. and Pager, D. (2001). The structure of disadvantage: Individual and occupational determinants of the black-white wage gap. *American Sociological Review*, 66(4):542–567.

Hardy, A. (2001). Rules and resources: negotiating the household registration system in vietnam under reform. *Sojourn: Journal of Social Issues in Southeast Asia*, pages 187–212.

Hazarika, G. and Otero, R. (2004). Foreign trade and the gender earnings differential in urban mexico. *Journal of Economic Integration*, pages 353–373.

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

Helpman, E., Melitz, M. J., and Yeaple, S. R. (2004). Export versus fdi with heterogeneous firms. *American economic review*, 94(1):300–316.

Hirsch, B. (2010). *Monopsonistic labour markets and the gender pay gap: theory and empirical evidence*, volume 639. Springer Science & Business Media.

Hirsch, B., Schank, T., and Schnabel, C. (2010). Differences in labor supply to monopsonistic firms and the gender pay gap: An empirical analysis using linked employer-employee data from germany. *Journal of Labor Economics*, 28(2):291–330.

Hoang, T., Mitra, D., and Pham, H. (2024). The effect of export market access on labor market power: Firm-level evidence from vietnam. Technical report, IZA Discussion Papers.

Huynh, N. and Nguyen, N. (2025). Trade, maternal time costs, and sex selection: Evidence from vietnam. *Available at SSRN 4582123*.

Huynh, Q. and Ku, H. (2025). Gender norms and female labor supply: Evidence from export shocks in vietnam.

Iyoha, E., Malesky, E., Wen, J., Wu, S.-J., and Feng, B. (2024). *Exports in Disguise?: Trade Rerouting during the US-China Trade War*. Harvard Business School.

Juhn, C., Ujhelyi, G., and Villegas-Sanchez, C. (2014). Men, women, and machines: How trade impacts gender inequality. *Journal of Development Economics*, 106:179–193.

Kelishomi, A. M. and Nisticò, R. (2024). Economic sanctions and informal employment. *Labour Economics*, 89:102581.

Keller, W. and Utar, H. (2022). Globalization, gender, and the family. *The Review of Economic Studies*, 89(6):3381–3409.

Kline, P., Rose, E. K., and Walters, C. R. (2022). Systemic discrimination among large us employers. *The Quarterly Journal of Economics*, 137(4):1963–2036.

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

Krueger, A. B. and Summers, L. H. (1988). Efficiency wages and the inter-industry wage structure. *Econometrica: Journal of the Econometric Society*, pages 259–293.

Lam, T. H. (2019). *Chinese FDI in Vietnam: Trends, status and challenges*, volume 34. ISEAS-Yusof Ishak Institute.

Lu, Y. and Yu, L. (2015). Trade liberalization and markup dispersion: Evidence from china's wto accession. *American Economic Journal: Applied Economics*, 7(4):221–253.

Mansour, H., Medina, P., and Velasquez, A. (2022). Import competition and gender differences in labor reallocation. *Labour Economics*, 76:102149.

Mayr-Dorn, K., Narciso, G., Dang, D. A., and Phan, H. (2026). Trade diversion and labor market adjustment: Vietnam and the u.s.-china trade war. *European Economic Review*, 185:105282.

McCaig, B. (2011). Exporting out of poverty: Provincial poverty in vietnam and us market access. *Journal of International Economics*, 85(1):102–113.

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. (2022). Fdi inflows and domestic firms: Adjustments to new export opportunities. Working Paper 30729, National Bureau of Economic Research.

McCaig, B., Pavcnik, N., and Wong, W. F. (2025). Foreign and domestic firms: Long run employment effects of export opportunities. *Review of Economics and Statistics*, pages 1–45.

McCall, L. (2005). The complexity of intersectionality. *Signs: Journal of women in culture and society*, 30(3):1771–1800.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *econometrica*, 71(6):1695–1725.

Menon, N. and Van der Meulen Rodgers, Y. (2009). International trade and the gender wage gap: New evidence from india's manufacturing sector. *World Development*, 37(5):965–981.

Neumark, D., Bank, R. J., and Van Nort, K. D. (1996). Sex discrimination in restaurant hiring: An audit study. *The Quarterly journal of economics*, 111(3):915–941.

Nguyen, A. and Lim, S. (2023). Structural transformation in the era of trade protectionism. *Available at SSRN 4426955*.

Nicita, A. (2019). *Trade and trade diversion effects of United States tariffs on China*. UN.

Niederle, M. and Vesterlund, L. (2007). Do women shy away from competition? do men compete too much? *The quarterly journal of economics*, 122(3):1067–1101.

Nunley, J. M., Pugh, A., Romero, N., and Seals, R. A. (2015). Racial discrimination in the labor market for recent college graduates: Evidence from a field experiment. *The BE Journal of Economic Analysis & Policy*, 15(3):1093–1125.

Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International economic review*, pages 693–709.

Oostendorp, R. H. (2009). Globalization and the gender wage gap. *The World Bank Economic Review*, 23(1):141–161.

Pacillo, G., Viet, C. N., Hafianti, S., Abanokova, K., Dang, H.-A., Achicanoy Estrella, H. A., and Läderach, P. (2020). Who bears the burden of climate variability? a comparative analysis of the impact of weather conditions on inequality in vietnam and indonesia. *AFD Research Paper*.

Parsons, C. A., Sulaeman, J., Yates, M. C., and Hamermesh, D. S. (2011). Strike three: Discrimination, incentives, and evaluation. *American economic review*, 101(4):1410–1435.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *The american economic review*, 62(4):659–661.

Pregelj, V. N. (2001). In *Normal-trade-relations (Most-favored-nation) Policy of the United States*. Congressional Research Service, Library of Congress.

Reich, M. (2017). *Racial Inequality: A Political-Economic Analysis*. Princeton University Press.

Rendall, M. (2017). Brain versus brawn: the realization of women's comparative advantage. Working Paper 491, University of Zurich, Institute for Empirical Research in Economics.

Reuben, E., Sapienza, P., and Zingales, L. (2014). How stereotypes impair women's careers in science. *Proceedings of the National Academy of Sciences*, 111(12):4403–4408.

Robinson, J. (1933). The theory of money and the analysis of output. *The Review of Economic Studies*, 1(1):22–26.

Robinson, J. (1969). *The economics of imperfect competition*. Springer.

Roca, J. D. L. and Puga, D. (2017). Learning by working in big cities. *The Review of Economic Studies*, 84(1):106–142.

Rodrik, D. (1999). Democracies pay higher wages. *The quarterly journal of economics*, 114(3):707–738.

Rotunno, L., Roy, S., Sakakibara, A., and Vezina, P.-L. (2024). Trade policy and jobs in vietnam: The unintended consequences of us-china trade tensions. Technical report, International Monetary Fund.

Sakakibara, A. (2023). Trade-induced structural transformation and household-level inequality: Lessons from vietnam. *Available at SSRN 4579918*.

Sarsons, H. (2017). Recognition for group work: Gender differences in academia. *American Economic Review*, 107(5):141–145.

Sauré, P. and Zoabi, H. (2014). International trade, the gender wage gap and female labor force participation. *Journal of Development Economics*, 111:17–33.

Schott, P. K. (2008). The relative sophistication of chinese exports. *Economic policy*, 23(53):6–49.

Schulze, T. and Xin, W. (2025). Demystifying trade patterns in a fragmenting world.

Seguino, S. (2000). The effects of structural change and economic liberalisation on gender wage differentials in south korea and taiwan. *Cambridge Journal of Economics*, 24(4):437–459.

Sharma, G. (2023). Monopsony and gender. *Unpublished manuscript*.

Solon, G., Haider, S. J., and 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. Technical report, The National Publishing House. Hanoi, Vietnam.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.

Topalova, P. (2007). Trade liberalization, poverty and inequality: Evidence from indian districts. In *Globalization and poverty*, pages 291–336. University of Chicago Press.

Tzioumis, K. (2018). Demographic aspects of first names. *Scientific data*, 5(1):1–9.

U.S. Department of Labor (2024). Earnings disparities by race and ethnicity. Accessed December 1, 2024. https://www.dol.gov/agencies/ofccp/about/data/earnings/race-and-ethnicity.

Utar, H. (2026). Elsewhere in north america: How us tariffs on china boosted mexico's manufacturing employment and output.

Utar, H., Zurita, A. C., and Torres, L. (2025). The us-china trade war and the relocation of global value chains to mexico. *Review of Economics and Statistics*, pages 1–47.

Verhoogen, E. A. (2008). Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. *The Quarterly Journal of Economics*, 123(2):489–530.

Vietnamese National Assembly (1994). Labour code of the socialist republic of vietnam. `https://vietnamembassy-usa.org/news/2002/10/labour-code-23-june-94`.

Vijaya, R. M. and Kaltani, L. (2007). Foreign direct investment and wages: a bargaining power approach. *Journal of World-Systems Research*, pages 83–95.

Weinberg, B. A. (2000). Computer use and the demand for female workers. *ILR review*, 53(2):290–308.

Williams, R. M. (1987). Capital, competition, and discrimination: A reconsideration of racial earnings inequality. *Review of Radical Political Economics*, 19(2):1–15.

Wolszczak-Derlacz, J. (2013). Mind the gender wage gap–the impact of trade and competition on sectoral wage differences. *The World Economy*, 36(4):437–464.

Wu, S.-J. (2024). Foreign profit shifting and the welfare responses to the us–china trade war: Evidence from manufacturers in vietnam (job market paper). Technical report, Technical report.