NAFTA and the Wages of Married Women*

Shushanik Hakobyan† John McLaren‡

International Monetary Fund University of Virginia

July 8, 2018

Abstract

Using US Census data for 1990-2000, we estimate effects of NAFTA on US wages, focusing on differences by gender. We find that NAFTA tariff reductions are associated with substantially reduced wage growth for blue-collar women far in excess of those for men, and the effect is driven mainly by *married* blue-collar women. We investigate several possible explanations for this finding. It is not explained by differential sensitivity of female-dominated occupations to trade shocks, differential elasticity of labor supply, or by household bargaining that makes married female workers less able to change their industry of employment than other workers. We find some support for an explanation based on an equilibrium theory of selective non-participation in the labor market, whereby some of the higher-wage married female workers in their industry drop out of the labor market in response to their industry's loss of tariff.

JEL Classifications: F16, F13, J31.
Keywords: NAFTA, Female and Male Wages, Local Labor Markets

*We are grateful for comments made by seminar participants at the the NBER Conference on Trade and Labor Markets, the World Bank, US International Trade Commission, Michigan State, the Pontifical Catholic University (PUC) of Rio de Janeiro, the International Trade Workshop at Duke and ELSNIT Conference. This research was supported by the Bankard Fund for Political Economy and by the Upjohn Institute. All remaining errors are ours. The views expressed in this paper are those of the authors and should not be attributed to the IMF, its Executive Board, its management, or its member country governments.

†E-mail: shakobyan@imf.org.

‡E-mail: jmclaren@virginia.edu.

1 Introduction

The effect of trade on gender inequality has not ever been a prime focus of trade economists, but gender gaps in income and labor-force participation are of intense policy and political interest, and can potentially be greatly affected by trade policy. This paper explores the impact of a major trade shock – the launch of NAFTA – on the wages of male and female workers in the US in the 1990s.

NAFTA, the most important trade policy change in the US over the last three decades, was launched on January 1, 1994 and featured a 10-year schedule of tariff phase-out between the US, Canada and Mexico. Hakobyan and McLaren (2016) examine the effect of US tariff reductions against imports from Mexico due to NAFTA on US workers' wage growth in the 1990s. The findings suggest that effects on US workers vary by educational class, location, and industry. In particular, NAFTA is associated with slower wage growth for less skilled workers employed in industries and residing in locations that were more vulnerable to imports from Mexico. But the paper did not consider the gender dimension of NAFTA's potential impact.

Much of the existing literature focuses on trade and the gender wage gap in developing countries (see Aguayo-Tellez (2012) and Papyrakis et al. (2012) for an extensive survey), but a number of studies look at the relationship in advanced economies. Black and Brainerd (2004) study the effects of increased import competition on the gender wage gap across industries and metropolitan areas in the US using the Current Population Survey (CPS) from 1977 to 1994 and the 1980 and 1990 Censuses. They find that the residual (after controlling for individual characteristics) gender wage gap narrowed more rapidly in concentrated industries that experienced a trade shock than in competitive industries, lending support to Becker's (1957) model of discrimination according to which increased market competition reduces employer discrimination in the long run.

Using US data from 1990-91 and 2006-07, Sauré and Zoabi (2014) examine the effects of a higher exposure to trade with Mexico on female employment shares and female relative wages across US states, and find that trade expansion had a negative impact on female employment relative to male in states with greater exports to Mexico. The results remain robust for married female workers, for female workers of all educational categories (less than high school, high school graduate and advanced education), and for workers in manufacturing only. They do not find a significant difference in female relative wage due to higher trade exposure, attributing it to a selection bias whereby the measured average wages of working women do not change, while the unmeasured potential wages of nonworking women decrease (as they leave labor force). Autor et al. (2013) examine the impact of rising Chinese import competition on U.S. labor market outcomes over the period of 1990-2007 and find that both male and female employment and the corresponding wages decreased but that these changes were more pronounced for women. Brussevitch (2018) shows that some portion of the declining gender wage gap in the US can be explained by differential labor adjustment costs. She estimates a structural econometric model of dynamic labor adjustment and finds that, although men tend to have overall lower adjustment costs than women, women have an advantage in moving into service-sector jobs following a shock to traded-sector labor demand. None of these papers however addresses the differences in income distribution by gender, marital status, education, industry of employment and location simultaneously, which is the focus of this paper.

Studies of trade and the gender wage gap in developing countries are more common and tend to conclude that trade liberalization improved labor market outcomes of women. Aguayo-Tellez et al. (2010) document increased female employment rates and female relative wages in Mexico during the 1990s due to NAFTA, and using establishment-level data for the Mexican manufacturing sector, show that female wage bill shares increased in response to reductions in US tariffs on Mexican goods, particularly for skilled blue-collar female workers. Using the same data from Mexico, Juhn et al. (2014) show that tariff reductions due to NAFTA increased the ratio of female blue-collar workers to male blue-collar workers as well as the relative wage of female blue-collar workers, with little evidence for white-collar women's share and relative wages. Gaddis and Pieters (2014) find that trade liberalization in Brazil reduced male and female labor force participation and employment rates, but the effects on men were significantly larger, leading to gender convergence in labor force participation and employment rates.

This paper borrows from several advances in the literature and builds on the methodological framework developed in Hakobyan and McLaren (2016) to study the differential impact of tariff reductions on men's and women's wage growth and labor force participation decisions over the 1990s by exploiting the exogenous nature of the NAFTA shock. We use publicly available US Census data from 1990 and 2000, taken from the IPUMS project at the Minnesota Population Center (www.ipums.org; see Ruggles et al. (2015)). The richness of our data allows us to estimate the differential impact of a trade shock within a location, industry, occupation and educational class.

To anticipate results, we find that reductions in blue-collar wage growth from NAFTA tariff reductions were much larger for women than for men, and much larger for married women than for single women, a finding that we refer to as a *married-woman differential*. We investigate four possible explanations for this finding: differential sensitivity to shocks across occupations; differential elasticity of labor supply; household bargaining within a marriage; and selective non-participation in the labor market on the part of married women. We are able to reject the first three with the data; the fourth appears to be plausibly a portion of the explanation, but it is unable to explain the full effect. We therefore conclude with a puzzle.

The rest of the paper proceeds in the following way. Section 2 briefly explains the methodological framework developed in Hakobyan and McLaren (2016), presents the basic results for the wage growth over 1990s for men and women, showing that they are robust to a wide range of specification changes, and then shows that the differential is driven mainly by married women. Section 2 concludes by laying out three stylized facts. Section 3 proposes four possible explanations for our basic findings in Section 2 and develops a simple theoretical model for each explanation followed by further empirical tests of the proposed theory. Section 4 considers a model of endogenous marriage and argues that it is not the reason for the married-woman differential. Finally, Section 5 summarizes.

2 Empirical approach and basic results

Our analysis of local labor markets requires a time-invariant definition of local labor markets in the US. We take advantage of consistently defined Public Use Microdata Areas (conspumas) constructed by and available from IPUMS-USA (Ruggles et al., 2015). There are 543 conspumas covering both urban and rural areas in the US.

We begin with the hypothesis that male and female workers are treated identically in the labor market, and so in our initial specification we treat them symmetrically. Following the empirical specification in Hakobyan and McLaren (2016), we construct a measure of conspuma's exposure to NAFTA as 1990 employment-share-weighted US tariff rate (against Mexican imports) in 1990, adjusted for Mexico's relative comparative advantage. The latter is defined as:

$$RCA^j = \frac{\left(\frac{x^{MEX}_{j,1990}}{x^{ROW}_{j,1990}}\right)}{\left(\frac{\sum_i x^{MEX}_{i,1990}}{\sum_i x^{ROW}_{i,1990}}\right)},$$

where respectively $x^{MEX}_{j,1990}$ is Mexico's and $x^{ROW}_{j,1990}$ is the rest of the world's exports of industry j in 1990 (excluding exports to the US). Clearly RCA^j exceeds unity if industry j is an industry that Mexico tends to export disproportionally. We weight tariffs by this measure on the reasoning that if Mexico does not export a product, the tariff on imports of that product from Mexico should be irrelevant. We use this as part of our measure of average local tariff or local vulnerability:

$$loc\tau^c_{1990} \equiv \frac{\sum_{j=1}^{N_{ind}} L^{cj}_{1990} RCA^j \tau^j_{1990}}{\sum_{j=1}^{N_{ind}} L^{cj}_{1990} RCA^j}, \tag{1}$$

where L^{cj}_t is the number of workers employed in industry j at conspuma c at date t, and

N_{ind} is the number of industries.[1]

The variation in this measure comes from three sources: differential concentration of employment across industries in each conspuma; specialization in industries in which Mexico has comparative (dis)advantage relative to the rest of the world; and the US-imposed differential tariff rates. Analogously, we define industry tariff rates, adjusted for Mexico's relative comparative advantage.

The use of the Census data collected in 1990 and 2000 dictates our empirical approach to identifying the effect of NAFTA. The agreement was framed as a gradual phase-out of tariffs between the three countries, starting in 1994 and continuing for 10 years (with a few tariffs continuing to 15 years). We focus on exposure to Mexican imports at the time of NAFTA's launch because the reduction of tariffs between the US and Canada had began much earlier with signing of a free trade agreement between the two countries in 1989 (Appendix Figures A1 and A2). The negotiated schedule of liberalization was different for each sector of the economy. As a result, for some industries, the period from 1990 to 2000 was the period of an announcement of tariff reductions, most of which occurred after 2000. For other industries, the same period saw rapid elimination of tariffs. As a result, we deal with variation in the timing of liberalization by controlling separately for both the initial tariff rates in 1990, which capture the potential vulnerability of a location or an industry to imports from Mexico, and actual change in tariffs between 1990 and 2000.[2]

Aside from the dynamics of the tariff cuts, a number of other considerations complicate the measurement of tariff effects. We found in our previous study that the wage effects of

[1] Kovak (2013) analyzes a model of local labor markets differentially affected by trade shocks which motivates an estimating equation quite similar to what we use. He discusses the tricky question of how to treat non-traded industries. The suggestion that emerges from his analysis is to use only the traded-industry employment when computing labor shares for local-average tariffs. Our formulation is compliant with that approach, because the RCA^j is set to zero for any non-traded industry.

[2] The initial tariffs and changes in tariffs are correlated but far from perfectly, and their coefficients are generally separately identified in the data. As a matter of theory, it is entirely possible that the effect of reducing a tariff from 6% to 3% over the sample period will be different from the effect of reducing a tariff from 3% to zero over the same period, even though those are both 3-percentage-point reductions, because in the first case workers and firms anticipate a further reduction in the near future but in the second case they do not. Therefore, failure to control for initial tariffs could create an omitted-variables bias. In robustness exercises below we show that the results are similar if the initial tariff is omitted.

the tariff reductions are very different for different educational groups, with substantial reductions in wage growth for blue-collar workers in vulnerable industries or locations but no effect for college-educated workers. Since these differences are crucial to understanding the welfare effects of the agreement, we keep those educational categories separate throughout, distinguishing between high school dropouts, high school graduates, workers with some college or associate degree, and college graduates. We have also found a strong and separately identified response of wages for each group of workers by industry and location, so we keep track of those distinctions in our estimation. We also allow for a different rate of wage growth for locations on the US-Mexico border. These considerations lead to a regression equation as follows:

$$\begin{aligned} log(w_i) &= \alpha X_i + \sum_j \alpha_j^{ind} ind_{i,j} + \sum_c \alpha_c^{conspuma} conspuma_{i,c} + \sum_n \alpha_n^{occ} occ_{i,n} \\ &+ \sum_{k \neq col} \gamma_{1k} educ_{ik} + \sum_k \gamma_{2k} educ_{ik} yr2000_i \\ &+ \sum_{k \neq col} \delta_{1k} educ_{ik} loc\tau_{1990}^{c(i)} + \sum_k \delta_{2k} educ_{ik} yr2000_i loc\tau_{1990}^{c(i)} \\ &+ \sum_{k \neq col} \delta_{3k} educ_{ik} loc\triangle\tau^{c(i)} + \sum_k \delta_{4k} educ_{ik} yr2000_i loc\triangle\tau^{c(i)} \\ &+ \sum_{k \neq col} \theta_{1k} educ_{ik} RCA^j \tau_{1990}^{j(i)} + \sum_k \theta_{2k} educ_{ik} yr2000_i RCA^j \tau_{1990}^{j(i)} \\ &+ \sum_{k \neq col} \theta_{3k} educ_{ik} RCA^j \triangle\tau^{j(i)} + \sum_k \theta_{4k} educ_{ik} yr2000_i RCA^j \triangle\tau^{j(i)} \\ &+ \mu Border_{c(i)} yr2000_i + \epsilon_i, \end{aligned} \tag{2}$$

where $conspuma_{i,c}$, $ind_{i,j}$ and $occ_{i,n}$ are dummy variables that take a value of 1 if worker i resides in conspuma c, works in industry j and has an occupation n, respectively; $c(i)$ is the index of worker i's conspuma, and $\triangle\tau^{j(i)}$ and $loc\triangle\tau^{c(i)}$ are the changes in tariff for industry j and location c, as defined at the beginning of this section.

The parameters of primary interest here are $\delta_{2,k}$ and $\delta_{4,k}$, which measure the initial-tariff effect and the impact effect, respectively, for the local average tariff; and $\theta_{2,k}$ and $\theta_{4,k}$, which measure the initial-tariff effect and the impact effect, respectively, for the industry tariff. If it is easy for workers to move geographically, so that local wage premiums are arbitraged away, but difficult for workers to switch industry, we will observe $\delta_{1,k}, \ldots, \delta_{4,k} = 0$ while $\theta_{1,k}, \ldots, \theta_{4,k} \neq 0$. In that case, industry matters, but location does not. On the other hand, if it is difficult for workers to move geographically but easy to switch industries within one location, we will see the opposite: $\delta_{1,k}, \ldots, \delta_{4,k} \neq 0$ while $\theta_{1,k}, \ldots, \theta_{4,k} = 0$. In reporting our results, we focus on the net effect when a *location* or an *industry* loses all of its protection within the sample period, thus the effect on wages within the sample period is equal to $\delta_{2,k} - \delta_{4,k}$ in a given location and $\theta_{2,k} - \theta_{4,k}$ in a given industry.

In the regressions below, we use a 5% sample from the US Census for 1990 and 2000, available from IPUMS-USA, selecting workers from age 25 to 64 who report a positive pre-tax wage and salary income in the previous year.[3] The sample includes workers from all sectors, including non-traded service-sector workers. In addition to constructed interaction terms and conspuma, industry and occupation fixed effects, we include the following personal characteristics: age, whether or not the worker speaks English, race, home ownership, presence of a school-aged child and educational attainment.

Table 1 shows descriptive statistics for the personal characteristics of workers by gender. In our sample, male and female workers have quite similar characteristics, with average age 41, 81-82% white, 71-72% home owners. Women are slightly less likely to have dropped out of high school and to have completed a four-year college degree, and they have slightly lower average wages.

Table 2 summarizes our measures of industry and location vulnerability. The 1990 RCA-adjusted industry tariff across 89 traded-goods industries ranges from 0 to 9%, with a mean of 1%. The initial local average tariff across 543 conspumas ranges from approximately 0.09 to 4.74%, with a mean just above one percent. It is worth pointing out that when computing

[3]The sample includes individuals who report being employed, unemployed or not in labor force in the census year. We use the last industry of employment for the unemployed and those not in labor force. Wage/income regressions omit those workers with no reported wage/income. Labor force participation regressions include all workers in the sample.

Table 1: Summary statistics by gender

	Male	Female	Total
Age	41.24	41.18	41.21
White	0.82	0.81	0.82
English speaking	0.991	0.993	0.992
Home owner	0.73	0.72	0.72
Child(ren)	0.68	0.71	0.41
High school dropouts	0.12	0.09	0.11
High school graduates	0.31	0.32	0.33
Some college	0.29	0.32	0.30
College graduates	0.29	0.27	0.27
Log wage income	10.2	9.59	9.91
N of Observations	5,366,329	4,862,010	10,228,339

local average tariff we omit agriculture by setting its tariff equal to zero because a coarse aggregation of industries in Census data applies the same tariff to all agricultural crops.[4]

Table 3 shows the difference between the initial-tariff effect and the change-in-tariff effect for the main specification in equation (2) with all right-hand-side variables and industry, conspuma and occupation fixed effects, run separately for men and women. The first two columns record the results with the worker's wage as the dependent variable; the next two columns capture the effect on employment (the dependent variable is a dummy variable that takes the value of 1 if the worker is employed); and the last two columns capture the effect on labor-force participation (the dependent variable is a dummy that takes the value of 1 if the worker is in the labor force). The first four rows display the net effect for the location variables, and the last four for the industry variables. Standard errors are clustered by conspuma, industry, and year, following Cameron, Gelbach and Miller (2006). The coefficients on personal characteristics generally have the expected signs and are not reported here.

[4]NAFTA was likely to be highly beneficial to US corn producers, but a serious blow to a wide range of fruit and vegetable producers. These, and all other crops, are lumped together in the useless category of 'Agriculture: Crops.' For further discussion, see Hakobyan and McLaren (2016).

Table 2: Summary Statistics for Industry and Local Average Tariffs

Variable	Mean	St. Dev.	Min	Max	N
Industry Tariff in 1990 (%)	1.0	2.0	0	8.8	89
Change in Industry Tariff (%)	-0.9	1.6	-7.0	0.01	89
Local Tariff in 1990 (%)	1.03	0.67	0.09	4.74	543
Change in Local Tariff (%)	-0.92	0.61	-4.30	-0.08	543

Notes: Industry level tariff variables are computed from 8-digit HS tariff data weighted by imports from Mexico and are mapped into 89 tradeable goods industries based on Census industry classification. RCA is Mexico's revealed comparative advantage in a particular industry as defined in the text. Conspuma level variables are weighted by employment in industries of a given conspuma.

Consider first the wage effects in the first two columns. Examining the location effect, the first row of Table 3 shows that among conspumas that lost their protection quickly under NAFTA, those that appeared to be very vulnerable had substantially lower wage growth for female high-school dropout workers than those with low initial tariffs. In particular, female workers with less than high school education, living in the most vulnerable conspuma with an initial local average tariff of 4.74 percent, would see a substantial drop in wage growth over the 1990s of around $4.74 \times 3.210 \approx 15$ percentage points. In a conspuma with mean initial tariff of 1.03 percent the drop in wage growth would be about 3.3 percentage points. However, for male workers with less than high school education the effect is one third as large. For workers with higher levels of educational attainment, the effect is smaller or insignificant, but where it is significant it is consistently substantially larger for female workers, particularly for the least educated. The industry variables in the bottom four rows show a similar pattern; for example, the effect for high-school dropout women is more than twice as large as that for men.

The next two columns show that a drop in the local average tariff is associated with a drop in employment in the conspuma for both men and women, for each education level, but always considerably larger for women. In the case of high-school dropouts, the effect is more than twice as large for women. The last two columns show that most of this effect can

Table 3: Labor Market Outcomes by Gender

	Wage		Employment		Labor force	
	Male	Female	Male	Female	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)
Location effect						
Less than high school	-1.086*	-3.210***	-1.080***	-2.463***	-0.350	-1.878***
High school graduate	-0.478	-0.984	-1.018***	-1.618***	-0.781***	-1.549***
Some college	-0.960**	-1.394*	-0.650***	-1.263***	-0.596***	-1.181***
College graduate	0.014	-1.546*	-0.417***	-0.677***	-0.405***	-0.692***
Industry effect						
Less than high school	-1.277***	-3.14**	-0.479	-0.53	-0.165	-1.038
High school graduate	-0.296	-2.327***	-0.274	-0.663	-0.271*	-1.001**
Some college	-0.146	-0.973	-0.126	-0.481	-0.275**	-0.877**
College graduate	-0.161	1.102	-0.168	-0.19	-0.21*	-0.475
N of Observations	5,366,329	4,862,010	6,165,584	5,789,584	6,165,584	5,789,584

Notes: The table reports the overall impact on male and female labor market outcomes (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. In Columns (1) and (2), the dependent variable is the log wage. Columns (3)-(6) report the estimates from a linear probability model of a worker being employed or in labor force. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

be accounted for by a drop in the population that is in the labor force, as opposed to transitions to unemployment. Again, the effects are consistently larger for women. The industry employment and labor-force effects are smaller across the board and often inconsistent.

To summarize, we find that the growth of wages of female workers in NAFTA-vulnerable locations and industries was substantially slower over the period in question compared to other female workers after controlling for a range of personal factors, as was the growth of the female labor force in those locations and industries; and that these effects were much stronger than the corresponding effects for men. For the least-educated women, the wage effect was as much as triple that for men.

2.1 Robustness exercises

In Tables 4 and 5, we explore various robustness exercises to ensure that the differentials we see in the main regression are not an artifact of our specification.

Simplifying the regression. One might wonder about multicollinearity between the local tariffs and the industry tariffs. In columns (1) and (2) of Table 4 we run the main regression without the industry tariffs, and in columns (3) and (4) without the local tariffs. The results are quite similar to the main regression, and the finding of stronger effects for women is preserved.

A related issue is that the initial tariffs and the changes in tariffs are highly correlated both across industries and across locations. This might lead to questions of multicollinearity, so we also run the main regression without the interaction terms with initial tariffs (columns (5) and (6) of Table 4). The main finding of a stronger effect for women is strengthened. The estimates for women in column (6) are larger and more statistically significant than the ones for men in column (5). In many cases the women's estimate is triple the size of the men's.[5]

Agricultural tariffs. We noted above that we omit agricultural tariffs because the coarseness of the Census categories make them all but useless for our purposes, and because the agricultural sector is the one sector where the rapid expansion of export opportunities is most of the story. By imposing the same tariff for different crops and localities with widely different experiences of NAFTA, this coarseness biases the results toward zero. We run the regression with the agricultural tariffs included anyway, reporting the results in columns (7) and (8) of Table 4. The results are all insignificant. In the last two columns of Table 4, we omit the agricultural workers altogether. The industry effects are essentially identical to the main results, but the location effects are gone, indicating that an important source of identification in the main specification is agricultural workers in locations with some

[5]Since this is the estimate for the change in tariff, a positive sign indicates that a drop in the tariff correlates with a drop in the wage, the reverse for the main regression. The economic interpretation is the same.

NAFTA-vulnerable manufacturing compared with agricultural workers in other locations. We continue to set agricultural tariffs to zero going forward.

The China shock. Our sample period predates the explosion of imports from China following its accession to the WTO, but Chinese exports had begun to grow in the 1990s, so it is natural to ask if the reduction in tariffs that is our focus might be correlated with that trend. In columns (1) and (2) of Table 5, we include change in share of imports from China as an additional control, handled in a way analogous to the tariffs (both industry imports and the employment-weighted average for each location, for example). The results are essentially the same as in Table 3. (We do not include the estimates of the China effect itself in the table since they are outside of our area of inquiry.)

Mexican tariffs. Of course, NAFTA lowered Mexican tariffs on US goods as well as US tariffs on Mexican imports, and the two could be correlated. In columns (3) and (4) of Table 5, we include the Mexican tariffs on US goods as controls, handled in a way analogous to the treatment of the US tariffs. The results are essentially the same as in Table 3. (We do not include the Mexican tariff coefficients in the table since they are outside of our area of inquiry.)

Placebo regression. The effects we are picking up could merely be the continuation of pre-existing trends in wages that are correlated with tariffs. Columns (5) and (6) of Table 5 report the results from a regression of wages from 1980 to 1990 on the tariff measures from our main regression (that is, from 1990 to 2000). Clearly, the results do not resemble the main regression results at all. The estimates for female workers are generally much smaller, in four cases with the opposite sign, and in the case of location effects, the male effects are much *larger* than for women. Clearly, our results do not come from a pre-existing trend.

Additional robustness exercises. To be sure that our results are not driven by the way we measure our dependent variable, or how we select the sample of workers, we estimate the same regression replacing the dependent variable with self-employment income for those with no wage income; replacing the dependent variable with weekly wage; excluding workers over

Table 4: Robustness Checks: Tariff Measures

	Only location effects		Only industry effects		No initial tariff		Includes agriculture		Excludes agric workers	
	Male	Female	Male	Female	Male	Female	Male	Female	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Location effect										
Less than high school	-1.2**	-3.77***			0.453*	1.985***	-0.184	-0.76	-0.128	-0.81
High school graduate	-0.495	-1.392**			0.236**	0.706**	0.236	0.833	0.274	0.8691
Some college	-0.964***	-1.473*			0.659***	1.538***	0.393	0.596	0.375	0.605
College graduate	0.023	-1.513**			0.447	1.738**	-0.405	-0.644	-0.45	-0.635
Industry effect										
Less than high school			-1.359***	-3.475**	0.904**	3.035***	1.138	-0.575	-1.137***	-3.211**
High school graduate			-0.326	-2.448***	0.188	1.984***	0.361	-1.49	-0.301	-2.555***
Some college			-0.219	-1.123*	0.265**	0.826*	0.566	-0.564	-0.206	-1.164
College graduate			-0.155	1.023	0.233	-0.314	0.379	0.929	-0.098	1.033
N of Observations	5,366,329	4,862,010	5,366,329	4,862,010	5,366,329	4,862,010	5,366,329	4,862,010	5,267,803	4,831,356

Notes: The dependent variable is the log wage. Columns (1)-(4) and (7)-(10) report the overall impact on male and female wages (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. Columns (1)-(2) control for only location effects, Columns (3)-(4) - only industry effects. Columns (5)-(6) report the results from a regression that only controls for change in tariffs and omits the initial tariffs. Columns (7)-(8) include agricultural tariffs which are set to zero in other columns. Columns (9)-(10) restrict the sample to non-agricultural workers. The regression results used to compute these point estimates are available from the authors upon request. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

Table 5: Robustness Checks: Omitted Variables and Placebo Regression

	China effect		Mexico effect		Placebo regression	
	Male	Female	Male	Female	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)
Location effect						
Less than high school	-0.898	-3.47***	-1.331*	-3.65***	7.41***	1.69***
High school graduate	-0.494	-1.411**	-0.416	-0.62	5.27***	0.88
Some college	-1.104*	-2.118***	-0.533	-1.133	6.37***	1.09
College graduate	-0.245	-2.345***	0.373	-0.508	4.52***	-0.63
Industry effect						
Less than high school	-1.47***	-3.294**	-1.072***	-2.634*	-0.145	-0.976**
High school graduate	-0.382*	-2.449***	-0.241	-2.361***	-0.212	-1.585***
Some college	-0.117	-0.981	-0.011	-0.75	-0.223	-1.004*
College graduate	-0.188	0.97	0.291	1.938*	0.452	3.335**
N of Observations	5,366,329	4,862,010	5,366,329	4,862,010	4,692,440	3,903,420

Notes: The dependent variable is the log wage. The table reports the overall impact on male and female wages (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. Additional controls include change in share of imports from China in Columns (1)-(2), and initial Mexican tariffs on US goods in 1990 and change in Mexican tariffs between 1990 and 2000 in Columns (3)-(4). Regressions in Columns (5)-(6) use 1980 and 1990 Census data. We omit agriculture by setting its tariff equal to zero. The regression results used to compute these point estimates are available from the authors upon request. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

55 years old; and excluding workers with spouses younger than 25 and older than 64. The results reported in Appendix Tables A1-A4 continue to be in line with the earlier findings in Table 3.

Conclusion. We conclude that our basic results are not driven by multicollinearity, a correlation with an omitted variable, a pre-existing trend, or a peculiarity of how we have measured wages. The NAFTA effect for blue-collar wage growth is much stronger for female workers than for men.

Table 6: Wage Growth by Gender and Marital Status

	Married workers		Single workers	
	Male	Female	Male	Female
	(1)	(2)	(3)	(4)
Location effect				
Less than high school	-0.99	-3.8***	-1.392*	-1.58
High school graduate	-0.643	-0.777	-0.509	-0.698
Some college	-1.113**	-1.56*	-0.877*	-0.442
College graduate	-0.19	-1.935*	0.032	-0.245
Industry effect				
Less than high school	-0.944*	-3.838***	-1.656***	-1.748
High school graduate	-0.113	-3.016***	-0.53	-0.435
Some college	-0.072	-1.598*	0.049***	0.612
College graduate	0.187	1.59	-1.0722	0.914
N of Observations	3,803,221	3,117,140	1,563,108	1,744,870

Notes: The table reports the overall impact on male and female wages (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

2.2 Splitting the sample by marriage

Having established that there is a large difference in wage responses between male and female workers, we wish to identify the reason. A first clue can be obtained by splitting the sample between married and unmarried workers. In Table 6, we run the main regression but in columns (1) and (2) the sample is restricted to married workers, while in columns (3) and (4) it is restricted to single workers. Of course, married or single status can be endogenous to trade policy. We discuss this endogeneity of marriage at length in Section 4. Our purpose here is purely statistical: To see if the gender differences we have identified above appear with equal force in the different subsamples as a clue to identifying the cause of the differential.

The effects for married women (column (2) of Table 6) are stronger than for women on average (column (2) of Table 3), while effects for single women (column (4) of Table 6) are considerably smaller and insignificant. Further, there is barely any difference between the male and female wage effects for single workers (columns (3) and (4) of Table 6).[6]

We conclude that the male-female differential response we have observed above is driven primarily by the differential effect of tariff changes on the wages of married women.

2.3 Main findings

To sum up the results so far, we find that: (1) There is a much stronger effect of NAFTA tariff reductions on blue-collar women's wages than men's – two to three times, depending on how it is measured. (2) In most specifications there is no real difference between the wage responses of single men and women. (3) There is, however, a large difference between the wage responses of married men and women. We call this finding the *married-woman differential.* Explaining this differential is our main goal.

Before proceeding, we should underline some limitations in our exercise. First, since we do not have panel data, but only two cross sections, we are unable to measure precisely the effect of the tariff changes on the wage growth for a given worker, holding constant the worker's time-invariant unobservable characteristics. Some portion of the changes in average wage reported in Table 3 may be due to changes in individual wages while some portion results from changes in the composition of the demographic group in question due to selection effects. In investigating possible explanations, these competing mechanisms will

[6]To test the difference in tariff effect across genders formally, we run an alternative version of the regression with male and female workers pooled, and additional interaction terms of the tariff measures, industry, location and occupation fixed effects with a female dummy added on the right hand side of (2). The results are detailed in Appendix Table A5, first for the married sample and then for the single sample, including results for the test of the hypothesis that the tariff coefficients are identical across genders. The estimates look similar in most respects to the main specification with a few differences. The difference between the gender coefficients is significant in three of the four blue-collar cases for married workers, and in one out of four for single workers. Both for the location effect and the industry effect, the largest effect by a wide margin is for married women with the least education, and in most categories the point estimate for married women is larger in magnitude than for single women. It should be noted that this specification imposes identical coefficients for male and female workers for all of the personal characteristics other than gender and marriage. These may vary quite a bit across genders. Consequently, this pooled regression implicitly imposes parameter restrictions for which there is no justification, and we trust the main specification, with separate regressions for male and female workers, more.

be at the forefront of the discussion. Second, in order not to lose focus on our main agenda of explaining the married-woman differential we are staying away from the related question of measuring the effect of tariff changes on the male-female wage gap.

3 Search for explanations

Below we investigate four different possible explanations for the results presented in previous section.

(i) *Heterogeneous occupations.* It could be that different occupations have different levels of sensitivity to industry-level trade shocks, for example because the cost of inter-industry mobility differs across occupations. If women are over-represented in the more sensitive occupations, that can lead to a larger wage effect on average for female workers than for men.

(ii) *Differential elasticity of labor supply.* If men and women are imperfect substitutes to employers and if they have different elasticities of labor supply, either at the industry level or at the level of a conspuma, then the same trade-related shock to labor demand can produce differential wage effects across genders.

(iii) *Household bargaining.* It could be that married women are less mobile than other workers because switching industries sometimes requires switching city of residence, which is a joint decision with her spouse. We investigate the possibility that if a husband has more bargaining power than a wife, this can result in asymmetries in moving frictions that result in larger wage impacts for married women than for single workers or married men. This can lead to differential mobility for married and single women.

(iv) *Selective non-participation.* It is possible that when an industry shrinks due to a trade shock that a certain fraction of workers choose to leave the labor force. If those leavers are disproportionally married women, and disproportionally the higher-paid workers in their industry, the selection effect can result in a larger drop for average wage for the remaining married female workers in the industry, compared to other groups. We present an equilibrium model in which exactly this prediction emerges.

3.1 Heterogeneous occupations

3.1.1 Theory

Occupations vary greatly in the gender composition of their workers, with some occupations dominated by female workers and others by men. As one example, 'textile sewing machine operators,' an occupation with more than a million workers in our dataset, has 10 female workers for every male worker. If occupations also differ in the portability of skills across industry, with some occupations very mobile across industries and others immobile, then it could be that female-dominated occupations happen to be, on average, less mobile across industries. This would imply a larger wage response to a trade shock for female workers on average even if all genders are treated equally.

A simple example can illustrate the point. Suppose that there are two industries indexed $i = 1, 2$ and two occupations indexed $j = 1, 2$. Production in each industry requires labor input from both occupations, so output of industry i is given by a concave linear homogeneous production function $f^i(L_1^i, L_2^i)$, where L_j^i is the number of workers in occupation j employed by industry i. Suppose that each worker is attached to an occupation and cannot change it.

To capture the idea that different occupations can have different degrees of mobility in a simple way, suppose that workers in occupation 2 cannot change their industry of employment, but workers in occupation 1 can change their industry freely. Perhaps occupation 2 requires mastering a particular part of a production process with particular machines that differ from one industry to another and so the skills required for it are not portable across industries (sewing machines, for example, are not useful outside of the apparel industry); while occupation 1 requires general production-floor activities that are similar across industries. Suppose that the price of output from both industries is given on world markets (for simplicity, assume that the economy in question is a small open economy), but the domestic price can differ from the world price due to trade policy. Letting good 2 be the numeraire, suppose that industry 1 is import-competing, and its domestic price, p, is equal to the world price plus an import tariff. All agents take all prices as given.[7]

Since occupation 1 is mobile, the wage w_1 paid to it must be the same in both industries. Since this will be equal to the marginal value product of labor, we have:

$$pf_1^1(L_1^1, L_2^1) = w_1 = f_1^2(L_1^2, L_2^2) = f_1^2(L_1 - L_1^1, L_2^2), \tag{3}$$

where subscripts on a function indicate partial derivatives and L_1 is the exogenous and fixed supply of workers in occupation 1. This determines the allocation of occupation-1 workers across the two industries, and also w_1. Further, the occupation-2 wages in the two industries must adjust to yield zero profits in both industries:

$$c^1(w_1, w_2^1) = p, \text{ and} \tag{4}$$

$$c^2(w_1, w_2^2) = 1, \tag{5}$$

where $c^i(\cdot)$ denotes the unit cost function for industry i and w_2^i is the occupation-2 wage in industry i.

Differentiating (3) with respect to p, allowing L_1^1 to adjust, shows that $\frac{dL_1^1}{dp} > 0$, so a reduction in the tariff on industry 1 will move labor from industry 1 to 2. This will reduce w_1 (from the industry 2 first-order condition) and increase $\frac{w_1}{p}$ (from the industry-2 first-order condition). If we write the elasticity of a variable X with respect to Y as ϵ_{XY}, then this implies:

$$0 < \epsilon_{w_1 p} < 1. \tag{6}$$

Differentiating the two zero-profit conditions then implies that a drop in the tariff will

[7] This simple structure gives the model the same form as the Ricardo-Viner model of trade (Jones, 1971).

require a more-than-proportional drop in w_2^1 to restore industry-1 zero profits, and an increase in w_2^2 to restore industry-2 zero profits:

$$\epsilon_{w_2^2 p} < 0 < 1 < \epsilon_{w_2^1 p}. \tag{7}$$

Conditions (6) and (7) together imply that the wage response for the immobile occupation in the import-competing industry will be much larger than for the mobile occupation. If it so happens that women are concentrated in occupation 2 and men in occupation 1, then a larger wage effect will be measured for female workers whose industry tariff is reduced than for men. This is true even with industry and occupation fixed effects, because the fixed effects will control for differences in the level of wage, not differences in the elasticity of wage with respect to the tariff change. We can now ask whether or not this story is consistent with the evidence.

3.1.2 Empirical test

To test whether the findings are driven by differential response of female-dominated occupations, we construct a dummy for female-dominated occupations and interact it with the industry and local tariff variables. To identify female-dominated occupations, we compute the ratio of women to men in each occupation in 1990. The ratio ranges from 0.01 (*Bus, truck, and stationary engine mechanics* – a highly male-dominated occupation) to 101 (*Secretaries* – a highly female-dominated occupation). Our dummy for female-dominated occupations takes the value of 1 if this ratio is greater than five, in other words the number of women in a given occupation is five times that of men in 1990, and zero otherwise.[8] Table 7 lists all such female-dominated occupations.

We add the dummy for female-dominated occupations to our main specification in equation (2) by interacting it with our industry and local tariff measures and year-2000 dummy.

[8]The ranking of occupations by female-to-male ratio barely changes when we use our entire sample or only 2000 Census.

Table 7: Top female-dominated occupations in 1990

Occupation	Ratio	Number of women
Secretaries	101.4	3,851,569
Dental hygienists	62.0	72,233
Kindergarten and earlier school teachers	60.0	256,903
Dental assistants	42.3	156,596
Receptionists	36.8	647,715
Child care workers	31.6	708,023
Home economics instructors	22.6	429
Typists	21.7	576,082
Private household cleaners and servants	20.3	342,895
Teacher's aides	20.2	527,236
Registered nurses	18.0	1,841,392
Dressmakers and seamstresses	17.6	99,349
Licensed practical nurses	16.1	418,852
Bank tellers	16.1	372,053
Health record tech specialists	14.8	48,605
Speech therapists	12.9	63,613
Dietitians and nutritionists	10.9	84,485
Bookkeepers and accounting and auditing clerks	10.8	1,706,530
Billing clerks and related financial records processing	10.0	181,137
Textile sewing machine operators	9.9	748,830
Stenographers	9.6	71,826
Eligibility clerks for government programs; social welfare	9.4	44,392
Data entry keyers	8.7	488,791
Hairdressers and cosmetologists	8.7	600,769
Payroll and timekeeping clerks	8.6	158,888
Nursing aides, orderlies, and attendants	8.4	1,634,812
Occupational therapists	7.9	33,858
Telephone operators	7.8	193,031
Sales demonstrators / promoters / models	7.7	42,690
Library assistants	7.1	84,999
Crossing guards and bridge tenders	6.7	33,675
Human resources clerks, except payroll and timekeeping	6.5	66,110
Kitchen workers	6.2	132,809
Welfare service aides	6.1	41,980
General office clerks	6.0	1,107,735
File clerks	5.9	157,802
Waiter/waitress	5.9	880,093
Housekeepers, maids, butlers, stewards, and lodging quarters cleaners	5.7	657,273
Cashiers	5.6	1,518,375
Special education teachers	5.1	50,671
Librarians	5.0	154,557

Table 8: Controlling for female-dominated occupations

	Married workers		Single workers	
	Male	Female	Male	Female
	(1)	(2)	(3)	(4)
Location effect				
Less than high school	-1.00	-3.878***	-1.645	-1.73
High school graduate	-0.65	-0.087	-0.491	-0.697
Some college	-1.112*	-0.996	-1.029	-0.78
College graduate	-0.193	-2.885***	-0.136	-1.153*
Industry effect				
Less than high school	-0.875*	-2.19*	-1.62*	-0.792
High school graduate	-0.082	-2.665**	-0.5	0.264
Some college	-0.032	-1.917	0.121	0.775
College graduate	0.17	1.875	-1.103	0.733
N of Observations	3,803,221	3,117,140	1,563,108	1,744,870

Notes: The table reports the overall impact on male and female wages (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

The summary results are reported in Table 8 analogous to Table 3. It is clear that the results are not affected in any substantive way after controlling for female-dominated occupations. We conclude that the differential effects of NAFTA by gender are *not* caused by the different occupational mixes shown by male and female workers.

3.2 Differential elasticity of labor supply

Another simple explanation for the differential wage effect is that men and women have different labor-supply elasticities, so that a similar trade-policy-induced shock to labor demand in a given industry reduces employment in that industry more, and wages less, for male workers than for female workers.[9] To evaluate this, first we need to note that it requires that male and female workers be regarded as imperfect substitutes by employers, even after controlling for industry and occupation. Otherwise, wages for both genders would move up or down together by an amount given by the combined labor-supply elasticity. This hypothesis can therefore be thought of as having two pieces: (i) imperfect substitutability between male and female workers, and (ii) a smaller elasticity of supply for female workers.

One way of testing (i) is to compute gender-specific measures of the tariffs and see if they have differential effects on male and female wages, where the local average male (female) tariff is the weighted average tariff (interacted with revealed comparative advantage as always), with weights given by the share of male (female) workers by industry as a fraction of total local employment. A disproportionate drop in the average male tariff indicates industry tariffs have fallen disproportionately in male-dominated industries.[10] If the two genders are perfect substitutes, a change in either gender-specific tariff will have the same effect on wages of either gender. Table 9 replicates the regressions from Table 3 but replaces the common local tariffs with the corresponding gender-specific local tariffs. The results are much stronger than in the main regression with common local tariffs for men and women, showing strong negative effects even for the wages of male workers. The difference between this regression and the main regression of the first two columns of Table 3 implies that the gender-specific tariffs matter and that male and female workers are likely not perfect substitutes.

However, if hypothesis (i) has some support, hypothesis (ii) does not. Recalling Table 3, columns (3) through (6) show that both the employment effect and the labor force effect of a given tariff change are much larger for women than for men, except for employment effects by industry, which are small and insignificant across the board. For example, the first row shows that in a location whose local average tariff falls to zero, both men and women exhibit

[9] We are grateful to an anonymous referee for pointing out this possibility.

[10] The tariffs show different geographic patterns of vulnerability. Appendix Table A6 shows the ten conspumas with the highest initial local average tariffs for both genders. The lists have a lot of overlap, and conspumas in the Carolinas and Georgia dominate, but the ranking is different and Danville, VA is in the top ten for women but not men, while Indiana locations show up for men but not women.

Table 9: Gender-Specific Local Tariffs

	Wage		Employment		Labor force	
	Male	Female	Male	Female	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)
Location effect						
Less than high school	-3.55***	-7.05***	-0.95***	-5.75***	0.12	-4.49***
High school graduate	-1.463**	-3.714***	-1.08***	-4.02***	-0.927***	-3.661***
Some college	-1.539***	-4.772***	-0.456*	-3.428***	-0.602***	-3.039***
College graduate	-0.31	-4.491***	-0.296	-1.669***	-0.311*	-1.629***
Industry effect						
Less than high school	-1.21***	-3.047**	-0.542*	-0.42	-0.215	-0.941
High school graduate	-0.269	-2.152***	-0.31	-0.569	-0.294*	-0.927*
Some college	-0.144	-0.872	-0.159	-0.432	-0.296**	-0.841**
College graduate	-0.132	1.118	-0.18	-0.195	-0.221**	-0.483
N of Observations	5,366,329	4,862,010	6,165,584	5,789,584	6,165,584	5,789,584

Notes: The table follows the same structure as Table 3, but uses gender-specific local tariffs. The table reports the overall impact on male and female labor market outcomes (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

declining employment rates, but the effect for women is double that for men. Hypothesis (ii) would require a *larger* effect for men. This point is made once again in Table 9; comparing columns (3) with (4) and (5) with (6), we can see that the labor-supply elasticities of men and women are either both small or the women's elasticity is much larger than the men's. Therefore, differential labor-supply elasticity might have been able to explain a larger wage effect for men, but it cannot explain our observed larger wage effect for women. (This point is brought home in a different form in our later discussion of Table 10.)

3.3 Household bargaining

3.3.1 Theory

We now consider the possibility that household bargaining, with asymmetric bargaining power within the household, may be driving the results.

For illustration of the main points in the simplest way possible, consider a model with two periods, two industries, and two towns. Suppose that industry 2 is the numeraire and produces an export good, and industry 1 produces an import-competing good, whose world price is P^w, which is taken as given, while the domestic price is $P = P^w + t$, where t is an import tariff. All economic agents have the same homothetic utility function, which produces a consumer price index $\phi(P)$. Denote the *real* price of good 1 by $p_1 \equiv \frac{P}{\phi(P)}$, which is increasing in the tariff; and the real price of good 2 by $p_2 \equiv \frac{1}{\phi(P)}$, which is decreasing in the tariff. Each worker can produce either good $i = 1, 2$ in either town $j = 1, 2$; no other factor than labor is required.

Each worker z has an inherent ability $a^{z,i,j}$ in industry i in town j. The worker's ability in a given industry is allowed to differ from one town to the next, which could occur because the worker has social networks or previous business associates in particular locations that allow him/her to find a more productive business arrangement than in other locations, even within the same industry (there is strong evidence for the importance of local social networks in finding employment; see Topa (2001)). We could think of the $a^{z,i,j}$ as representing worker z's local "opportunities" in industry i in town j. Worker z's real wage is then $w^{z,i,j} = p_i a^{z,i,j}$ if he or she works in industry i in town j.[11]

In addition to the wage, each worker z expects a utility benefit $\epsilon^{z,j}$ from being in city j. This could be due to idiosyncratic tastes for climate, amenities, friends or enemies who

[11]This structure is of the type known as an 'assignment model' (Costinot and Vogel, 2015). It would be much more realistic to assume that each industry produces with labor and at least one other factor, for example, a specific factor which is in fixed and exogenous supply in each town. Specifically, each industry i in each town $j \in \{1, 2\}$ could have an endowment of a specific factor denoted $K^{i,j}$. This would allow for the two towns to have different employment patterns. Those features create complications that are not germane to the point being made here, however, so we omit them.

happen to live in each town. Both a worker's ability in each industry and town, and that worker's preference for each town, are fixed for that worker's lifetime, and the distribution of these two traits across workers is independent. Suppose that the utility the worker receives is a function $v(c^z, \epsilon^{z,j})$ of consumption c^z and amenity preferences $\epsilon^{z,j}$.

Now, suppose that during period 1, it is announced that the tariff t will be reduced, lowering real price of output in industry 1, and hence lowering the real wage for every worker employed in that industry. Workers in each industry have the option of switching to the other industry and/or town at the end of period 1. If a worker switches, he/she will receive the period-2 wage and idiosyncratic town utility benefit in the new industry/town combination.[12]

Assume that the workers are composed in equal numbers of male and female, and that some fraction are paired up in heterosexual marriages. The distribution of abilities and town preferences is the same for each gender and also for married and single workers.

First, consider single workers. A worker with no family attachments will simply choose the industry and city combination (i, j) in each period to maximize $v(w^{z,i,j}, \epsilon^{z,j})$, since for such a worker consumption c^z will be equal to the real wage. As a result of the movements of workers out of industry 1, the drop in wages to industry-1 workers caused by the tariff reduction will be mitigated by a selection effect: The workers who leave the industry are on average those who are less productive in industry 1 than the average worker in the industry. This selection effect means that the average wage for single workers in industry 1 will fall by less than the output price p_1.[13] Importantly, some of the workers who switch industries will also switch towns in order to do so, because they have an attractive industry-2 opportunity in the other town but not in their original one ($a^{z,2,j} > a^{z,2,j'}$).

[12]The idiosyncratic abilities and town benefits will imply that only a fraction of workers will switch industries or move following the trade shock. In this way, they act like switching costs or moving costs. A full model would need to include direct costs of moving and switching industries, such as retraining and the like. We omit those here for simplicity of exposition.

[13]If we had a richer model with a fixed factor in each industry, there would be a second mitigating effect: The reduction in the labor supply to industry 1 would push up the marginal physical product of labor in that industry, increasing the price of effective labor there, and so increasing the wage received by any industry-1 worker conditional on ability.

Now, consider married workers. Assume that both partners in a marriage must live and work in the same town; that all workers are employed in equilibrium, regardless of gender or marital status; and that marriages do not either form or break up. (We discuss the endogeneity of marriage in Section 4.) Within each marriage, intra-household allocation issues are dealt with by bargaining, as for example in Browning et al. (1994). Suppose that at the beginning of Period 1, each couple finds itself exogenously located in one of the two towns,[14] and must bargain to choose the town in which to live and work in Period 1, and again bargain at the beginning of Period 2 after the policy has been revealed.[15] The threat point takes the form of continuing to live in the initial town and each partner in the marriage consuming his/her real wage.

Analyzing this model in detail would be beyond the scope of this paper, but clearly for either married worker the relocation decision is more complicated than it is for a single worker. There will be cases in which a married worker in industry 1 will not switch industries, because to do so for that worker would require switching cities, which would be costly for the spouse, while the same worker would have switched as a single worker. As a result, married workers may be less able to move out of the shrinking industry, and the selection effect that attenuates the drop in average wage for single workers will be weaker.

For these reasons, a model of this sort can easily rationalize a larger wage effect for tariff reduction for married workers than for single workers, but in order to rationalize a greater effect for married *women* than for married men, some special assumptions need to be made. In the case of linear preferences, in which $v(c, \epsilon) = c + \epsilon$, it can be shown that the mobility behavior of married men and married women will be the same, now matter how much bargaining power the husband has (see our working paper for details). The reason is that the household chooses its location to maximize the joint surplus, which takes both spouse's incomes into account equally, and then bargains over intro-household division of

[14] In a fuller model of dynamic adjustment, such as in Artuç and McLaren (2015), this initial allocation would be determined endogenously as the per-shock steady state.

[15] We assume that the change in tariff at the start of Period 2 is a surprise, so does not factor into Period-1 bargaining.

the surplus. Relative bargaining power matters only for the latter of those two choices. This is clearly negated by the data; our motivating finding from the main regression shows much larger wage effects for married women. However, alternative specifications are possible; for example, if $v(c, \epsilon) = c\epsilon$, a member of the household will enjoy consumption spending more while located in a town that he or she enjoys. This allows for bargaining power to affect the couple's location decisions, and allows for the possibility that the household will move in response to a husband's industry shock more readily than a wife's, if the husband has greater bargaining power.

3.3.2 Empirical test

To test this theory, we run a set of regressions where the dependent variable is the share of employed married or single women or men of educational class k in each industry j and conspuma c in total labor force or working age population of conspuma c between 1990 and 2000. Our aggressors include industry- and location-specific initial tariffs and change in tariffs. The results are reported in Table 10, in which Panel A records the results for women and Panel B for men, and the first two rows of each panel show the results for the share in the working-age population and the following two for the share in the labor force.

According to the household-bargaining theory, the employed married women's share in each industry/conspuma should *rise* relative to other groups when hit by a trade shock since other groups are leaving the industry/conspuma but at least a fraction of the married women cannot leave. This is not at all what we find. For example, focusing on high-school dropouts, negative coefficients in the first column show that the share of high-school-dropout married women workers in a town or industry that loses its tariff completely falls, indicating that married women in that educational group leave that town, that industry, or the labor force; and the coefficient is about twice the size of the corresponding value in column 5, for single women, indicating that married women workers adjust *more* than single women. The estimates are all strongly statistically significant. The results for high-school graduates in

Table 10: Change in share of employed married/single women/men in working age population and labor force

Panel A	Married women				Single women			
	Less than HS	HS	Some College	College Grad.	Less than HS	HS	Some College	College Grad.
Working age population								
Location effect	-0.0022***	-0.0008*	0.0002	-0.0001	-0.0011***	-0.0001	0.0003	-0.0002
Industry effect	-0.0013***	-0.0009***	-0.0015***	-0.0016***	-0.0007***	-0.0010***	-0.0012***	-0.0008***
Labor force								
Location effect	-0.0023***	-0.0004	0.0005*	0.0001	-0.0011***	0.0002	0.0006**	-0.0002
Industry effect	-0.0014***	-0.0011***	-0.0018***	-0.0019***	-0.0008***	-0.0012***	-0.0014***	-0.0010***

Panel B	Married men				Single men			
	Less than HS	HS	Some College	College Grad.	Less than HS	HS	Some College	College Grad.
Working age population								
Location effect	-0.0025***	0.0014***	0.0002	0.0007**	-0.0002	0.0002	0.0005***	0.00005
Industry effect	-0.0013***	-0.0022***	-0.0015***	-0.0011***	-0.0004***	-0.0009***	-0.0007***	-0.0004***
Labor force								
Location effect	-0.0025***	0.0022***	0.0006	0.0011***	-0.0001	0.0005	0.0007***	0.0001
Industry effect	-0.0015***	-0.0026***	-0.0018***	-0.0014***	-0.0004***	-0.0011***	-0.0008***	-0.0006***

Notes: N = 116,750. The number of observations is not equal to the number of conspumas (543) times the number of industries (239) because we only include industries for which we observe at least one of these groups (single/married women/men) to have positive employment. The table reports the overall impact on the change in the share of female/male workers (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

the next column are similar but not as strong. Comparing the first two columns of Panel A with those of Panel B, we find that the effects for married men and women are quite similar, with the male responses the same as for married women or slightly larger. But the married women would have to have a clearly smaller response than all other groups in order for the household bargaining explanation to be the right one. We conclude that household bargaining does not explain the married-woman differential. This explanation would require married female workers to be less elastic in their response than other demographic groups, but the evidence rejects this prediction.

3.4 Selective non-participation

3.4.1 Theory

Some studies, such as Autor et al. (2013) and Saurian and Zoabi (2014), have found evidence of workers withdrawing from the labor force in response to a loss of tariff protection.[16] We argue here that under some conditions selection decisions by some women to withdraw from the labor market in response to a trade shock hitting their industry could produce magnified wage responses for married women compared to other workers. The way this could happen is as follows.

Suppose that single workers have no option to stay out of the labor market, and suppose that cultural norms prevent a married man from doing so except in case of disability or retirement age (this will of course depend on the time and place and local culture, but is probably a reasonable assumption to impose for our data period). Under these assumptions, the only group of workers with an option to leave the labor market is married women. Suppose that a married woman will choose to remain in the labor market if her wage is high enough or her husband's wage is low enough; then if an import-competing industry is hit with a trade shock that lowers wages for all workers in the industry, a certain fraction

[16]By contrast, Dai et al. (2018) find in Chinese Census data that local labor markets hit with tariff reductions see an *increase* in female labor-force participation.

of the married women will respond by leaving the labor market. Now, if those women are the most productive women in that industry, their departure will lead to a selection effect that will magnify the effect on average wages of married women still in the industry. This is exactly what will happen if two conditions are satisfied: (i) The departing women have higher-income spouses – as they will tend to do because only a worker with a sufficiently highly-paid spouse can afford to leave the labor market. (ii) Partners in marriage with highly-paid spouses tend to be highly-paid themselves, since the marriage market features positive assortative matching. These two features together tend to lead to the departing women being higher-wage workers than the ones they leave behind, pushing average wages down beyond the effect of the initial trade shock.[17]

We can formalize a simple model as follows. Suppose that unmarried workers simply consume their own wages, but married workers share their earnings. Suppose that all married couples have the same utility function, an increasing, concave, twice-differentiable function $U(\cdot)$, which is a function of the couple's combined real wage. If a married couple have a wife whose real wage is w^w and the husband's real wage is w^h, and if they both work, their utility is $U(w^w + w^h)$. On the other hand, if the wife chooses not to work, their utility is $U(w^h)+F$, where $F > 0$ is extra utility they share from the wife's extra time for non-market activities, a parameter that is the same for all households. If $U(w^w + w^h) - U(w^h) \geq F$, the wife will work, and otherwise she will leave the labor market. (For all workers, for the moment assume that there is no other alternative employment; there is only one choice, and that is to be in or out of the labor force for married women.)

Clearly, for a given w^h, a married female worker will remain in the labor market if and only if w^w is above a given threshold. Denoting that threshold as $\tilde{w}^w(w^h)$, and taking the

[17] Note that the effects we find empirically are all within an educational class, since we control for education. Consequently, the theory should be interpreted as one in which higher-wage women *within their educational class* withdraw from the labor market in response to a trade shock.

derivative of $U(\tilde{w}^w(w^h) + w^h) - U(w^h) = F$ with respect to w^h, we obtain:

$$\frac{d\tilde{w}^w}{dw^h} = \frac{U'(w^h) - U'(\tilde{w}^w(w^h) + w^h)}{U'(\tilde{w}^w(w^h) + w^h)} > 0. \tag{8}$$

Therefore, we can draw a figure with w^w on the horizontal axis and w^h on the vertical axis, with an upward-sloping curve representing the threshold between the region in which the female worker stays in and leaves the labor force. This curve is represented in Figure 1, which measures the wife's real wage on the horizontal axis and the husband's on the vertical axis. Any point in the figure to the right of the curve represents a couple for whom the wife's wage is high enough and the husband's wage is low enough that the wife remains in the labor market. Any point to the left of the curve represents a couple for whom the wife will leave the labor market. An assumption on the curvature of U allows us to characterize the shape of the curve:

Proposition 1. *If the coefficient of relative risk aversion associated with U is everywhere greater than 1, then the curve defined by $U(\tilde{w}^w(w^h) + w^h) - U(w^h) = F$ goes through the origin. Further, any ray through the origin that intersects the curve will intersect it from below, and only once.*

The implication of Proposition 1 is that, if relative risk aversion exceeds 1, other things equal, a woman will be more likely to work, the higher is her wage (the farther to the right is the couple in the diagram), and the *lower* is her husband's wage (the farther down is the couple in the diagram).

To fill in the rest of the model, suppose that there are many industries, one of which is the import-competing industry 1, initially protected by a tariff. The price of industry-1 output is denoted p. We wish to compare outcomes before and after a trade shock. To make the analysis as simple as possible, consider a two-period model, and suppose that in Period 1 workers select their industries, and a fraction λ of male and female workers choose a spouse and marry, expecting the same trade policy to prevail in Period 2. Importantly,

Figure 1:

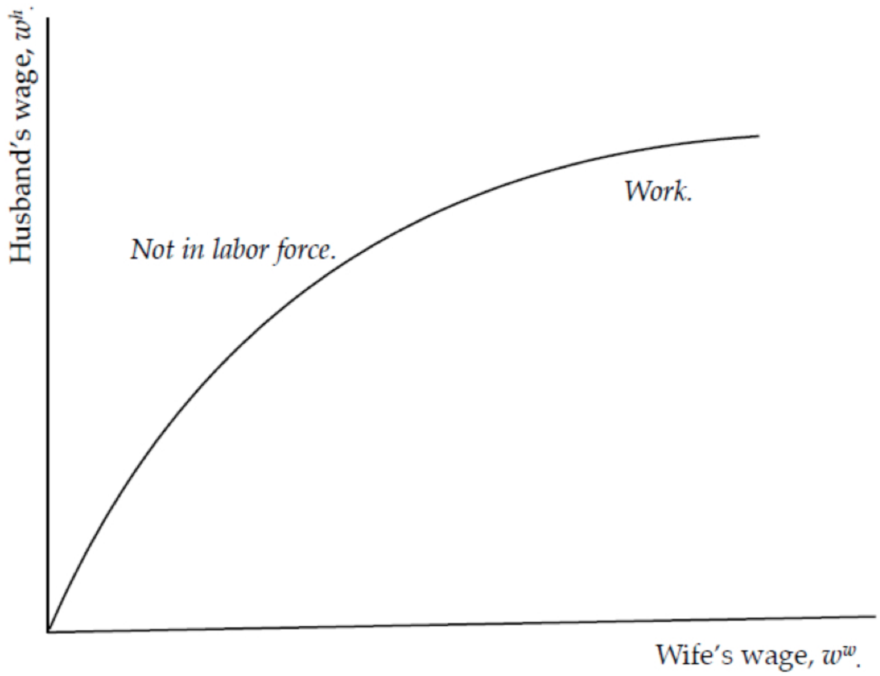

for this discussion, λ is an exogenous constant. We discuss the important question of how the marriage rate may respond to trade policy, and how that may affect the married-woman differential, in the next section. In Period 2, all agents are surprised by a change in trade policy that lowers the value of p. Workers are unable to change their choice of industry or spouse in Period 2. Denote the initial-equilibrium value for the industry-i output price by p^i and the Period-2 value by $\tilde{p}^i$.

Each worker z has ability level $a^{z,i}$ in industry i, which is a constant for each worker. The distribution of $a^{z,i}$ values is the same for male and female, married and unmarried workers. The wage received by worker z in industry i is $w^z = a^{z,i}p^i$, so each worker z will have a wage given by $w^z = \max_i\{a^{z,i}p^i\}$.

Now, suppose that a randomly selected fraction λ of male and female workers marry in Period 1, sorting according to positive assortative matching. Given the symmetry of the model, this implies that within each marriage the male wage and the female wage are equal. As a result, every married couple will occupy a point along the 45^o line, portrayed as the solid ray, in Figure 2. Some fraction will have the husband and the wife both in industry

Figure 2:

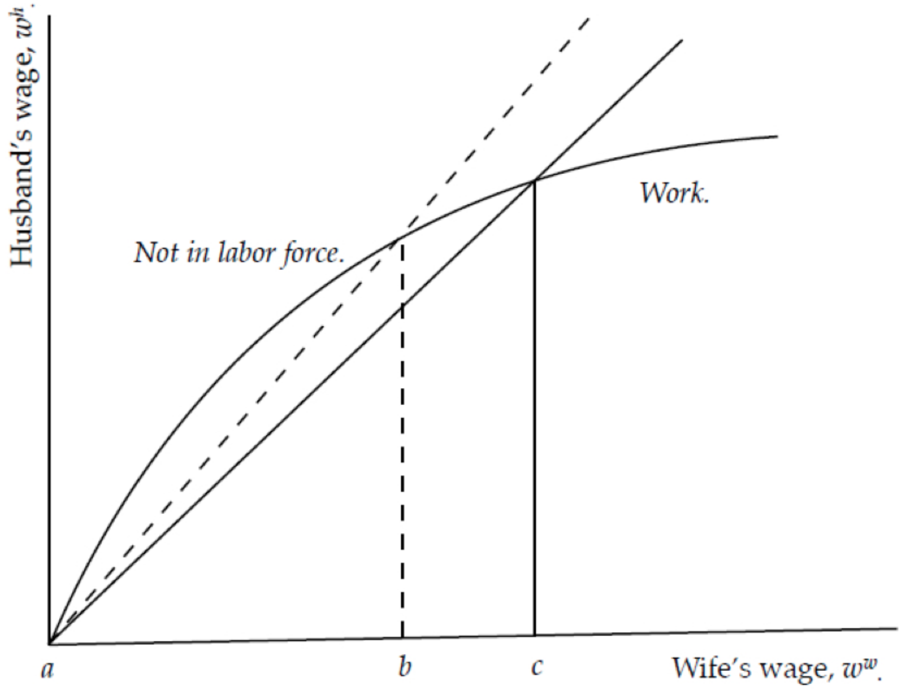

1; some will have husband in 1 and wife in 2, and so on; and some fraction will be located above the curve so that the wife leaves the labor market. The range of wages for this subset of married female workers in industry 1 is given by ac.

Now, consider a married couple with the wife in industry 1 and the husband in some other industry. When the Period 2 shock arrives, since p^1 will fall to $\tilde{p}^1$, the ray showing the wage pairs for this subset of married couples will rotate as shown in the broken ray in Figure 2. Consequently, a fraction of the women in this set will leave the labor market, and only ab will remain. Since the portion of workers who remain in industry 1 will see a wage reduction of $\frac{p^1-\tilde{p}^1}{p^1}$, and the portion who leave the labor market, bc, are at the higher end of the wage distribution, the average wage for married female workers in this industry will fall by *more* than $\frac{p^1-\tilde{p}^1}{p^1}$.

On the other hand, a couple with both members in industry 1 will see both wages fall proportionally, a move down and to the left along the solid ray in Figure 2. If both spouses were initially in the labor market, they will continue to be so after the shock. Consequently, average wages for married women in this category will fall by $\frac{p^1-\tilde{p}^1}{p^1}$, the same as unmarried workers or married male workers.

This can all be summarized as follows.

Proposition 2. *Assume that the coefficient of relative risk aversion associated with U is everywhere greater than 1. Then as a result of the trade shock, the wages of all workers in industry 1 fall in the same proportion, except for married women whose husband is in a different industry. Their average wage in industry 1 falls by more than the other groups, and their share of employment in industry 1 falls.*

3.4.2 Empirical test

To see if this theory is consistent with the data, we first look at some basic correlations implied by the model. Specifically, we ask (i) whether our data exhibit positive assortative matching, and (ii) whether, as implied by Proposition 1 (combined with relative risk aversion greater than unity), a higher husband's wage lowers the probability that a married woman will stay in the labor force. To test these hypotheses we need a model that deals with the sample-selection issue raised by the fact that the wage of a worker who has left the labor force is not observed.[18]

Suppose that each female worker i has a latent wage w_i^*. This is a function of the worker's personal, industry, geographic, and occupational characteristics X_i including potentially the husband's wage because of positive assortative matching:

$$w_i^* = X_i\beta + \epsilon_i. \tag{9}$$

The wage is observed if and only if the worker chooses to be in the labor force. This occurs if:

[18]We have last year's wages for workers who are not currently in the labor force but who were employed last year. We do not use those wages here.

$$w_i^* \geq Z_i\beta + \eta_i. \tag{10}$$

The Z_i should in principle could contain all of the variables in X_i plus some variables that plausibly affect the decision to be in the workforce but not the wage conditional on being in the labor force, such as cultural factors, family information and home ownership. Assume that η_i is distributed as $N(0,1)$ and ϵ_i is distributed as $N(0,\sigma)$, with a correlation of ρ between the two. We restrict the sample to married female workers whose husbands work at least 35 hours per week and report positive wages, and estimate these two equations together using the Heckman two-step procedure.

The results are reported in Table 11. The first column shows the results for equation (10). This equation contains standard demographic and educational variables, plus a dummy for home ownership (as a proxy for wealth), for presence of a school-aged child (which can raise the opportunity cost of going into the workforce), and for immigrant status, which may be correlated with cultural views on female labor-force participation. These last three are omitted from the second column, which shows the results for equation (9). The second-to-last row of the table shows the estimate of $\lambda \equiv \sigma\rho$. One important note to make is that, as in all of our wage regressions, the specification of equation (9) includes conspuma dummies, which are (as always) suppressed in the table, but equation (10) does not have conspuma dummies since a probit is inconsistent as the number of fixed effects becomes large, and the estimation does not converge in our case.

Controlling for all other observables, a female worker in our data is less likely to be in the labor force if she has less than a college education, if she has a school-age child at home, and if she is an immigrant. A small surprise is the positive coefficient on home ownership; one might have expected this to have a negative sign, if it proxies for higher household wealth. Other things equal, the probability of being in the labor force peaks at age 45, while the wage takes a minimum at that age.

The main coefficients for our purposes are in the first row. *Ceteris paribus*, a higher

Table 11: Selection model

Dependent variable:	"In Labor Force" Dummy	Wife's Logwage
Spouse's logwage	-0.164***	0.114***
	(0.00140)	(0.00321)
Age	0.150***	-0.0766***
	(0.000920)	(0.00325)
Age squared	-0.00175***	0.000957***
	(1.08e-05)	(3.77e-05)
White	-0.158***	0.0600***
	(0.00339)	(0.00499)
Speak English	-0.0121***	0.00450*
	(0.00196)	(0.00271)
Less than high school	-0.520***	-0.0421***
	(0.00418)	(0.0129)
High school graduate	-0.222***	-0.150***
	(0.00288)	(0.00603)
Some college	-0.131***	-0.156***
	(0.00287)	(0.00477)
Own house	0.133***	
	(0.00281)	
School-aged child	-0.0322***	
	(0.00233)	
Immigrant	-0.133***	
	(0.00418)	
Year 2000 dummy	0.0137***	0.424***
	(0.00216)	(0.00303)
Lambda	-2.234***	
	(0.0640)	
N of Observations	2,919,674	2,919,674

husband's wage implies a higher wage for the female worker, which is evidence of positive assortative matching (item (i) above). In addition, *ceteris paribus*, a higher husband's wage is correlated with a lower probability of being in the labor force, which is in line with Proposition 1 (item (ii) above). Both of these findings hold with a high degree of statistical significance.

Having checked for the basic correlations implied by the model, we turn now to testing the main predictions of the theory. First, we test the prediction of Proposition 2, that some fraction of married women choose not to participate in the labor force in response

Table 12: Labor Force Participation

	Married workers		Single workers	
	Male	Female	Male	Female
	(1)	(2)	(3)	(4)
Location effect				
Less than high school	-0.36	-1.673***	-0.311	-2.003***
High school graduate	-0.777***	-1.462***	-0.883***	-1.449***
Some college	-0.636***	-1.125***	-0.533***	-1.039***
College graduate	-0.408***	-0.613***	-0.495**	-0.684***
Industry effect				
Less than high school	-0.12	-1.148*	-0.219	-0.826
High school graduate	-0.352**	-1.055**	-0.159	-0.733
Some college	-0.396***	-1.162**	-0.005	-0.284
College graduate	-0.184	-0.342	-0.327***	-0.541**
N of Observations	4,349,897	3,792,160	1,815,687	1,997,424

Notes: The table reports the overall impact on male and female wages (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

to a trade shock, in two ways. The first is to estimate a linear probability model of labor force participation where the dependent variable is a dummy that takes a value of 1 if the individual is in the labor force and zero otherwise. The second test examines a subsample of married men and women where both spouses are employed in the same or different industries. According to our theory, for those couples that work in the same industry there would be no effect on labor force participation, and the effect of the trade shock on wife's and husband's wages would be the same, because if the husband is hit with the same wage shock as the wife, the couple cannot afford to lose her income.

Table 12 reports the results from a linear probability model of labor force participation for each of four groups of workers. The right-hand-side variables are the same as in the wage regression, and the results are arranged in the same way as in Table 6. The coefficients for location effects are negative for all groups of workers at almost all educational levels, with the effects being larger for both married and single women and decreasing in the level of educational attainment. This implies that, for example, high-school dropout female workers in a location that had high protection and lost it by 2000 are less likely to be in the labor force by 2000. The industry effects are less strong and imprecisely measured, but the overall story is the same in that women are more likely to drop out of labor force than men in highly-protected industries that lost their protection by 2000. The overall message of Table 12 is that NAFTA-driven tariff reductions did disproportionately push female workers, both married and single, out of the labor force in the hardest-hit communities.

Another implication of the marriage market/selection model is that among married couples, controlling for all other factors, wages should fall for female workers when their own industry tariff falls but should *rise* when their husband's industry tariff falls, since if the husband's tariff falls, the female worker is more likely to remain in the labor force, and under conditions of Proposition 1 this will raise the average productivity of working women in her industry. Table 13 examines this prediction. This is a wage regression for the married women in the sample with employed husbands, with all of the controls of earlier regressions and the same format, except that we control not only for the female worker's industry tariff but also for her husband's industry tariff. (The location tariffs are the same for both spouses.) Recalling our basic estimating equation (2), this creates a large number of interactions, which prove to be too many to estimate, so we collapse the four educational categories to two, 'Less educated' (high-school or below) and 'Highly educated' (some college or college graduate).

The first four rows of Table 13 show results in line with the main specification as in Table 6; a married female worker's wage falls when her town's average tariff or her industry's tariff falls. The last two lines show the effect of the husband's industry tariff, and these are in line with the prediction of the selection model: A drop in the husband's industry tariff increases the wife's expected wage, with high statistical significance. It is difficult to imagine

Table 13: Own and Spouse's Industry Tariffs:

Dependent variable:	Wife's logwage
Location effect	
Less educated	-1.132**
Highly educated	-1.55**
Industry effect	
Less educated	-4.708***
Highly educated	-2.054**
Less educated (husband)	2.111***
Highly educated (husband)	2.79***
N of Observations	3,052,775

a mechanism by which this finding could be rationalized except through the sort of selection mechanism we have sketched here.

Finally, recall that according to the labor-force-participation model, if both spouses are in the same industry, a drop in the tariff should have no effect on the wife's participation decision, while if they are in different industries, a drop in the wife's industry tariff will make her more likely to leave the labor force. To test this hypothesis, we next limit our sample to married individuals with both spouses being employed, and run separate wage and labor force participation regressions for spouses employed in the same and different industries. Although the sample size for the former type of couples (employed in the same industry) is small and the estimates are imprecisely measured, we do find support for our theory (Table 14). In particular, there is no distinguishable difference in labor force participation response of men and women when both spouses are employed in the same industry (columns 5 and 6); for this group, the labor-force participation effects are mostly small and insignificant. The earlier differential results across men and women are completely driven by women that are employed in a different industry then their spouse.

The main elements of the selective non-participation story are therefore consistent with the data. However, we should note that this theory does not provide a full explanation. Table 14 shows that even for couples with both spouses employed in the same industry, the effect of industry tariff reductions on the wife's wage (-4.326) is negative, statistically

Table 14: Selective Non-Participation

	Wage growth				Labor Force Participation			
	Same industry		Different industries		Same industry		Different industries	
	Male	Female	Male	Female	Male	Female	Male	Female
Location effect								
Less than high school	2.827**	-1.876	-0.8	-3.91***	0.408	0.483	-0.508***	-2.171***
High school graduate	-0.601	1.146	-0.268	-0.454	0.029	-0.208	-0.321***	-1.194***
Some college	1.802	-1.022	-1.586***	-1.518**	-0.17	-0.956***	-0.44***	-0.975***
College graduate	1.177	-1.192	-0.533	-2.193*	-0.252**	0.14	-0.172*	-0.539***
Industry effect								
Less than high school	0.207	-4.326*	-1.028*	-3.503**	0.08	-0.526	-0.325**	-1.304**
High school graduate	-1.357**	-3.28**	-0.315	-2.689***	-0.631***	-0.222	-0.242*	-0.994**
Some college	-0.268	-1.263	-0.001	-1.502	-0.13	-0.169	-0.205**	-1.177**
College graduate	1.132	1.12	-0.253	2.17	0.415*	-0.092	-0.165	-0.251
N of Observations	279,714	294,311	2,204,347	2,348,297	326,036	360,951	2,474,287	2,817,623

significant, and much larger than the statistically insignificant effect on the husband's wage (0.207). This would not be the case if our selective non-participation story was the sole force driving the results.

4 Endogenous marriage

To this point, we have assumed that the decision to marry or to remain married does not itself respond to changes in trade policy. This is obviously not a safe assumption, and if the marriage decision is affected by trade policy the married-woman differential could be a mix of a change in the pool of married female workers (which we have ignored so far) together with changes in individual workers' wages and changes in the labor force participation (which were the subject of the previous section). If we had panel data, we could separate these effects cleanly, measuring the effect of a tariff change on a worker conditional on initial marital status, holding fixed unobserved time-invariant personal characteristics, but with our repeated cross sections we are unable to do so. Therefore, although the possible effect of

NAFTA on marriage is outside of our focus in this paper, the topic needs to be addressed.

Indeed, Autor et al. (2017) present evidence that increases in manufactured imports from China have had a strong influence on marriage rates and the marriage market. Autor et al. (2017) show that locations in the US that saw a larger China import shock (because they had larger employment shares in industries that saw large increases in imports from China) saw substantial reductions in the income of male workers relative to female workers. These changes were correlated with a drop in the fraction of workers who were married. This is interpreted in terms of a model in which women are reluctant to marry a man who is likely to have low income relative to the woman's income, an assumption that has broad empirical support in the literature on the economics of the family, and which we can call the 'relative-income hypothesis.' Because of this hypothesis, a shock to the labor market that lowers men's incomes relative to women's will tend to reduce the number of marriages in equilibrium.

To explore whether or not this sort of mechanism could possibly help explain the married-woman differential finding that is our focus, consider the following stylized model of the marriage market, based on an extremely simplified version of the relative-income hypothesis.[19] Suppose that a given local labor market has a unit endowment of male workers and a unit endowment of female workers. Two industries produce with Ricardian technology. For cultural reasons, men all work in industry 1 and women in industry 2. This stark assumption allows for the lowering of tariffs to have a differential effect on the two genders' incomes, which Autor et al. (2017) show to be important in the case of the China shock. Each worker z has a productivity a^z in his or her industry. These productivities have a uniform distribution on $[a_{min}, a_{max}]$. Industry 2 is the numeraire sector, and industry 1 output has a domestic output price given by p. Therefore, a man's income will be given by pa^z and a woman's income by a^z. We close down the mechanism discussed in the previous sub-section

[19] Autor et al. (2017) present a model in which childbearing has a central role, but we omit such considerations because they are not essential to make the point about wage movements that is our interest here.

by supposing that non-participation in the labor market is not an option for any worker.

Each worker wishes to marry a worker of the opposite sex, but a woman will refuse a match with a man whose income is not at least equal to $\kappa > 1$ times her own; if no potential spouse is available who satisfies that constraint, she will prefer to remain unmarried. This is a crude way of adding the relative-income hypothesis into the analysis, but it suffices for our purposes. Aside from this constraint, each worker will prefer a spouse with a higher income to one with a lower income. If $p \geq \kappa$, every worker will marry and positive assortative matching will apply, so in order to allow for a non-trivial outcome, assume that $p < \kappa$. In this case, the highest-income male worker (with income equal to pa_{max}) will be unacceptable to the highest-income female worker (with income equal to a_{max}), so she will choose to remain unmarried. The highest-income female worker to marry will have $a^z = pa_{max}/\kappa$, and will marry the highest-income male worker. If male and female workers match assortatively this implies that the lowest-income female worker will marry a male worker with productivity $a_{min} + a_{max}\left(1 - \frac{p}{\kappa}\right)$. If the inequality

$$\frac{a_{max}}{a_{min}} > \frac{\kappa}{p} \tag{11}$$

holds, then it is easy to confirm that this marginal male worker's income exceeds κ times the income of the lowest-income female worker, so this is an equilibrium.[20]

In equilibrium, then, (i) the most productive female workers in their industry remain unmarried; (ii) the least productive men in their industry remain unmarried; (iii) the fraction of both groups who remain unmarried is equal to:

$$\left(\frac{a_{max}}{a_{max} - a_{min}}\right)\left(1 - \frac{p}{\kappa}\right), \tag{12}$$

which is a decreasing function of p. As a result, if industry 1 is import-competing so that

[20] If the stated inequality does not hold, the equilibrium is qualitatively the same but the lowest-income female worker is indifferent between remaining unmarried and marrying the marginal male worker.

lowering tariffs result in a drop in p, male worker incomes will fall relative to female worker incomes, resulting in a decline in marriage rates in this local labor market (all as described in Autor et al. (2017)). More subtly, the pool of married female workers will become less productive, because the marginal women choosing to switch to unmarried status would have been the most productive married female workers; but the pool of *unmarried* female workers will *also* become less productive, because those marginal women will now be the *least* productive unmarried female workers. The opposite effects will hold for married and unmarried male workers. To summarize:

Proposition 3. *In this model of endogenous marriage rates, a reduction in the import-competing price p causes: (i) A shift to the left in the distribution of incomes of men relative to women. (ii) A reduction in marriage rates, with a drop in the productivity cutoff above which women remain unmarried and an increase in the productivity cutoff below which men remain unmarried. (iii) A reduction in the average productivity and average wage of married female workers and also of unmarried female workers. (iv) An increase in productivity of married men and also of unmarried men, and so a reduction in the average wage for each of those two groups that is smaller than the drop in p.*

It is conceivable that this mechanism in and of itself could give rise to a married-woman differential, but it is by no means clear that the drop in married-woman wages will exceed the drop in unmarried-woman wages. However, for this to be the driver of the married-woman differential, a number of other predictions must hold, including (i) falling male incomes relative to female incomes in response to falling local average tariffs; (ii) falling marriage rates in response to falling local average tariffs; (iii) relative effects for unmarried wages, and (iv) relative effects for male wages. We address these in turn.

4.1 Male incomes relative to female incomes

A straightforward way of measuring a shift in the distribution of male incomes relative to female incomes is a quantile regression. Table 15 reports the results of conspuma-level regressions to study this question.[21] In the first three columns the dependent variable is the first difference in the x-percentile male wage in a conspuma minus the x-percentile female wage in the same conspuma, for $x = 25\%$, 50%, or 75% (the first difference is measured between 2000 and 1990). In these calculations the wages of all working-age persons are included, including people with zero wage and people who are not in the labor force. The regressors are the initial local average tariff in 1990 and the change in the local average tariff between the two years. The 'Overall effect' is the difference between the two coefficients and measures the expected change in the dependent variable per unit of tariff for a conspuma that lost all of its tariff protection by 2000. These effects are all positive and significant. For example, a one-standard-deviation increase in the initial local average tariff, if the tariff went to zero by 2000, would increase the 25^{th} percentile male wage relative to the 25^{th}-percentile female wage by $0.67\% \times \$40,188$, or about \$270. For the last three columns the dependent variable is the difference in male and female wage quantiles as a percentage of the male wage. The effect is again positive, but significant only for the higher wage quantiles. For the 25^{th} percentile, the point estimate implies that a one-standard-deviation higher local average tariff eliminated by 2000 would be associated with a $0.67\% \times 264 = 1.8$ percentage-point increase in the male-female wage differential. The effects for higher quartiles are much smaller.

The conclusion of this table is that local-average tariff reductions due to NAFTA were not correlated with large changes in the local wages of men relative to women, and to the extent that there was any change it was a small *increase* in the relative wages of men. This appears to be a major contrast with the later China shock, and conceivably the reason could be the NAFTA-induced relative decline in married women's labor market participation discussed above. Therefore, the first element of the endogenous-marriage-rate explanation for the married-woman differential is not present.

[21] This table is analogous to Table 3 of Autor et al. (2017), which is a much simpler version of a quantile regression studied by Chetverikov et al. (2016).

Table 15: Change in male-female annual wage differential by percentile of conspuma's wage distributions

	Male-female wage differential ($)			Male-female wage differential in % of male wages		
	P25	Median	P75	P25	Median	P75
Local tariff in 1990	-101,390	-69,360	-178,231	-2,887	156.4	160.0
	(146,529)	(110,579)	(145,518)	(1,904)	(288.6)	(165.5)
Change in local tariff	-141,578	-127,061	-274,631	-3,151	2.653	96.53
	(153,114)	(119,537)	(165,571)	(2,121)	(318.5)	(180.8)
Overall effect	40,188*	57,701***	96,400***	264.0	153.7***	63.5*
Mean Outcome Variable	1,446	1,108	2,461	-9.34	-9.61	-7.08
Level of Male Earnings 1990	13,259	26,912	41,351			
Level of Female Earnings 1990	4,560	14,792	24,982			
N of Observations	543	543	543	540	543	543
R-squared	0.136	0.138	0.301	0.003	0.178	0.100

Notes: Standard errors are clustered by state. Regressions are weighted by working age population of conspuma in 1990. Additional controls include 1990 shares of conspuma population that are white, college educated, and fraction of employed women, as well as border dummy. The number of observations is different in column 4 because 25th percentile male wage in three conspumas was zero in 1990.

4.2 Marriage rates

The next prediction to check is that marriage rates fell in local labor markets that lost protection relative to others. This is evaluated in Table 16, which is constructed as in the previous table except with a dependent variable that is a measure of the marriage rate in the conspuma. The marriage rate in a conspuma can be computed in different ways; the numerator in these regressions is either the number of married persons in the working age population, among the employed, or among the labor force; and the denominator is either the total number of persons employed or the size of the labor force. Computed in any of these ways, the result is a small positive net effect, in most cases not significant. For example, from the specification with the rate computed as the number married in the working age population divided by the total labor force (second column), the effect of a one-standard-deviation increase in the local average tariff that is eliminated by 2000 is a rise in the marriage rate equal to $0.67\% \times 0.911 = 0.61\%$, or less than a one-percentage-point increase in the marriage rate.

Table 16: Change in share of married in total employment or labor force (includes population controls)

	Married in working age population		Married among employed		Married in labor force	
Change in share of married in	total employment	labor force	total employment	labor force	total employment	labor force
Local tariff in 1990	-2.205	-1.365	-0.621	0.0566	-1.130	-0.397
	(2.097)	(1.413)	(0.827)	(0.503)	(1.240)	(0.661)
Change in local tariff	-3.252	-2.276	-0.870	-0.104	-1.456	-0.629
	(2.532)	(1.732)	(1.007)	(0.620)	(1.492)	(0.812)
Overall impact	1.047**	0.911*	0.249	0.1606	0.326	0.232
N of Observations	543	543	543	543	543	543
R-squared	0.438	0.417	0.343	0.266	0.370	0.346

Notes: Standard errors are clustered by state. Regressions are weighted by working age population of conspuma in 1990.

The conclusion of this table is that local-average tariff reductions due to NAFTA were not correlated with large changes in local marriage rates, and to the extent that there was any change it was a small *increase* in marriage rates. Again, this appears to be a contrast with the experience of the China shock, and therefore, the second element of the endogenous-marriage-rate explanation for the married-woman differential is not present.

4.3 Relative wages of other demographic groups.

Recall from Proposition 3 that under this endogenous marriage-rate model, a larger drop in local average tariff would be associated with a drop in the average wage for unmarried workers relative to their average wage in other conspumas. However, we already know that this is not found in the data. The first four lines of the last column of Table 6 show us that there is no statistically-significant difference in the growth of average wages for unmarried female workers in conspumas that lost a lot of protection compared to others.

We conclude that although marriage rates should assuredly be viewed as endogenous, and in principle they could affect the wage response of married women in ways similar to the results we observe, it is unlikely that endogenous marriage rates are the explanation for the married-woman differential that we have observed in the data.

5 Conclusion

We have documented a sharp difference in labor-market response to NAFTA across gender and marital status: The largest effects of NAFTA, by far, are shown in the wages of married female workers whose industry of employment lost its tariff. We have shown that this cannot be explained by the different occupation mix of male and female workers, or by household bargaining in which husbands with disproportionate bargaining power within the household prevent their wives from adjusting to shocks as they otherwise would wish to do. In addition, the patterns in the data cannot be explained by marriage rates that respond to changes in tariffs. We do find some support for an interpretation based on selective non-participation, in which some married female workers adjust to a trade shock by leaving the labor market; under this interpretation, because of positive assortative matching in the marriage market, the ones who do so tend to be the women with higher wages. However, this does not account for all of the features of the data, so we are left with a puzzle.

References

[1] Aguayo-Tellez, Ernesto (2012). "The Impact of Trade Liberalization Policies and FDI on Gender Inequalities." World development report, Washington, DC: World Bank.

[2] Aguayo-Tellez, Ernesto, Jim Airola, and Chinhui Juhn (2010). "Did Trade Liberalization Help Women? The Case of Mexico in the 1990s," NBER Working Paper No. 16195.

[3] Artuç, Erhan, Shubham Chaudhuri, and John McLaren (2008). "Delay and Dynamics in Labor Market Adjustment: Simulation Results." *Journal of International Economics* 75(1): 1-13.

[4] Artuç, Erhan and John McLaren (2015). "Trade Policy and Wage Inequality: A Structural Analysis with Occupational and Sectoral Mobility." *Journal of International Economics* 97(2): 278-94.

[5] Autor, David H., David Dorn, and Gordon H. Hanson (2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103(6): 2121-68.

[6] Autor, David H., David Dorn, and Gordon H. Hanson (2017). "When Work Disappears: Manufacturing Decline and the Falling Marriage-Market Value of Men." NBER Working Paper #23173 (February).

[7] Becker, Gary S. (1957). "The Economics of Discrimination," Chicago: University of Chicago Press.

[8] Black, Sandra E. and Elizabeth Brainerd (2004). "Importing equality? The impact of globalization on gender discrimination," *Industrial and Labor Relations Review* 57(4): 540-559.

[9] Browning, Martin, Francois Bourguignon, Pierre-André Chiappori, and Valerie Lechene (1994). "Income and outcomes: A structural model of intrahousehold allocation," *Journal of Political Economy* 102(6): 1067-96.

[10] Brussevich, Masha (2018). "Does Trade Liberalization Narrow the Gender Wage Gap? The Role of Sectoral Mobility." *European Economic Review*, March.

[11] Cameron, A. Colin, Gelbach, Jonah B., Miller, Douglas L. (2006). "Robust inference with multi-way clustering." NBER Technical Working Paper 0327.

[12] Chetverikov, Denis, Bradley Larsen, and Christopher Palmer (2016). "IV Quantile Regression for Group-Level Treatments, With an Application to the Distributional Effects of Trade," *Econometrica* 84(2): 809-833.

[13] Clements, Benedict, Kamil Dybczak, and Mauricio Soto (2016). "Older and Smaller," *Finance & Development* 53 (1).

[14] Costinot, Arnaud, and Jonathan Vogel (2015). "Beyond Ricardo: Assignment Models in International Trade." *Annual Review of Economics* 7:1, pp. 31-62.

[15] Dai, Mi, Wei Huang, and Yifan Zhang (2018). "How Do Households Adjust to Trade Liberalization? Evidence from China's WTO Accession," IZA Discussion Paper No. 11428.

[16] Gaddis, Isis and Janneke Pieters (2014). "The Gendered Labor Market Impacts of Trade Liberalization: Evidence from Brazil," World Bank Policy Research Working Paper 7095.

[17] Hakobyan, Shushanik and John McLaren (2016). "Looking For Local Labor Market Effects of NAFTA." *The Review of Economics and Statistics* 98(4).

[18] Hakobyan, Shushanik and John McLaren (2018). "NAFTA and the Wages of Married Women." NBER Working Paper #24424 (March).

[19] Jones, Ronald I. (1971). “A Three-factor Model in Theory, Trade and History.” In *Trade, Balance of Payments and Growth: Essays in Honor of C.P. Kindleberger*, ed. by Bhagwati et al. Amsterdam: North Holland.

[20] Juhn, Chinhui, Gergely Ujhelyi, and Carolina Villegas-Sanchez (2014). “Men, Women, and Machines: How Trade Impacts Gender Inequality.” *Journal of Development Economics*, 106: 179-193.

[21] Kovak, Brian K. (2013). “Regional Effects of Trade Reform: What is the Correct Measure of Liberalization?” *American Economic Review* 103(5): 1960-76.

[22] Papyrakis, Elissaios, Arlette Covarrubias, and Arjan Verschoor (2012). “Gender and Trade Aspects of Labour Markets,” *The Journal of Development Studies* 48(1): 81-98.

[23] Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. Integrated Public Use Microdata Series: Version 6.0 [Machine-readable database]. Minneapolis: University of Minnesota, 2015.

[24] Sauré, Philip, and Hosny Zoabi (2014). “International Trade, the Gender Wage Gap and Female Labor Force Participation,” *Journal of Development Economics* 111: 17-33.

[25] Topa, Giorgio (2001). “Social interactions, local spillovers and unemployment,” *The Review of Economic Studies* 68(2): 261-295.

Appendix

Proof of Proposition 1.

Proof. For values of w^w and w^h along a ray through the origin, $U(w^w + w^h) - U(w^h)$ can be written as $U(y) - U(\kappa y)$, with $\kappa \in (0, 1)$ a constant. For the proposition, it is sufficient to show that this function is a decreasing function of y. First:

$$\frac{d}{dy}\left[U(y) - U(\kappa y)\right] = U'(y) - \kappa U'(\kappa y). \tag{13}$$

This derivative is negative if and only if $U'(y) < \kappa U'(\kappa y)$. For this, it is sufficient that $\kappa U'(\kappa y)$ be a decreasing function of κ (since $\kappa < 1$). Note:

$$\frac{d}{d\kappa}\left[\kappa U'(\kappa y)\right] = U'(\kappa y) + U''(\kappa y)\kappa y. \tag{14}$$

This is negative if and only if

$$\frac{U''(\kappa y)\kappa y}{U'(\kappa y)} < -1, \tag{15}$$

which is the stated condition. $\square$

Figure A1: Weighted Average US Tariff Rates: MFN and NAFTA

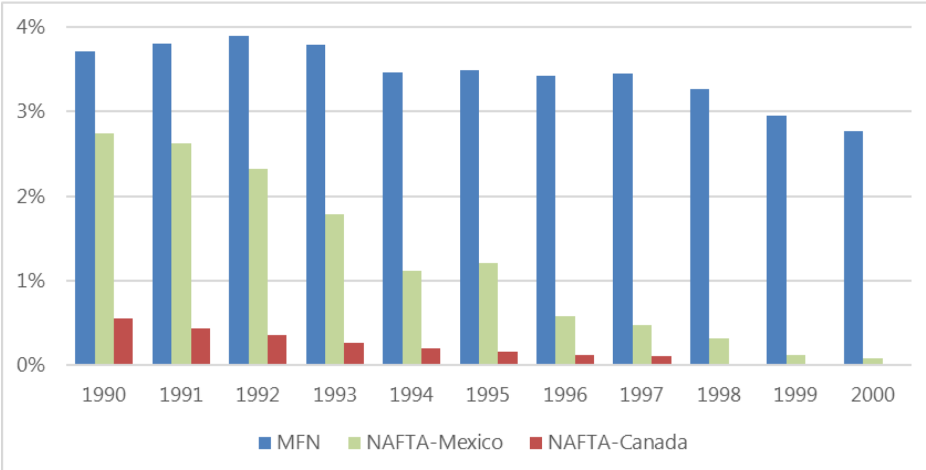

Figure A2: Weighted Average Canadian Tariff Rates: MFN and NAFTA-US

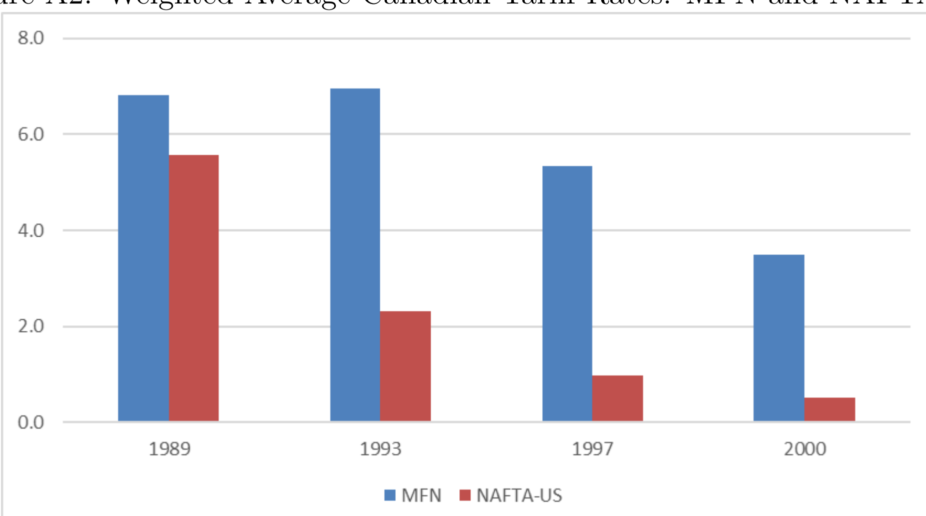

Table A1: Total Income (positive business and farm income for zero-wage earners)

	Married workers		Single workers	
	Male	Female	Male	Female
	(1)	(2)	(3)	(4)
Location effect				
Less than high school	-0.79	-3.52***	-1.276	-1.46
High school graduate	-0.583	-0.785	-0.416	-0.82
Some college	-1.042*	-1.576**	-0.818*	-0.506
College graduate	-0.206	-2.169*	-0.44	-0.386
Industry effect				
Less than high school	-1.099**	-4.004***	-1.898***	-1.906*
High school graduate	-0.089	-3.267***	-0.587	-0.487
Some college	-0.001	-1.81*	0.1996	0.479
College graduate	0.067	2.446*	-0.9589	1.0307
N of Observations	4,129,994	3,300,173	1,671,062	1,810,199

Notes: The table reports the overall impact on male and female income (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

Table A2: Weekly Wage

	Married workers		Single workers	
	Male	Female	Male	Female
	(1)	(2)	(3)	(4)
Location effect				
Less than high school	0.463	-1.453***	1.394***	-0.03
High school graduate	0.125	0.552	0.917*	0.808*
Some college	-0.418	-0.535	-0.434	-0.028
College graduate	0.231	-0.974	0.129	0.128
Industry effect				
Less than high school	-0.358	-2.91***	-0.801*	-1.631**
High school graduate	-0.006	-2.494***	0.117	-0.085
Some college	0.122	-1.058**	0.556*	0.5951
College graduate	0.661	2.351***	0.28	1.188**
N of Observations	3,803,221	3,117,140	1,563,108	1,744,870

Notes: The table reports the overall impact on male and female weekly wages (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

Table A3: Excluding workers over 55

	Married workers		Single workers	
	Male	Female	Male	Female
	(1)	(2)	(3)	(4)
Location effect				
Less than high school	-1.47*	-4.23***	-2.14***	-2.23
High school graduate	-0.556	-0.97*	-0.057	-0.595
Some college	-0.902	-1.607**	-0.804	-0.2337
College graduate	0.037	-2.059	-0.657	-0.64
Industry effect				
Less than high school	-1.332**	-4.178***	-2.157***	-2.627**
High school graduate	-0.357	-3.458***	-0.734	-0.433
Some college	-0.079	-1.813	0.026	0.844
College graduate	0.225	2.62*	-0.33	0.908
N of Observations	3,533,254	2,937,650	1,553,640	1,593,447

Notes: The table reports the overall impact on male and female wages (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

Table A4: Married workers between ages 25 and 64

	Married workers	
	Male	Female
	(1)	(2)
Location effect		
Less than high school	-0.98	-3.95***
High school graduate	-0.845*	-0.595
Some college	-1.121**	-1.389
College graduate	-0.348	-1.923*
Industry effect		
Less than high school	-0.791**	-3.753**
High school graduate	-0.0916	-3.263***
Some college	-0.112	-1.563*
College graduate	0.188	1.616
N of Observations	3,555,277	2,914,875

Notes: The table reports the overall impact on male and female wages (computed as a difference between initial-tariff and impact effects) for each education group when a location or an industry loses all of its protection within the sample period. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

Table A5: Pooled Regression for Male and Female Workers

	Married workers			Single workers		
	Male	Female	Difference Statistically Significant?	Male	Female	Difference Statistically Significant?
	(1)	(2)		(3)	(4)	
Location effect						
Less than high school	-1.19	-3.272***	**	-0.477	-2.664**	***
High school graduate	-1.37**	0.049	*	-1.022	-0.324	
Some college	-1.383*	-1.18		-1.195*	-0.193	*
College graduate	0.238	-2.478*	*	0.718	-0.956	**
Industry effect						
Less than high school	-0.981**	-3.578**	*	-1.56*	-1.737	
High school graduate	-0.244	-2.911***	***	-0.524	-0.437	
Some college	-0.184	-1.66*		0.053	0.627	
College graduate	0.15	1.353		-1.093	0.882	**
N of Observations	6,920,361			3,307,978		

Notes: The table reports the overall impact on male and female wages for each education group when a location or an industry loses all of its protection within the sample period, obtained from a single regression where tariff measures, industry, location and occupation fixed effects are interacted with female dummy and female effect is computed as the sum of the net effect for male and female workers. The statistical significance of the difference between the overall impact on male and female wages is reported in the respective column for each regression. The regression results used to compute these point estimates are available from the authors upon request. We omit agriculture by setting its tariff equal to zero. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

Table A6: Top 10 Most Vulnerable Conspumas (by gender)

Rank	State	Counties/Cities	$loc\tau^c_{1990}$ (%)	$loc\triangle\tau^c$
		Panel A: Female local tariffs		
1	North Carolina	Alamance, Randolph	2.27	-2.21
2	South Carolina	Oconee, Pickens	2.09	-2.02
3	North Carolina	Alexander, Burke, Caldwell	2.07	-1.91
4	Georgia	Catoosa, Dade, Walker	2.03	-1.72
5	Missouri	Douglas, Howell, Texas, Oregon, Wright, Ozark, Shannon	1.91	-1.21
6	Virginia	Danville, Pittsylvania	1.82	-1.75
7	North Carolina	Catawba	1.79	-1.75
8	South Carolina	including Cherokee, Chester, Chesterfield, Clarendon, Darlington	1.74	-1.69
9	North Carolina	Cabarrus, Rowan	1.69	-1.65
10	North Carolina	Cleveland, McDowell, Polk, Rutherford	1.68	-1.62
		Panel B: Male local tariffs		
1	Georgia	Catoosa, Dade, Walker	2.71	-2.24
2	Indiana	Lake	2.71	-2.14
3	Indiana	Lake	2.38	-1.91
4	South Carolina	Anderson	2.18	-2.06
5	Indiana	Porter	2.16	-1.74
6	South Carolina	Oconee, Pickens	2.15	-2.09
7	North Carolina	Alamance, Randolph	2.14	-2.08
8	Indiana	Lake	2.07	-1.67
9	Michigan	Genesee	1.95	-1.93
10	South Carolina	including Cherokee, Chester, Chesterfield, Clarendon, Darlington	1.93	-1.84