The Gender Wage Gap and the Child Penalty*

Giulia Olivero

Cornell University

November 2024

Job Market Paper

Click here for the most recent version

Abstract

Child penalties account for most of the remaining gender inequality in the labor market. Yet, we still do not know much about why they remain so large and persistent. I start by documenting a novel fact, which is the presence of heterogeneity in child penalties in the US by measures of intra-household comparative advantage. Then, I investigate the effect of the closing of the gender wage gap on employment penalties for mothers over the years 1980-2010. To do so, I leverage gender differences in occupational choices and combine gender-specific local labor market shocks with pseudo-event studies around childbirth. I find evidence of a greater fall in child penalties in local labor markets with a faster convergence in the wage rate of women and men. I explore possible mechanisms and find evidence of an increase in education of women (relative to men), delayed childbirth, and suggestive evidence of a shift in gender attitudes.

*I am grateful for precious guidance and feedback from Doug L. Miller, my dissertation committee chair. I am further grateful for comments and/or help from Francine Blau, Megan Hyland, Hyewon Kim, Henrik Kleven, Eleonora Patacchini, Evan Riehl, Seth Sanders, Na'ama Shenhav, Revathy Suryanarayana, Meredith Welch, Yalin Xu, and participants in the Cornell Labor Works in Progress Seminar and the Cornell Labor Workshop. I am especially thankful to Francesco Billari for his insightful suggestions on the theoretical framework. I acknowledge institutional support through a research support grant from the Cornell Population Center. All errors are my own.

1 Introduction

Despite outperforming men in most measures of educational attainment, women still earn significantly less (Blau and Kahn, 1996, 2017; Adda, Dustmann and Stevens, 2017; Albanesi, Olivetti and Petrongolo, 2022). Recent evidence suggests that child penalties (the causal effect of parenthood on employment outcomes of women relative to men) account for most of the remaining gender inequality in the labor market[1] (Kleven, Landais and Søgaard, 2019a; Kleven, Landais, Posch, Steinhauer and Zweimuller, 2019b; Cortés and Pan, 2023; Kleven, 2024; Goldin, Kerr and Olivetti, 2022). The mechanisms through which parenthood has a long-lasting effect on women's labor market trajectories have been largely explored and include, among others, the reduction in hours worked, the loss of skills during interruptions, lower accumulation of experience, and even selection into lower-pay child-friendly occupations (Adda, Dustmann and Stevens, 2017).

However, a still unanswered crucial question is why child penalties remain so large and persistent, even in modern societies. One candidate explanation is rooted in the idea of comparative advantage. According to Becker's model of household specialization, as women's labor market opportunities converge to men's, there is less reason to expect specialization to remain starkly gendered. In line with the model's predictions, recent work shows that when the gender wage gap decreases (whether due to increased economic independence of women, or adverse economic shocks to men), the marriage rate declines and couples increasingly match on similarities, rather than potential gains from trade (Shenhav, 2021; Anelli, Giuntella and Stella, 2021; Autor, Dorn and Hanson, 2019; Juhn and McCue, 2017). What about the arrival of children,

[1] I use the term "penalty" for consistency with a long-established body of research. The term specifically refers to the drop in employment experienced by women right after giving birth. It encompasses various mechanisms, which can include both supply-side choices (e.g., women choosing to work fewer hours, take career interruptions, or select into more flexible but lower-paying jobs) and demand-side factors (e.g., employer discrimination against mothers).

though? Although Becker's model predicts declining specialization upon childbirth as well, descriptive work does not seem to support this hypothesis (Juhn and McCue, 2017; Kleven, Landais and Søgaard, 2021). However, the existing evidence is limited to specific contexts, and efforts to address this question within a causal framework have been scarce. This is partly due to data limitations that afflict the measurement of child penalties, as well as the complexities arising from endogeneity concerns.

In this paper I contribute to the literature on the determinants of child penalties by asking: 1) Do child penalties differ by measures of intra-household comparative advantage?; and 2) Did historical shifts in the average female-to-male wage ratio across US labor market have a causal effect on the evolution of the child penalty in employment? My analysis is composed of two main building blocks. First, I document a novel fact, which is the presence of heterogeneity in the size of the child penalty by degree of intra-household comparative advantage in the US. My results are robust to alternative definitions of intra-household comparative advantage and differ from existing evidence which, so far, has suggested that comparative advantage is not a meaningful determinant of child penalties (Kleven, Landais and Søgaard, 2021). Second, I look at the causal effect of the evolution of the female-to-male wage ratio within local labor markets on the child penalty over the period 1980-2010. Figure 1 shows the evolution of the average female-to-male hourly wage in the US, as well as the evolution in child penalty in employment. Both trends show substantial progress towards greater gender equality up until the mid 1990s, after which convergence seems to plateau. The goal of this paper is to establish what is the direction and magnitude of the causal relationship between these two trends.

To answer these questions, I must address two main empirical challenges. The first challenge is that measuring child penalties typically requires longitudinal data; however, traditional panel datasets often have limited sample sizes, making it difficult to estimate penalties at granular geographic levels. Kleven (2024) proposes a method to estimate child penalties using multiple adjacent cross-sections of data from large datasets, such as the ACS or CPS. Unfortunately, these datasets have limitations: the CPS lacks sub-state geographical identifiers, and the ACS is only available starting in 2005. Here, I build on this existing method by proposing and validating a way to measure child penalties using one single cross-section of data. I demonstrate that this method yields accurate estimates and, crucially, can be applied to Decennial Census data, which provides sub-state geographical identifiers. This approach allows for the computation of child penalties at the level of local labor markets, which I define as a commuting zone (CZs)-by-year.

The second challenge is that relative wages of women and men are not randomly distributed across labor markets. To address this, I construct a sex-specific proxy for potential wages in each local labor market using a Bartik-style instrument. By keeping the industrial and occupational structure of each local labor market fixed at its 1970 level (prior to the period of analysis), this proxy captures the component of wages driven by national shifts in labor demand of men and women over the period from 1980 to 2010. During this time, there was disproportionate wage growth for cognitive and people-oriented occupations, in which women held a comparative advantage. As a result, wages increased more for women (relative to men) in local labor markets where certain service sector industries were prevalent (Shenhav, 2021).

I find evidence of a larger decline in child penalties in employment in local labor markets that experienced a greater convergence in hourly wages of women and men. Specifically, a 10% increase in the potential relative wage of women and men (which is approximately the size of the great gender convergence over this period) is associated with a 15-16pp decrease in the child penalty in employment for mothers. I explore potential mechanisms and find that the effect is partially explained by a reduction in the educational gap of parents (the average mother-father difference in years of education increases), and by a delay in childbirth. I also find evidence of a shift in gender norms in those labor markets in which employment opportunities improved more for women relative to men.

This paper adds to the recent but fast-growing literature on the determinants of child penalties which, so far, seems to rule out explanations rooted in biology and government policy. Kleven et al. (2021) find that, in Denmark, penalties are equally large for adoptive mothers as they are for biological ones. In another study based on Austrian administrative data, (Kleven, Landais, Posch, Steinhauer and Zweimuller, 2019b) show that, in the long-run, parental leave and childcare policies have little or no effect on child penalties, although increasing the duration of paid and job-protected leave implies larger penalties in the short-run. Similarly, Albanesi et al. (2022) compare data from 24 countries and find very limited evidence of beneficial effects of longer parental leave on maternal participation and earnings, although more generous support for childcare seems to play a more important role for female participation.

What, then, drives child penalties? Recent studies increasingly highlight gender norms as key determinants of child penalties, both across and within countries (Boelmann, Raute and Schonberg, 2021; Kleven, 2024). This raises important questions: What are gender norms, how are they measured, and what do they capture? While much of the existing literature has defined gender norms using elicited measures from survey questions - often treating them as static or slow-moving cultural traits - recent work by Kleven, Olivero and Patacchini (2024) conceptualizes gender norms through behaviors rather than attitudes, and shows that highly localized variation in exposure to gender norms can have significant impacts.

My work contributes to the literature on the fundamental determinants of child penalties by looking at a specific feature of local labor markets which, so far, has not received much attention within this literature: wages. Precisely, this paper investigates whether child penalties respond to economic incentives, more precisely to shifts in the female-to-male wage ratio within a local labor market. The evidence on whether the labor supply of mothers responds to shifts in labor demand is extremely limited. One notable exception is Kuka and Shenhav (2023) who leverage variation in post-birth work incentives created by timing of birth and EITC eligibility and find evidence of a reduction in the child penalty for affected mothers.

Additionally, by finding suggestive evidence that shifts in gender norms mediate the relationship between economic incentives and child penalties, I bridge the classic labor economics literature on household specialization and labor supply with the more recent literature on child penalties, which emphasizes gender norms as the ultimate determinant of household behavior and labor market outcomes for parents.

My work also contributes to the literature on the effect of the relative wage for women and men which, so far, has focused on outcomes related to household formation and, more specifically, decision to marry (Anelli, Giuntella and Stella, 2021; Shenhav, 2021; Autor, Katz and Kearney, 2008). This paper contributes to that literature by moving the question one step further and assessing whether, for those households that do still form, an increase in the relative wage has any effects on the degree of household specialization upon childbirth.

Finally, I contribute to the emerging literature on the measurement of child penalties by proposing and validating a method that uses one single cross-section of data, building on Kleven (2024)'s method that leverages multiple adjacent cross-sections. To the best of my knowledge, work is the first to produce measures of child penalties in the United States at the sub-state geographical level.

This paper proceeds as follows. Section 2 lays out the theoretical motivation. Section 3 provides an overview of my data sources. I talk about the estimation of child penalties in section 4. My main results are divided into two main parts. In section 5 I document a novel fact, which is the presence of heterogeneity in child penalties in employment by measures of intra-household comparative advantage. In Section 6 I leverage exogenous shocks in the potential wages of women and men across different labor markets by adopting a Bartik-style empirical framework, and I shows how the closing of the gender wage gap affected child penalties in employment at the CZ-by-year level. I also explore potential mechanisms. Section 7 concludes.

2 Hypothesis Development

To provide an intuition on the relationship between the effect of the female-to-male wage ratio on the intra-household division of market and non-market work, and how that relationship can be mediated by the presence of norms, I propose a simple version of a model in the spirit of Becker, building on existing work (Jones et al., 2015; Shenhav, 2021; Bertrand et al., 2021).

I introduce a utility maximization problem for a married household consisting of two members: one male (m) and one female (f). I assume that households are homogeneous, and that the bargaining problem is resolved efficiently within the household. The household members are endowed with one unit of time each, and must decide how to allocate it between working in the labor market (t_f^W, t_m^W), the production of goods at home - which includes housework and child-rearing (t_f^H, t_m^H), and leisure (t_f^L, t_m^L). The utility of the household depends on the consumption of market goods (c_f^W, c_m^W), a joint home-produced good (c^H), and the amount of leisure (t_f^L, t_m^L). Additionally, I assume that a household's utility is influenced by a societal gender norm t^*, which dictates an expected level of female labor supply. Deviations from this norm are penalized, reflecting societal pressure to conform, with parameter γ capturing the strength of the penalty.

Hence, the household maximizes this utility function:

$$U = \log(c_f^W) + \log(t_f^L) + \log(c_m^W) + \log(t_m^L) + \log(c^H) - \gamma(t_f^W - t^*)^2$$

Subject to the following constraints:

$$c_f^W + c_m^W \le (1-\tau_f)wt_f^W + (1-\tau_m)wt_m^W \tag{i}$$

$$c^H \le At_f^H \tag{ii}$$

$$t_f^H + t_f^W + t_f^L \le 1 \tag{iii}$$

$$t_m^W + t_m^L \le 1 \tag{iv}$$

i *Budget constraint:* The wage rate w is assumed to be equal for simplicity. However, I assume that household members face sex-specific tax rates τ_f and τ_m. Specifically, $1-\tau_f = (1-\tau_d)(1-\tau_m)$ where τ_m represents the common labor income tax rate and τ_d represents the additional wedge faced by a female in the labor market (Jones, Manuelli and McGrattan, 2015). This wedge can be interpreted as a proxy for direct wage discrimination, or the shadow value on a constraint that restricts a woman's job opportunities (such as a glass ceiling). Although the details are not explicitly modelled here, a decrease in this wedge can be the result of changes in regulations relating to discriminatory practices, as well as the result of changes in sex-specific productivity.

ii *Home production constraint:* Where t_f^H is the time the wife spends producing the public good, and A is the productivity parameter. Note that, to keep things simple, and in line with empirical evidence, I assume that $t_m^H = 0$ which means that husbands only allocate their time between market work and leisure.

iii-iv *Time constraints:* both agents are endowed with one unit of time.

In order to derive predictions, I simplify this problem and obtain the two following first-order conditions (see Appendix A for the derivation):

$$t_f^W: \quad \frac{2(1-\tau_f)w}{(1-\tau_f)w \cdot t_f + (1-\tau_m)w \cdot t_m} - \frac{2}{1-t_f} - 2\gamma(t_f - t^*) = 0 \tag{FOC1}$$

$$t_m^W: \quad \frac{2(1-\tau_m)w}{(1-\tau_f)w \cdot t_f + (1-\tau_m)w \cdot t_m} - \frac{1}{1-t_m} = 0 \tag{FOC2}$$

From these first order conditions, the following facts can be derived (see Appendix A for more detail):

- Changes in w, despite changing levels of consumption, do not affect the spouses' allocation of hours to any of the activities (market work, home production, or leisure).
- If either τ_m or τ_f change, keeping the other constant, hours adjust. In particular, t_f^W (t_m^W) increases if either τ_f (τ_m) decreases, or if τ_m (τ_f) increases. In other words, in response to a reduction in market discrimination by gender, a wife works more in the market, a man works less. If τ_m and τ_f both change proportionally, with $(1-\tau_f)/(1-\tau_m)$ fixed, no change in hours take place.
- As the penalty associated with deviation from gender norms (γ) increases, the response of t_f^W to τ_f diminishes. This means that, when societal pressure to conform to the gender norm is extremely strong, changes in taxes have virtually no effect on female labor supply because the household is primarily concerned with avoiding deviations from the norm.

Hence, in this model, the presence of the gender wage gap causes wives to specialize and allocate a substantial fraction of their time to home production.

Figure 2 graphically shows how the female labor supply (y-axis) changes for different levels of female (τ_f) and male (τ_m) tax rates, as well as for different penalties when deviating from the gender norm (γ). The norm about female labor supply is set at 0.1 (horizontal black dotted line). The first thing to note is that, as the male tax rate increases (moving from left panel to right panel), the overall female labor supply increases. Second, for a given level of male tax rate (τ_m), the female labor supply decreases as the female tax rate (τ_f) increases (moving from left to right on the horizontal axis). Third, the responsiveness of the female labor supply to the female tax rate (the slope of the curve) decreases as the penalty associated with deviations from the gender norm increases (moving from blue line, to red, to yellow). For very high levels of γ, the female labor supply is basically unresponsive to changes in tax rates.

According to this model's predictions, a decrease in the gender wage gap $(1 - \tau_d)$ could result in two possible empirical results, depending on how large is the penalty associated with deviating from gender norms (γ):

1. *For low enough values of γ, when the female-to-male wage ratio goes up, child penalties in employment decrease, in line with Becker model of household specialization.*

2. *For higher values of γ, when the female-to-male wage ratio goes up, child penalties in employment remain the same, because norms are all that matters, in line with evidence from recent literature.*

Hypothesis (1) draws on Becker principles of household specialization and is grounded on the idea that, as women's labor market opportunities increase and become more similar to men's, the opportunity cost of staying home increases and there is less of a reason to expect specialization to be starkly gendered. However, consistently with (2), existing work shows that child penalties are strongly correlated with social norms, which are regarded as slow to change and resisting economic forces (Fortin, 2005; Bertrand, Kamenica and Pan, 2015; Kleven, 2024). In line with this, Kleven, Landais and Søgaard (2021) descriptively show that the size of the child penalty in Denmark does not depend on the earnings potential of the two spouses.

Next, I move to the empirical test of these two alternative hypotheses.

3 Data

In my main analysis I combine data from the IPUMS Decennial Census 1970-2000, and ACS data for the year 2010 (Ruggles et al., 2010). Additionally, I use the National Longitudinal Survey of Youth 1979 (NLSY79) to validate my event study estimates. Finally, I use data from the General Social Survey (GSS) to compute measures of elicited gender norms. I run my analyses at the local labor markets level, which I define as a CZ-by-year (commuting zone by year). Commuting zones (CZs) are an aggregation of county-level data intended to reduce spatial auto-correlation and to reflect more closely the geographic interrelationships between employers and labor supply (Tolbert and Sizer, 1996). This is how I construct my main variables:

- *Actual wages:* To create a measure of actual hourly wages in each CZ-by-year, I use 5% Decennial Census data from IPUMS for the years 1980-2000 and ACS data for the year 2010.Hourly wages are defined as total yearly pre-tax wage and salary income divided by the product of weeks worked per year and usual hours worked per week. The analysis focuses on a sample of working-age individuals (18-64) who reported positive earned income and positive hours worked. I exclude self-employed individuals, those with imputed wage or hours data, individuals in group quarters, and those who report being in school. In addressing top-coding, I closely follow the approach of Autor et al. (2008) (see Appendix B.1 for more details).

- *Potential wages:* To create a measure of potential wages I construct a Bartik-style instrument following previous work (Shenhav, 2021; Bertrand et al., 2015; Autor et al., 2019). Specifically, I obtain baseline industry and occupation-by industry shares from the IPUMS 1% Decennial Census 1970 (see Section 5 for an explanation of how I construct this instrument). Occupation-by-industry shares and wages are obtained from 5% Decennial Census data for the years 1980-2000 and ACS data for the year 2010.

- *Child penalties:* Child penalties are computed from 5% Decennial Census data from IPUMS for the years 1980-2000 and the 5-years ACS data for the year 2010[2] (Ruggles et al., 2010). I only keep individuals observed between age 18 and 55 who are not in school and such that their age at first birth was between 22-45. Event-study methodology is validated using pooled data from the National Longitudinal Survey of Youth 1979 (NLSY79).

- *Gender norms:* An index of gender norms progressivity is constructed using information about elicited gender norms from the General Social Survey (GSS), following the procedure described in Kleven (2024).

4 Measuring Child Penalties

The first step in my empirical analysis is to estimate child penalties. In Section 4.1, I provide an overview of standard approaches for estimating child penalties using longitudinal data. Section 4.2 summarizes the recent "pseudo-panel" technique for estimating child penalties with multiple cross-sections of data and extends it to a single cross-section. Finally, in Section 4.3, I validate the method introduced in Section 4.2 using the NLSY79 longitudinal dataset.

4.1 Traditional Estimation

Child penalties are defined as the causal impact of having a child on the labor market outcomes of women relative to men. Typically, the measurement of child penalties follows the event

[2] The 5-years ACS data pools individuals surveyed over the years 2006-2010. This is to increase sample size, since with one only year of data the estimates can get fairly noisy.

study approach proposed by Kleven et al. (2019b). Intuitively, the goal is to capture the extent to which employment outcomes for women and men diverge after the birth of their first child, netted out of their pre-birth difference. Empirically, this means estimating the following equation on a panel of parents, separately by gender:

$$Y^g_{itj} = \sum_{j=-5, j\neq -1}^{j=10} \alpha^g_j \cdot D_{i,t-j} + \sum_k \beta^g_k \cdot \mathbb{I}[k = age_{it}] + \delta^g_t + \epsilon_{itj} \tag{1}$$

where Y^g_{itj} is the outcome for individual i of gender g observed in year t, j years away from having a child. The index j indicates the time relative to the event (the birth of the first child) which, in this case, ranges between 5 years before to 10 years after. δ^g_t represents calendar-year fixed effects, while $\sum_k \mathbb{I}[k = age_{it}]$ is a vector of age-in-years fixed effects. These two sets of fixed effects control non-parametrically for life-cycle trends and time trends. The conditions for causal identification in this framework were laid out and validated against IV-approaches in Kleven et al. (2019a).

The main coefficients of interest, α^g_j, dynamically trace the change in the employment outcomes at each event time, relative to the event time of reference (in this case -1, the year before the birth of the child) conditional on year and age fixed effects[3].

Then, for each gender g and event-time j, the estimated level effects are converted into percentages effects by calculating:

$$P^g_j = \frac{\hat{\alpha}^g_j}{E[\tilde{Y}^g_{itj}|j]} \tag{2}$$

where $\tilde{Y}^g_{itj}$ is the predicted counterfactual outcome (obtained as the predicted outcome when omitting the contribution of the coefficient of the relative time $D_{i,t-j}$). Finally, the penalty is

[3]In a traditional panel fixed-effects event-study framework in which only treated individuals (parents) are used in the estimation, one additional restriction needs to be implemented in order to restore identification. In this specification, additional restrictions are not strictly necessary since individual fixed effects are not included.

obtained as the difference in average effects for women and men across positive event-times:

$$ChildPenalty = \underbrace{E[P_j^m - P_j^w | j \geq 0]}_{\text{post-birth difference}} \tag{3}$$

where P_j^m and P_j^w are event-time specific penalties for men and women, respectively.

4.2 Estimation with sparse cross-sectional data

Estimating child penalties from equation (1) traditionally requires large-scale panel data that track individuals throughout their adult life and childbearing years. The intensive data requirement is one of the main reasons why this literature has been relatively limited until recent years. Kleven (2024) proposed a novel method to estimate child-penalties using a pseudo-panel approach based on multiple adjacent waves of cross-sectional data, such as the Current Population Survey (CPS) and the American Community Survey (ACS) . The key idea is to link individuals that exhibit similar demographic characteristics across adjacent survey years to create pseudo-longitudinal observations. An advantage of this approach is that it leverages much larger sample sizes than traditional panel datasets, such as the PSID, SIPP, or NLSY, which allows for more precise estimates. This method opened new doors to cutting-edge research on child penalties and marked the first effort to produce state-level child penalty estimates in the U.S. (and country-level estimates elsewhere).

In this paper, I build on this method and propose a way to measure child penalties using only a single cross-section of data, specifically the decennial census. One key reason for using the decennial census is that, compared to higher-frequency cross-sectional surveys such as the CPS and ACS, the census provides sub-state geographical identifiers it provides sub-state geographical identifiers and covers pre-2005 years. These identifiers make it theoretically possible to construct child penalty estimates at a more granular geographical level, such as commuting zones (CZs) which, to the best of my knowledge, has not been explored before.

An additional benefit of estimating child penalties from one cross-section of data follows naturally from its design. Since child penalties are inherently longitudinal processes, defining them in a specific year t can be ambiguous - for example, should t refer to the year of birth, the year before birth, or the year when the decision to have a child is made? Cross-sectional child penalties avoid this ambiguity by providing a cleaner and more straightforward way to study how certain factors X_t affect child penalties at time t. While this feature was not the original goal of the method, it turns out to be a useful property that simplifies the analysis. Lastly, another practical advantage of this approach (compared to the pseudo-panel approach based on multiple cross-sections) is that it significantly reduces computational complexity, making it faster and more efficient for large-scale analysis.

The main challenge with cross-sectional data is that I can only observe parents after the birth of their child, along with childless individuals (some of whom may have children in the future). As a result, I do not directly observe parents' pre-birth employment outcomes. Following the approach of Kleven (2024), I match parents with childless individuals who share similar demographic characteristics. However, since I rely on a single cross-section, I perform the matching within the same census year rather than combining data from multiple years.

In practice, this is how it works. For each census year (e.g. 1980) I first select individuals who report having a child. I use the age of their eldest child to infer the event-time at which they are being observed (where event-time=0,...,10) , and their age-at-birth. I limit myself to individuals observed between the age of 18 and 55, and who had their first child between the ages of 22 and 45. Then, I take parents observed at event-time 0 (i.e. those with a newborn) and age a, and I match them to childless individuals in that same CZ with similar demographic characteristics and age $a - l$ where l=1,...,5. Since each parent can be matched to several childless individuals, each match is attributed a weight of $\frac{1}{k}$ where k is the total number of matches for that parent. The following variables are used for matching: gender, race/ethnicity, educational attainment, and marital status[4]. For married individuals, I also include a matching variable that captures their degree of intra-household comparative advantage, based on their spouse's characteristics[5]. After constructing a pseudo-panel of parents observed from five years before to ten years after the birth of their child, I estimate the following equation, separately by gender:

$$Y_{itj}^{g} = \sum_{j=-5, j\neq -1}^{j=10} \alpha_j^g \cdot D_{i,t-j} + \sum_k \beta_k^g \cdot \mathbb{I}[k = age_{it}] + \epsilon_{itj} \tag{4}$$

where equation (4) is defined in the same way as equation (1) with the only exception that year fixed-effects have been removed.

There are two important points worth highlighting here. The first concerns how this new specification subtly, but fundamentally, changes the interpretation of the event study coefficients. Traditionally, child penalties are understood as changes in employment trajectories for a real cohort of parents, where the cohort is defined by the year of the child's birth. In contrast, my approach estimates the child penalty for a synthetic cohort of individuals—i.e., the changes in employment trajectories they would experience if they went through their childbearing years *all at once, at time t.*

Borrowing from basic demography concepts, it may be helpful to draw an analogy to two alternative methods of computing life expectancy from life tables: the cohort approach and the period approach. A cohort life table shows the probability of death for individuals born in a

[4]Following Kleven (2024), I match on these categories: education (some college or less, and a college degree or more), Race (White non-Hispanic, Black non-Hispanic, Hispanic), marital status (married with spouse present, and everyone else (divorced, widowed, or never-married)

[5]See Section 5 for a more detailed description of how this is measured. Also note that intra-household comparative advantage can not be estimated for respondents who do not have a spouse in the household, hence they are still included in the overall analyses, but they are assigned a separate NA value for comparative advantage.

given year, based on mortality rates observed throughout their lifetimes. In contrast, a period life table uses age-specific mortality rates from a single year (or group of years) and assumes that those rates will remain constant for the remainder of a person's life.

Although the period approach has limitations—both when estimating child penalties and life expectancies—because it does not account for future changes, it offers two key advantages. First, it enables me to construct child penalties at a more granular geographical level than previously possible. Second, the cross-sectional aspect of this approach is particularly useful and conceptually straightforward when estimating a panel fixed-effects regression of child penalties (as the dependent variable) on relative wages (as the independent variable) in a given year t. Finally, the validity of this approach rests on the following conceptual assumption:

Conditional on having the same observable demographic characteristics (gender, race, education, marital status, CZ of residence), a childless individual observed in the same CZ in year t at age $a - l$ provides a good proxy for the employment outcome of a parent who had their first child in year t at age a, l years priors to t.

Although there is no perfect way to test this assumption, I validate my approach using a real longitudinal dataset: the NLSY79.

4.3 Method Validation using NLSY79

To validate my event study method, I first pool together all the waves from the NLSY79 (1978-2019). Then, I select individuals (mothers and fathers) who have children anytime throughout the panel years and I only keep those who have their first child between the age of 22 and 45, in line with my previous sample selection. I create a variable "employed" equal to 1 if they report working a positive number of weeks (and 0 otherwise) and I collect it longitudinally throughout the survey waves. Then, I run the event study indicated in equation (1), on the real panel of parents. Figure 3 plots the estimated event study coefficients for mother and fathers, as well as the robust standard errors . The figure shows flat pre-trends both for mothers and fathers. As expected, while the employment trajectory of fathers is basically unaffected by the birth of their first child, the left panel of Figure 3 shows a sharp drop in the employment of mothers, which is exactly what we would expect.

Next, I apply the matching algorithm. In each survey year, I only keep parents post-birth, and I match them to childless individuals observed in that same survey year, who have similar demographic characteristics. I pool together several survey years to maximize my sample size and keep the sample selection as similar as possible to the one used in the estimation on the real panel. It is worth noting that I only allow matches with the never-parents observed in the dataset[6]. This might be overly conservative, but helps me mitigate two main concerns. First of all, since the NLSY79 has a relatively small sample size, I might end up with “overly-good" matches. Although it is mechanically impossible for an individual to be matched to themselves (since matches are created with younger individuals in the same survey wave) there is a concern that my overall pseudo-panel would end-up being too similar to the actual-panel and perform overly-well. Second, this should provide a worst-case scenario for the quality of my estimates from decennial census data (i.e. the case in which none of the matched childless individuals would end up becoming a parent in the future).

Figure 4 compares event study estimates obtained for the real panel and the pseudo panel obtained with my matching algorithm. Coefficients in the pseudo-panel tend to overestimate slightly the real-panel coefficients.

In particular, the estimation error seems to be the largest for fathers during their first post-child years (as indicated by that slight upward bump in the α_0^m,α_1^m, α_2^m estimates) . Existing work has consistently documented that, in the United States, fathers tend to have higher education and income than the average male population. What that small upward bump seems

[6]The idea, when matching individuals in ACS/CPS/Decennial Census, is that most (if not all) of the childless individuals will likely have a child in the future.

to suggest, is that the matching on observables I am currently performing might not fully capture that positive self-selection into fatherhood. However, the overall child penalty estimates are extremely similar: the real panel estimate is 20.25% while the pseudo-panel estimate is 19.31%, which is a very small margin of error.

5 Part I: Child Penalties by Intra-Household Comparative Advantage

Having created a pseudo-panel of parents and having validated the quality of my child penalty estimates, I proceed with the empirical analysis. In this first part, I examine whether child penalties vary according to the relative economic advantage of spouses within the household. Since I do not have a direct measure of comparative advantage, I construct a proxy using each spouse's share of the household's predicted potential earnings. I use potential earnings rather than actual earnings because actual earnings can be distorted by career interruptions related to motherhood, while a reliable analysis of comparative advantage requires an estimate of each partner's earning capacity that is independent of the presence of children.

To estimate potential earnings, I take decennial census 1980-2000 and 2010 5-years ACS and I run sex-by-year-by-CZ Mincer regressions that predict total earnings based on education, race, and experience. I limit the sample to individuals aged 18 to 55 who report positive earned income. The model is specified as:

$$\ln Y_{it} = \alpha \, \text{YrsEdu}_{it} + \beta_1 \, \text{Exp}_{it} + \beta_2 \, \text{Exp}_{it}^2 + \beta_3 \text{Race}_i + \nu_{it}, \tag{5}$$

where Y_{it} is the observed earnings of individual i in year t, YrsEdu_{it} refers to years of education, Exp_{it} is years of experience (calculated as $\text{Age}_{it} - \text{YrsEdu}_{it}$), and Race_i are individual race dummies. These regressions are conducted separately by gender. I predict earnings for men using the full sample of men, as their earnings are less likely to be affected by parenthood. For women, I restrict the sample to childless women to avoid distortions from motherhood-related career interruptions. Using information on household composition, I calculate a wife's share of household potential earnings as: $\frac{Pot\_Earnings\_wife}{Pot\_Earnings\_wife+Pot\_Earnings\_husband}$. Households are then split into tertiles based on the wife's share of potential earnings[7]. Figure 5 shows the distribution of households' comparative advantage[8], separately by Census year. Each distribution is approximately normal, with the bulk of observations ranging between $[0.45, 0.52]$. The compressed nature of the distribution is expected, as variation in potential wages is considerably smaller than the variation observed in actual wages.

Then, to quantify the heterogeneity of child penalties by comparative advantage, I estimate the following event-study regression:

$$Y^g_{itj} = \sum_{j=-5, j\neq -1}^{j=10} \alpha^g_j \cdot D_{i,t-j} \cdot 1[E_i = e] + \sum_k \beta^g_k \cdot \mathbb{I}[k = age_{it}] + \gamma_e + \eta_s + \theta_t + \epsilon_{itj} \tag{6}$$

Where the main event-study coefficients are interacted with tertiles of comparative advantage $[E_i = low, medium, high]$. I include CZ fixed effects η_s and year fixed effects θ_t to account for spatial and temporal variation. γ_e are fixed effects for tertile of exposure, and errors are clustered at the commuting zone level. Figure 6 presents the estimated interacted coefficients, revealing substantial heterogeneity in the employment drop for women based on their intra-household earnings potential. Women who earn a higher share of potential household earnings experience roughly a 30% reduction in employment, significantly smaller than the 38% drop observed for women with a lower share. These differences are precisely estimated and statistically significant.

[7]Note that the pseudo-panel is created by matching married individuals on comparative advantage, in addition to the other demographic characteristics. This is to ensure consistency in in the "type" of household in the pre-birth and post-birth observations

[8]Comparative advantage is defined in terms of the labor market. A low comparative advantage wife is one whose potential earnings represent a smaller share of the household's total earnings.

One reasonable concern is that using childless women to estimate female potential earnings might introduce bias. Some studies address this by using a sample of men to predict earnings for both genders (Kleven, Landais and Søgaard, 2021), arguing that childless women's earnings could still be influenced by future motherhood plans or fertility-related considerations. At the same time, given the persistent gender wage gap in the U.S. (even among individuals with similar characteristics) it is unclear whether men's earnings are a realistic counterfactual for women's. To address this concern, I repeat the analysis using spouses' difference in total years of education as an alternative measure of comparative advantage. Figure 7 shows the distribution of households by spousal educational gap. The distributions are centered around 0 (indicating no educational gap between spouses), with most observations falling within the range $[-10, +10]$. I split my sample into three groups: households where wives have less education than their husbands ($\Delta \leq 0$), households where they have the same level of educations ($\Delta = 0$), and households where wives are more educated than their husbands ($\Delta \geq 0$). Figure 8 shows event-study estimates from Equation (6) using this alternative measure of comparative advantage. The results remain consistent, and the heterogeneity in the effects across groups appears even more pronounced.

Another concern is that the observed patterns might reflect differences in educational attainment. For instance, women contributing a smaller share of the household's potential earnings may have, on average, lower human capital and weaker labor market attachment, making them more likely to leave the workforce after having children. To address this, I estimate child penalties using a variation of Equation (6) where, instead of allowing for heterogeneity by the degree of comparative advantage, I interact the main event-time dummies with an education category. I distinguish between women with a college degree or more and those with less than a college degree. As shown in Figure 9, the child penalties for these two groups are nearly identical, indicating that differences in educational attainment do not drive the observed patterns. If anything, women with lower educational attainment exhibit a marginally smaller child penalty.

While this finding may appear intuitive in hindsight, it constitutes a significant contribution to the literature. This result contrasts with (Kleven, Landais, Posch, Steinhauer and Zweimuller, 2019b), who finds no evidence that comparative advantage influences child penalties in Denmark. To the best of my knowledge, this effect has not been formally estimated in other contexts. Although prior research has explored heterogeneity in child penalties, the prevailing consensus suggests limited or no variation by comparative advantage. My results challenge this view, providing robust evidence that comparative advantage plays a non-negligible role in shaping child penalties. These findings offer important insights into how economic incentives interact with household decision-making and labor market outcomes.

Nonetheless, this analysis remains descriptive. To rigorously investigate this association, I extend the analysis by leveraging exogenous shocks that differentially affected the wages of women and men across labor markets. This approach enables me to isolate the causal impact of wage shifts on child penalties, yielding deeper insights into the ways in which labor market conditions influence family dynamics and gendered labor market outcomes.

6 Part II: Leveraging exogenous shocks to the female-to-male wage ratio

Did historical shifts in the average female-to-male wage ratio across US labor market have an effect on the evolution of the child penalty? Conceptually, I would like to estimate the causal effect of shifting the female-to-male wage ratio on the size of the child penalty in employment over the years 1980-2010 in the United States. My empirical specification is a panel fixed-effects regression of the following type:

$$Y_{cz,t} = \beta lnRelativePotentialWage_{cz,t} + \alpha_{cz} + \gamma_t + \phi X_{cz,t} + \nu_{cz,t} \tag{7}$$

where $Y_{cz,t}$ is a measure of child penalty in employment in a given commuting zone (CZ) at time t, where $t = 1980, 1990, 2000, 2010$. The main independent variable is a measure of gender wage gap in a local labor market. Note that I focus on potential rather than actual wages for two main reasons. First, previous research has shown that potential wages are the relevant metric for household production decisions (Pollak, 2005; Shenhav, 2021). Second, actual wages are likely subject to reverse causality; for example, in regions where women exhibit a lower preference for working, their observed wages may be lower due to reduced labor force attachment. Consider a CZ where women prefer to stay home or work part-time after having children—observed wages for women would consequently be lower. To address this issue, I use hourly wages instead of yearly wages and focus on wages for a sample of individuals working full-time and full-year. Empirically this is captured by $(lnW_f - lnW_m)$, where W_f and W_m are average potential hourly wages for women and men, respectively.

I include CZ fixed effects α_{cz} to control for constant differences across CZs which may explain initial differences in occupational choices, preferences for market vs non-market work, or family formation. Year fixed effects γ_t absorb national changes in labor market opportunities for women. In all my specifications I include a vector of controls for CZ demographics $X_{cz,t}$. I cluster standard errors at the commuting zone level and weight all the regressions by size of the working-age population.

However, this is not enough. The coefficient of interest β from equation 7 captures the causal effect of relative wage on the size of the child penalty in employment if, conditional on the control variables included, the variation in relative potential wages is plausibly exogenous (i.e., orthogonal to the error term). However, wages (both actual and potential) are not randomly assigned. For example, hourly potential wages for women and men in a given region may still reflect unobserved skills or characteristics inherent to that labor market, such as regional differences in educational attainment, industry composition, or cultural norms surrounding work. These unobserved factors can introduce biases if not properly accounted for. Hence, I proceed by constructing an instrument. In Section 6.1 I talk about the construction of an instrument for relative wages. In Section 6.2 I show the validity of my identification strategy, and in Section 6.3 I show the results.

6.1 Proxy for relative wage

I construct a proxy for relative potential wages using an application of the Bartik approach borrowed from previous work (Shenhav, 2021; Bertrand et al., 2015; Autor and Dorn, 2013) which leverages three main facts: 1) Local labor markets differ in their industry composition; 2) Women and men tend to work in different industries, and occupations within industries; and 3) The computerization wave of the 1980s reoriented demand towards occupations in which women were largely employed. As a result, female and male wages converged much faster in local labor markets that had higher presence of service sector jobs, compared to labor markets that were dominated by agriculture, construction or manufacturing.

Empirically, this is done by creating a proxy for potential wages for each CZ-by-year-by-gender that takes the following form (Bartik, 1991; Bertrand, Kamenica and Pan, 2015; Shenhav, 2021):

$$\widehat{w_{cz,g,t}} = SHARE_{1970} \times SHIFT_t \tag{8}$$

where the $SHARE_{1970}$ component jointly captures initial differences in industries and occupations between men and women and across local labor markets in the year 1970 (prior my period of analysis), and the $SHIFT_t$ term is composed of national occupation- and industry-specific average wages in each year t. Specifically, I break-down the "share" element into the following components:

$$SHARE_{1970} = \sum_{ind} \underbrace{\frac{E_{ind,cz,g,1970}}{E_{cz,g,1970}}}_{\text{Between-industry exposure}} \times \sum_{o} \underbrace{\frac{E_{o,ind,g,1970}}{E_{ind,g,1970}}}_{\text{Within-industry exposure}} \tag{9}$$

where the *"Between-industry"* component is the share of the working population of sex g employed in each industry within a CZ, and the *"Within-industry"* component captures the share of each occupation within each industry by sex. Note that the *"Within-industry"* component is obtained at the national level, rather than being CZ-specific. This is to minimize the amount of noise in the shares, since the occupation-by-industry cells are quite small in the 1970 census. Appendix B.2 and B.3 show industry and occupation categories, as well as their distribution in the 1970 decennial census. Finally, I expand the *"Within-industry"* component by adding the following occupational updating term:

$$\pi^{W*}_{o,ind,t,-cz} = (\pi^{W}_{o,ind,t,-cz}) \times (\frac{1}{\pi_{ot,-cz}}) \tag{10}$$

where $\pi^{W}_{o,ind,t,-cz}$ is the ratio of within-industry share of an occupation in t, compared to the same share in 1970; and $\pi_{ot,-cz}$ is the ratio of share of an occupation in t, compared to the same share in 1970. Conceptually, this updating term leverages the differential growth in the importance of occupations across industries and captures deviations in the growth of the within-industry employment share of an occupation, from the growth in the national employment share of the occupation. This allows to increase the predictive power of the wage proxy, without compromising its validity: indeed, this source of growth is likely to reflect industry productivity or changes in technology rather than being driven by changes in labor supply (Shenhav, 2021).

$$\widehat{w_{cz,g,t}} = \sum_{ind} \underbrace{\frac{E_{ind,cz,g,1970}}{E_{cz,g,1970}}}_{\text{Between-industry exposure}} \times \sum_{o} \underbrace{\frac{E_{o,ind,g,1970}}{E_{ind,g,1970}} \times \pi^{W*}_{o,ind,t,-cz}}_{\text{Within-industry exposure}} \times SHIFT_t \tag{11}$$

The final piece in this equation is the $SHIFT_t$ element which I obtain as the weighted average of national occupation-by-industry-specific wages for a sample of wage workers working full-year full-time, excluding the CZ for which I am constructing the wage proxy (to avoid a mechanical correlation between the proxy and the observed wage in a local labor market). Hence the proxy for potential wages becomes:

$$\widehat{w_{cz,g,t}} = \sum_{ind}\sum_{o} \underbrace{\frac{E_{ind,cz,g,1970}}{E_{cz,g,1970}}}_{\text{Between-industry exposure}} \times \underbrace{\frac{E_{o,ind,g,1970}}{E_{inf,g,1970}} \times \pi^{W*}_{o,ind,t,-cz}}_{\text{Within-industry exposure}} \times w_{o,ind,t,-cz} \tag{12}$$

To gain intuition for the approach, consider a simple case of two industries j_1 and j_2, and that wages vary from low to high over time only for j_1. In that case, the identification strategy would simplify to a difference-in-differences design, comparing the effect of the wage change across counties that have a greater presence of j_1 to those with a lesser presence of j_1, in the base year of reference.

6.2 Validity of the identification strategy

Before presenting my main regression results, I verify that the identification strategy is not undermined by any significant threats to its validity. The proxy for potential wages must satisfy the two key assumptions of an instrumental variable (IV) approach. First, the proxy must be correlated with the endogenous variable. Column 1 of Table 1 shows that the correlation is both economically and statistically significant. Additionally, while it is not a necessary condition, we fail to reject the null hypothesis that the coefficient is equal to 1, which provides reassurance about the quality of the instrument. In Columns 2 and 3, I run the same regression separately for women and men, and the coefficients are statistically significant in both cases. Figure

10 illustrates the correlation between the long-run change in actual and potential wages from 1980 to 2010. Figure 11 displays the same relationship, disaggregated by gender. In all cases, the correlation is reassuringly positive. Interestingly, the instrument appears to be a slightly better predictor of female wages than female wages. Following Shenhav (2021), I also perform a falsification test. While observed male and female wages may be correlated due to similar market conditions, the male (female) observed wage should not be correlated with the female (male) proxy if the proxies are driven by exogenous national variation. In other words, we don't want the current potential wage for women or men to be predictive of opposite-sex wages. In Column 4 and 5 of Table 1 I show that only the coefficient on the same-gender proxy is significant. This allows me to rule out the possibility that the potential wage is spuriously correlated with a shift in general labor market conditions, such as a local resource boom or greater enthusiasm for working.

6.3 Reduced-form Results

Next, I proceed to obtain my main results: Table 2 shows the estimates from reduced-form regressions. Note that I follow previous work in running reduced-form regressions of outcomes on a proxy for potential wages, rather than instrument observed wages with the proxy, which allows for potential wages to have an impact on marriage decisions through multiple channels, such as through higher bargaining power as well as through a higher realized wage (Bertrand, Kamenica and Pan, 2015; Shenhav, 2021; Aizer, 2010).

The coefficients have been rescaled to represent the effect of a 10% increase in relative potential wages, which is approximately the size of the convergence over this time period[9]. The coefficient of interest is negative and significant, and indicates that a 10% increase in the relative wage for

[9]The percent growth between 1980 and 2010 for potential $\log(W_f - W_m)$ is approximately: Change $=$ $\text{Value}_{2010} - \text{Value}_{1980} = -2.286251 - (-2.596048) = 0.309797$ from which we obtain: $\left(\frac{0.309797}{2.596048}\right) \times 100 \approx 11.93\%$. This approach quantifies the relative change in the log-transformed difference between female and male potential wages over the period.

women and men is associated with roughly a 15-16pp decrease in the size of the child penalty in employment. In 1980 the average child penalty in employment was 43% and it declined to about 23% in 2010, although the largest change happened between 1980 and 1990 (from 43% to 27%) while the decline was much smaller over the subsequent decades. My results seem to suggest that around three quarters of the decline in the child penalty over my period of analysis might be explained by improvement in economic opportunities for women, relative to men.

Column 2 of Table 2 shows the same result when controlling for the average education of the CZ, separately by gender, as well for the CZ racial composition. Column 3 adds the sex-ratio of the working-age population. In Panel B of Table 2 I examine the sensitivity of my results when controlling for average potential earnings in the market (measured as the average of female and male potential wages). This allows me to separate the effect of the relative wage from absolute wages. The estimated effects are only marginally reduced when I introduce this control variable. The insensitivity of the point estimate indicates that there is enough variation in the relative wage measure independent of the average wage measure.

To better understand this effect, I test whether higher relative wages have a stronger effect on the employment of women before or after having a child. Indeed, keeping fathers' employment constant, there are theoretically two possible ways in which the child penalty could increase: 1) if the post-birth employment of mothers goes down; or 2) if the pre-birth employment of mothers goes up (see Figure 12 for a toy illustration). To partially test for this, I split my estimation sample into pre-birth and post-birth observations. Table 3 shows regression results where the dependent variable is CZ-by-year employment rate for women with children (column 1) and women without children (column 2). The results suggest that, although a higher relative wage increases employment for all women, the effect is substantially larger for mothers. This seems to make sense based on well-documented facts: female labor supply underwent a substantial surge in the decades following World War II (and especially in the 1960s-1970s), and by the beginning of my period of analysis (1980) it had become quite common for childless women to seek employment. Hence, mothers have become the most marginal workers, and those more likely to make labor supply decisions based on intra-households earnings potential.

6.4 Mechanisms

How do relative wages affect the child penalty? One potential mechanism is through changes in family formation. Based on the descriptive evidence from Section 5, one could hypothesize that, as wages converge, spouses are more likely to match based on similar levels of human capital. To test this, I construct a measure of the average education gap between spouses (calculated as the wife's years of education minus the husband's) for each CZ-by-year. The results in Table 4 show that as wages converge, the education gap between spouses narrows. Specifically, a 10% increase in relative wages is associated with a 0.092 to 0.194-year increase in the average wife-husband education gap.

Another possibility is that, as women's economic stature improves, they may choose to delay childbirth. Indeed, delaying childbirth might imply that careers' interruptions happen at a less critical point in a woman's life, or even just at a point in which career investments are high enough that the incentives to get back to work are stronger. This, in turn, could lead to stronger attachment to the labor market. To explore this, I construct a CZ-by-year measure of age at first birth (conditional on being a mother) as the new dependent variable. Table 5 shows that an increase in the female-to-male wage ratio is associated with a postponement of childbirth. Specifically, a 10% increase in the female-to-male wage ratio corresponds to approximately 1.2 to 1.4 additional years at first birth. Given that the mean age at first birth in the U.S. was 22.7 in 1980 and increased to 25.4 by 2010 (according to CDC statistics), this suggests that about half of that increase could be attributed to the rise in women's relative economic stature.

Finally, a large body of research shows that child penalties are closely correlated with gender norms. One possibility, therefore, is that gender norms play a key role. In regions where employment opportunities have historically favored women, a gender ideal that reconciles motherhood and career may have developed.

To test this hypothesis, I pool wage data from the General Social Survey (GSS) and I use GSS survey questions to construct a measure of gender bias for each adult respondent (ages 18–64) following previous work (Kleven, 2024). Specifically, I take answers to the following questions:

- *It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family*

- *A working mother can establish just as warm and secure a relationship with her children as a mother who does not work*

- *A pre-school child is likely to suffer if his or her mother works*

I standardize each answer to be $N(0, 1)$ and sum them up to obtain an index of gender progressivity bias. Figure 13 shows the evolution of this constructed measure of gender bias over time. Interestingly, norms seem to evolve similarly to child penalties over time, and mirroring relative wages. Unfortunately, the GSS does not have county identifiers before 1993 (when most of the action takes place). Hence, I construct a measure of gender bias at the state-by-year level, and then run a regression with state and year fixed effects. Column 1 of Table 6 shows regression results, suggesting that there is some evidence of a decrease in elicited gender bias, when relative wages converge. In Columns 2 and 3 of Table 6 I first run a state-level regression only focusing on post-1993 observations (for which I can identify CZs), and then a CZ-level regression. Results go away when we lose the first years of data, suggesting that it is hard to make inferences at the CZ level. However, although suggestive, these results suggest that higher potential relative wages are associated with a reduction in gender bias, supporting the idea that economic opportunities for women can shape social norms around gender roles and work.

7 Conclusions

Child penalties account for most of the remaining gender inequality in the labor market and their study has recently attracted the attention of researchers across different disciplines. Despite the great advances that have been produced, the existing literature still suffers from two main shortcomings. The first is the lack of large scale longitudinal data which is traditionally required to study child penalties within an event-study framework. The second is the limited understanding of their fundamental determinants: child penalties are large and persistent, and we do not know much about what causes them. Another limitation of existing work is that while child penalties are known to be correlated with gender norms, these norms often remain a "black box". In particular, gender norms are frequently considered static and slow-moving, with little exploration into their evolution or responsiveness to economic and social factors.

To address the first gap, I propose a novel method for measuring child penalties using sparse cross-sectional data, building on pre-existing approaches that leverage multiple cross-sections. I validate this method against real panel data and apply it to generate the first sub-state-level (CZ-by-year) estimates of child penalties in the U.S. In response to the second gap, I provide evidence that comparative advantage is a potential determinant of child penalties. My analysis reveals a novel finding, which is the presence of significant heterogeneity in the size of child penalties based on intra-household earnings potential. Furthermore, I find that child penalties in employment have declined more significantly in local labor markets where wages for women and men converged most between 1980 and 2010. These results align with Becker's model of household specialization, suggesting that as economic opportunities between genders equalize, the division of labor between market and home production becomes less gendered.

Given that child penalties capture differences in parental labor supply after the birth of the first child relative to pre-birth levels, I further decompose the effects on women's labor supply into changes pre- and post-birth. My findings indicate that higher relative wages for women in a local labor market are associated with increased labor supply for mothers post-birth, rather than pre-birth. I also identify mechanisms driving this relationship, including a relative increase in women's educational attainment (compared to men), and a rise in age-at-first-birth. I also find suggestive evidence that economic opportunities can shape gender norms, indicating a potentially dynamic relationship where changes in economic conditions influence societal norms around gender roles and labor market participation.

This work is not without limitations; I am currently testing the robustness of my findings using alternative measures of earnings potential and variations in relative wages. Future research should further explore the link between child penalties and gendered economic opportunities, potentially aided by the availability of improved data.

References

Adda, Jérôme, Christian Dustmann, and Katrien Stevens, "The career costs of children," *Journal of Political Economy*, 2017, *125* (2), 293–337.

Aizer, Anna, "The gender wage gap and domestic violence," *American Economic Review*, 2010, *100* (4), 1847–1859.

Albanesi, Stefania, Claudia Olivetti, and Barbara Petrongolo, "Families, Labor Markets, and Policy," Technical Report, National Bureau of Economic Research 2022.

Anelli, Massimo, Osea Giuntella, and Luca Stella, "Robots, Marriageable Men, Family, and Fertility," *Journal of Human Resources*, 2021, pp. 1020–11223R1.

Autor, David, David Dorn, and Gordon Hanson, "When work disappears: Manufacturing decline and the falling marriage market value of young men," *American Economic Review: Insights*, 2019, *1* (2), 161–78.

Autor, David H and David Dorn, "The growth of low-skill service jobs and the polarization of the US labor market," *American economic review*, 2013, *103* (5), 1553–1597.

_ , Lawrence F Katz, and Melissa S Kearney, "Trends in US wage inequality: Revising the revisionists," *The Review of economics and statistics*, 2008, *90* (2), 300–323.

Bartik, Timothy, "Who benefits from state and local economic development policies?," 1991.

Bertrand, Marianne, Emir Kamenica, and Jessica Pan, "Gender identity and relative income within households," *The Quarterly Journal of Economics*, 2015, *130* (2), 571–614.

— , **Patricia Cortes, Claudia Olivetti, and Jessica Pan**, "Social norms, labour market opportunities, and the marriage gap between skilled and unskilled women," *The Review of Economic Studies*, 2021, *88* (4), 1936–1978.

Blau, Francine D and Lawrence M Kahn, "International differences in male wage inequality: institutions versus market forces," *Journal of Political Economy*, 1996, *104* (4), 791–837.

— **and** — , "The gender wage gap: Extent, trends, and explanations," *Journal of economic literature*, 2017, *55* (3), 789–865.

Boelmann, Barbara, Anna Raute, and Uta Schonberg, "Wind of change? Cultural determinants of maternal labor supply," 2021.

Cortés, Patricia and Jessica Pan, "Children and the remaining gender gaps in the labor market," Technical Report 2023.

Fortin, Nicole M, "Gender role attitudes and the labour-market outcomes of women across OECD countries," *oxford review of Economic Policy*, 2005, *21* (3), 416–438.

Goldin, Claudia, Sari Pekkala Kerr, and Claudia Olivetti, "When the Kids Grow Up: Women's Employment and Earnings across the Family Cycle," Technical Report, National Bureau of Economic Research 2022.

Jones, Larry E, Rodolfo E Manuelli, and Ellen R McGrattan, "Why are married women working so much?," *Journal of Demographic Economics*, 2015, *81* (1), 75–114.

Juhn, Chinhui and Kristin McCue, "Specialization then and now: Marriage, children, and the gender earnings gap across cohorts," *Journal of Economic Perspectives*, 2017, *31* (1), 183–204.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard, "Children and gender inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 2019, *11* (4), 181–209.

_ , _ , and _ , "Does biology drive child penalties? evidence from biological and adoptive families," *American Economic Review: Insights*, 2021, *3* (2), 183–98.

_ , _ , Johanna Posch, Andreas Steinhauer, and Josef Zweimuller, "Child penalties across countries: Evidence and explanations," *AEA Papers and Proceedings*, 2019, *109*, 122–26.

_ , Giulia Olivero, and Eleonora Patacchini, "Child Penalties and Parental Role Models: Classroom Exposure Effects," Technical Report, National Bureau of Economic Research 2024.

Kleven, Henrik Jacobsen, *The geography of child penalties and gender norms: A pseudo-event study approach*, National Bureau of Economic Research, 2024.

Kuka, Elira and Na'ama Shenhav, "Long-run effects of incentivizing work after childbirth," Technical Report, National Bureau of Economic Research 2023.

Pollak, Robert A, "Bargaining power in marriage: Earnings, wage rates and household production," 2005.

Ruggles, Steven, J Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B Schroeder, Matthew Sobek et al., "Integrated public use microdata series: Version 5.0 [Machine-readable database]," *Minneapolis: University of Minnesota*, 2010, *42.*

Shenhav, Na'ama, "Lowering standards to wed? Spouse quality, marriage, and labor market responses to the gender wage gap," *Review of Economics and Statistics*, 2021, *103* (2), 265–279.

Tolbert, Charles M and Molly Sizer, "US commuting zones and labor market areas: A 1990 update," 1996.

Figures

Figure 1. Evolution of relative wage and child penalty in employment

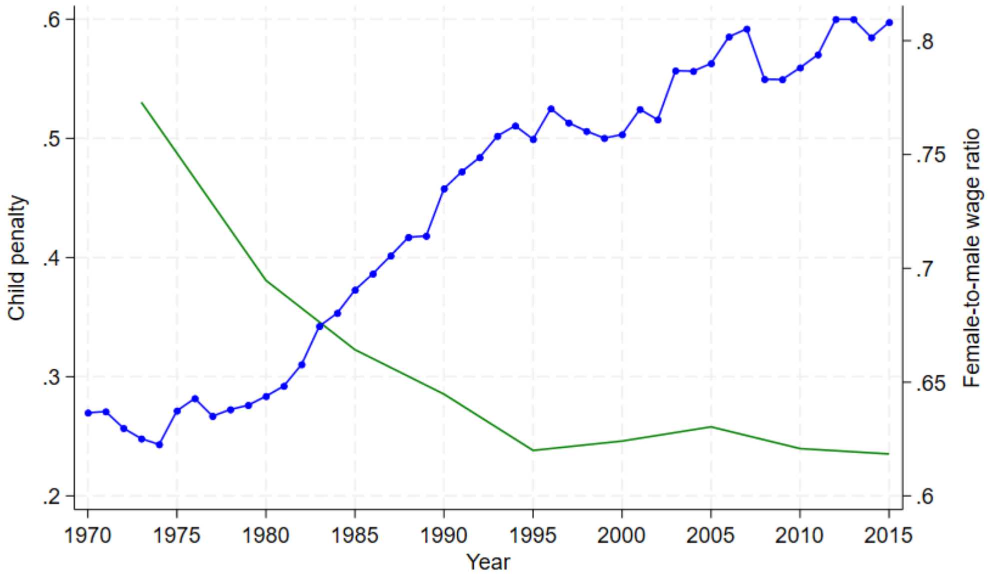

Green solid line*: represents the evolution of child penalty in employment computed by Kleven et al. (2022). Penalties are computed with an event study comparing differences in labor market outcomes for mothers and fathers over 10 years after giving birth to their first child, as compared to their pre-child levels.*
Blue dotted line*: the dotted line represents the average female-to-male hourly wage. Computed from CPS data for the years 1970-2015 using a sample of individuals aged 18-64 working full-year full-time.*

Figure 2. Changes in female labor supply for different levels of τ_f, τ_m, γ

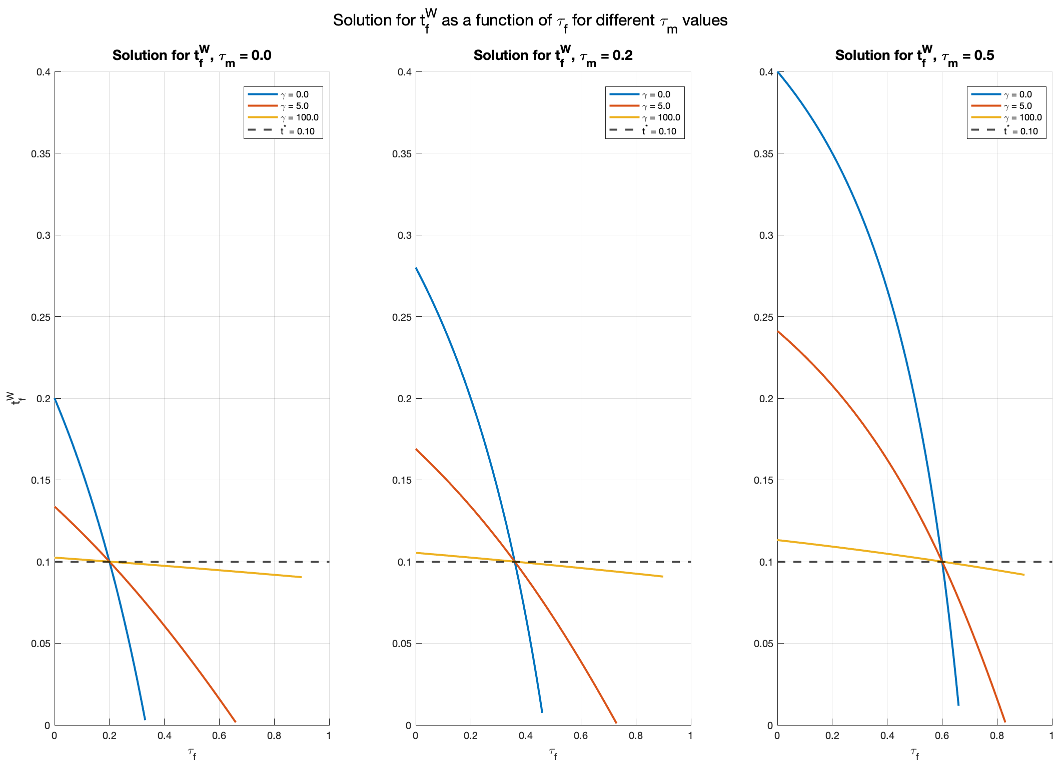

Notes: Changes in female labor supply (t_f^W, on the y-axis), for different levels of female tax rate (τ_f, on the x-axis), male tax rate ($\tau_m = 0$ left panel, $\tau_m = 0.2$ center panel, $\tau_m = 0.5$ right panel) and penalty associated with deviations from gender norms (γ=0, 5, 100). The "norm" on female labor supply is set at $t^* = 0.1$.

Figure 3. Child penalty in employment - NLSY79 real panel

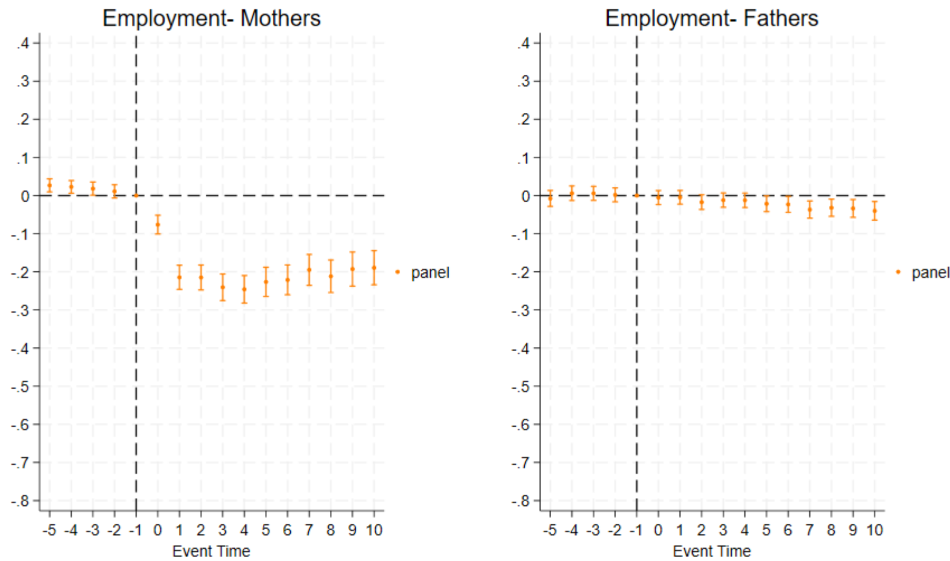

Event study on the child penalty in employment for a real sample of mothers and fathers observed between age 18-55 who had their first child between the age of 22 and 45. Sample obtained from pooling together waves from the NLSY79. Estimation includes age and years fixed-effects. Event time=-1 indicates the year before the birth of the first child. Robust standard errors.

Figure 4. Child penalty in employment - NLSY79 real vs pseudo panel

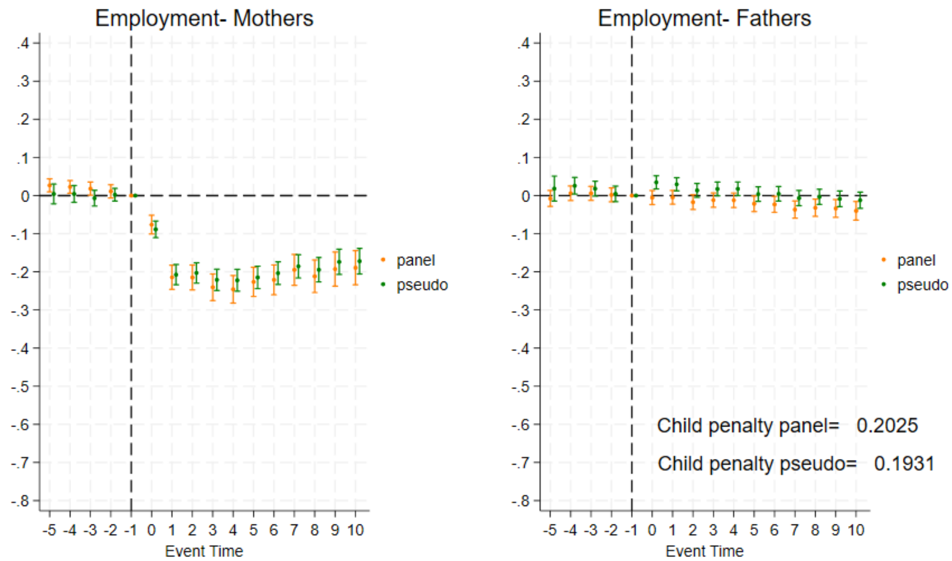

Orange solid line:*Event study on the child penalty in employment for a real sample of mothers and fathers who had their first child between the age of 22 and 45.*
Green solid line:*Event study on the child penalty in employment for a pseudo-panel of mothers and fathers who had their first child between the age of 22 and 45. Sample obtained from pooling together waves from the NLSY79. Estimation includes age and years fixed-effects. Event time=-1 indicates the year before the birth of the first child. Robust standard errors.*

Figure 5. Distribution of household by comparative advantage of the wife and Census year

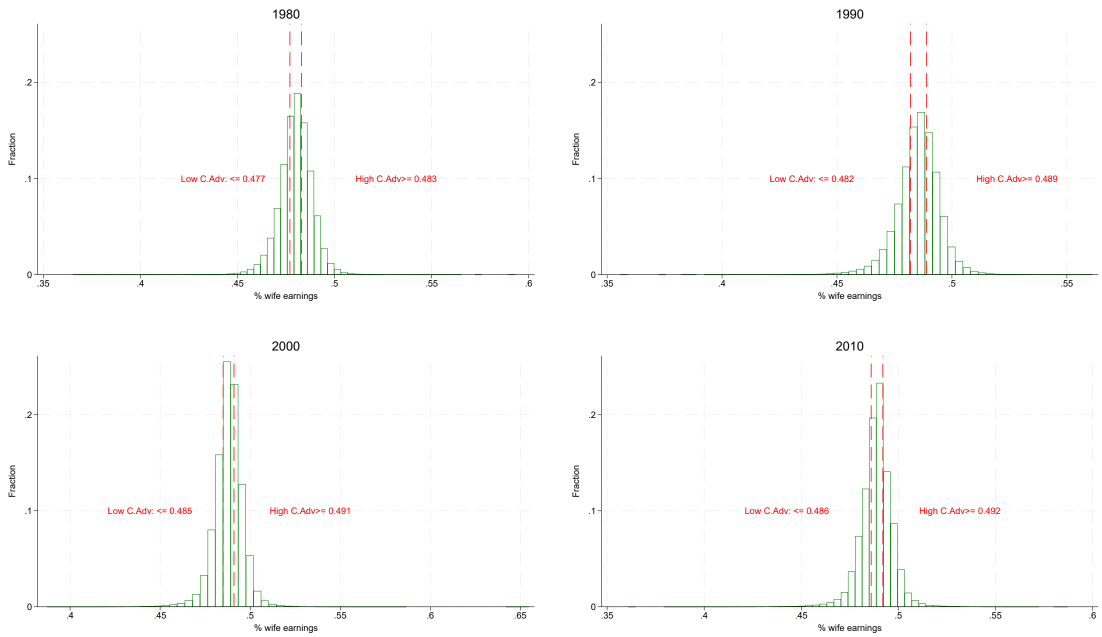

Notes: Distribution of households by share of potential earnings earned by the wife. Each observation represents one married Census household. Vertical lines identify the tertiles of the distribution (left: low comparative advantage, center: medium comparative advantage, right: high comparative advantage). Data: Decennial Census 1980-2000, ACS 5-years 2010.

Figure 6. Heterogeneous child penalty estimates by tertile of comparative advantage - potential earnings

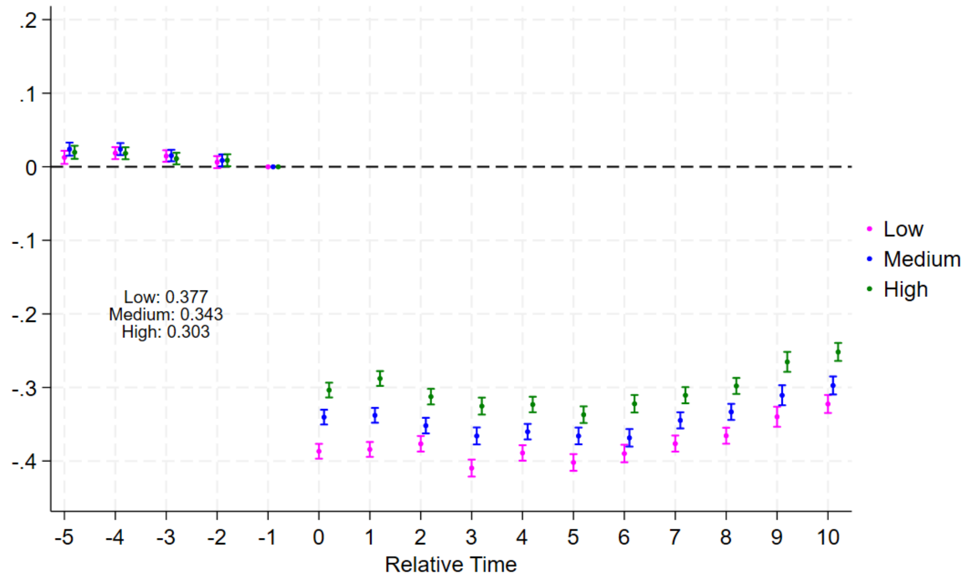

Notes: Event-study estimates, separately by tertile of comparative advantage, where comparative advantage is estimated based on intra-household earnings potential. Coefficients obtained from estimating Equation (6). Data: Decennial Census 1980-2000, ACS 5-years 2010. Errors clustered at the CZ level.

Figure 7. Distribution of household by education gap between the two spouses and Census year

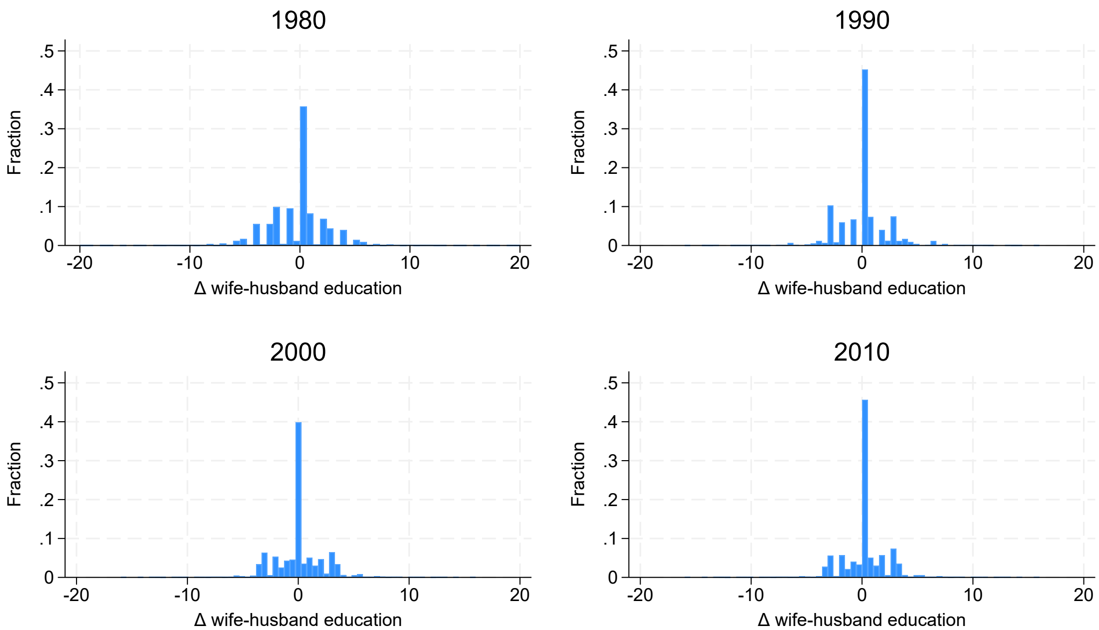

Notes: Distribution of households by wife-husband education gap. Each observation represents one married Census household. (left: low comparative advantage, center: medium comparative advantage, right: high comparative advantage). Data: Decennial Census 1980-2000, ACS 5-years 2010.

Figure 8. Heterogeneous child penalty estimates by tertile of comparative advantage - education

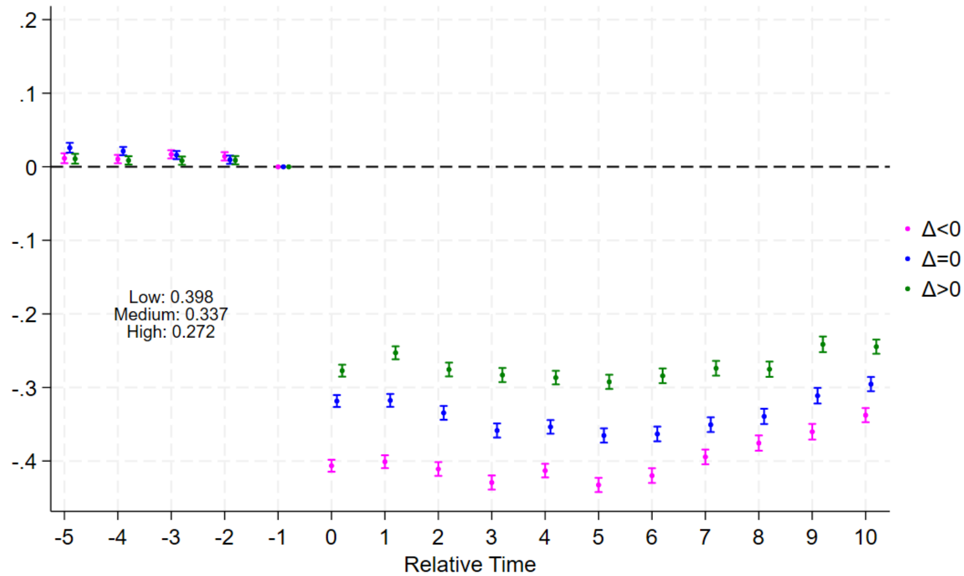

Notes: Event-study estimates, separately by educational level (less than a college degree, vs college or more). Coefficients obtained from estimating Equation (6). Data: Decennial Census 1980-2000, ACS 5-years 2010. Errors clustered at the CZ level.

Figure 9. Heterogeneous child penalty estimates by education

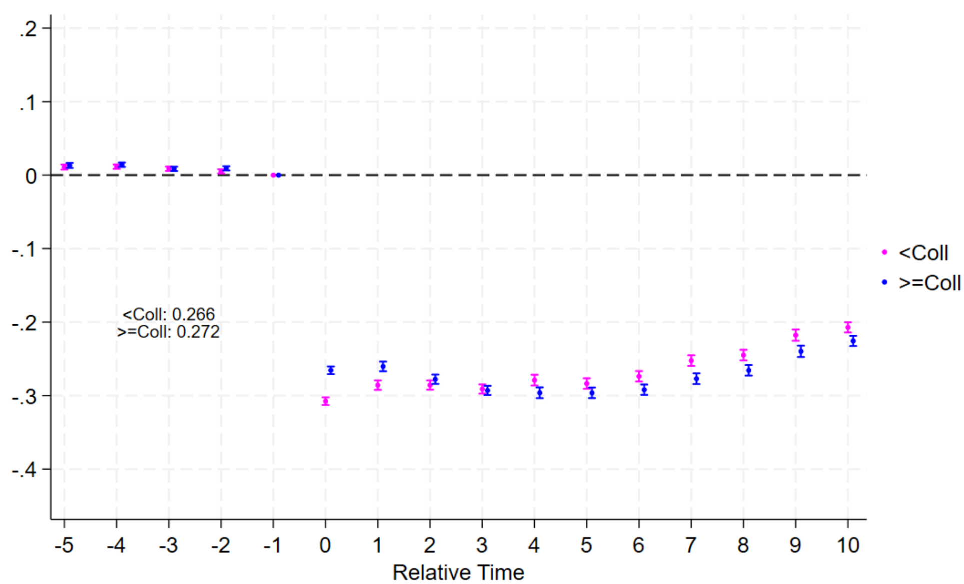

Notes: Event-study estimates, separately by tertile of comparative advantage and Census year.

Figure 10. Change in observed and potential log relative wage 1980-2010

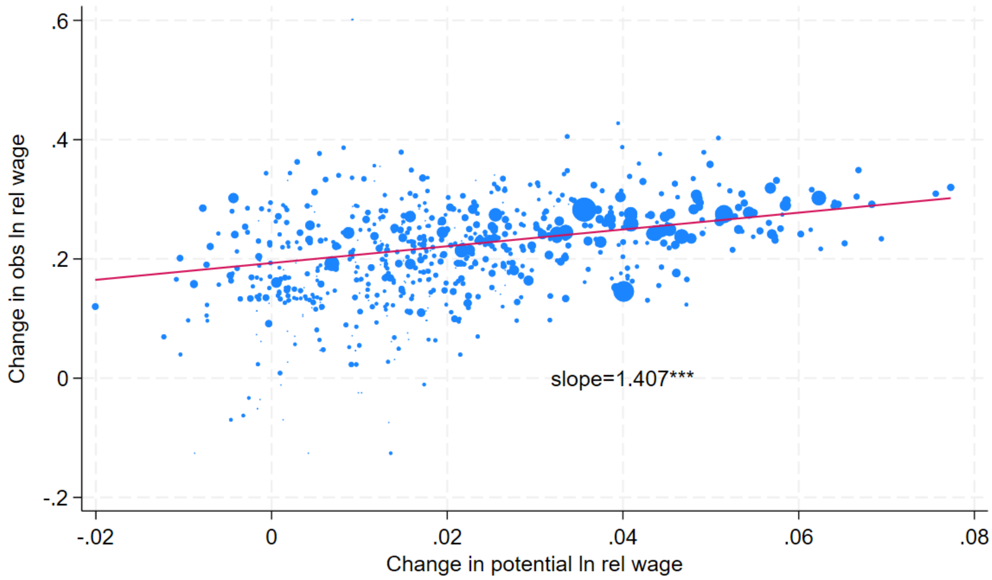

This figure represents the correlation in the long change of potential log relative wage, and actual log relative wage (1980-2010). Each dot represents a CZ, the size of each dot is proportional to the size of the working age population in the CZ (and weights are applied accordingly). The slope is positive and significant.

Figure 11. Change in observed and potential log female and log male wage 1980-2010

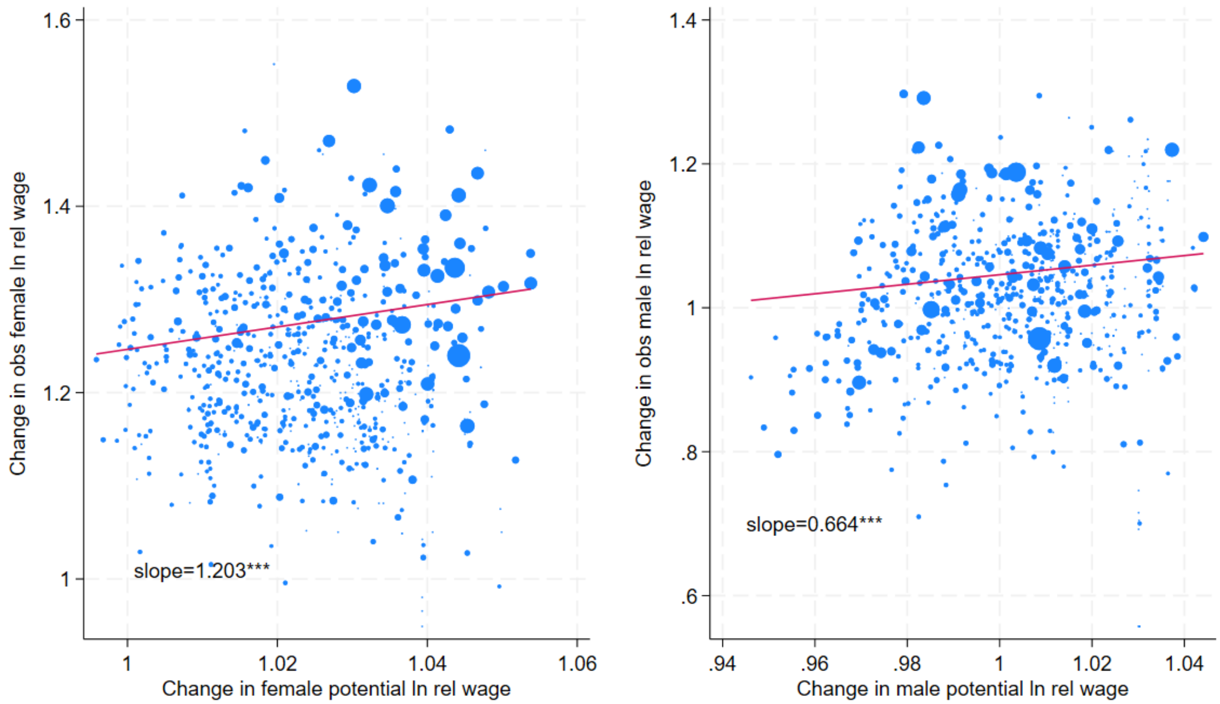

This figure represents the correlation in the long change of potential log relative wage, and actual log relative wage (1980-2010) separately for women (left) and men (right). Each dot represents a CZ, the size of each dot is proportional to the size of the working age population in the CZ of that gender (and weights are applied accordingly). The slope is positive although it is only significant for men.

Figure 12. Change in pre- versus post- birth employment

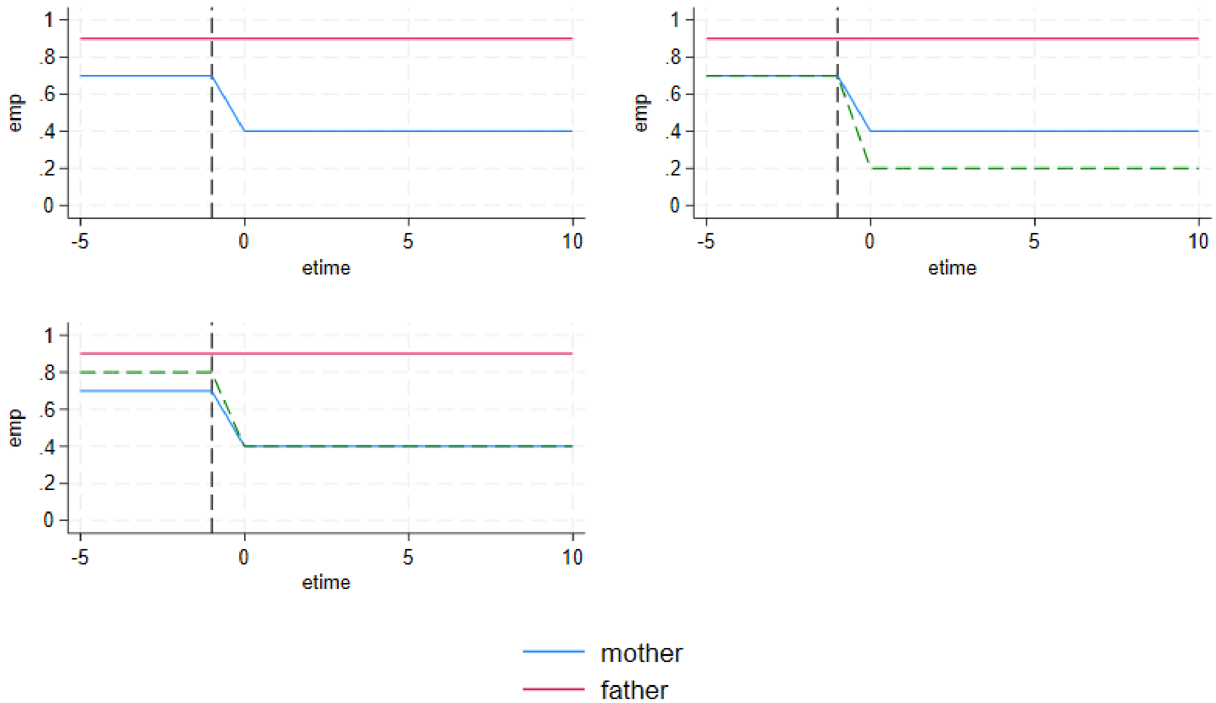

Toy example showing how a child penalty could mechanically increase. The red line represents employment trajectory for fathers. The blue line represents employment trajectory for mothers. The green dotted line shows that the child penalty can increase if post-birth employment of mothers decreases (figure in top right corner) or if the pre-birth employment of mothers increases (figure in bottom-left corner).

Figure 13. Evolution of relative wage, child penalty in employment, and elicited gender norms

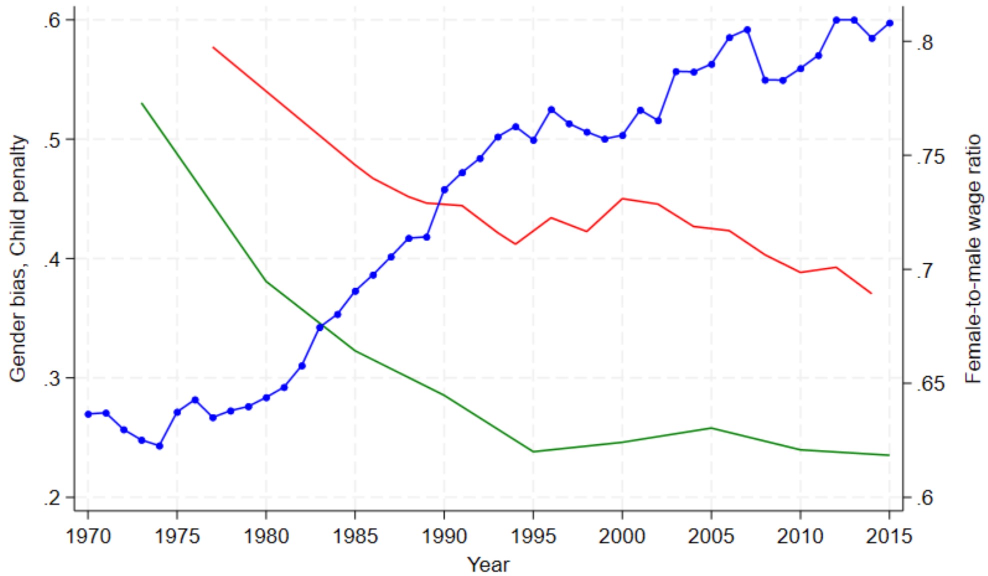

Green solid line*: represents the evolution of child penalty in employment computed by Kleven et al. (2024). Penalties are computed with an event study comparing differences in labor market outcomes for mothers and fathers over 10 years after giving birth to their first child, as compared to their pre-child levels.*

Blue dotted line*: the dotted line represents the average female-to-male hourly wage. Computed from CPS data for the years 1970-2015 using a sample of individuals aged 18-64 working full-year full-time.*

Tables

Table 1. First-stage results

	Corr. w/ Actual			Cross-Effects?	
	Relative	Female	Male	Female	Male
ln potential rel. wage	1.193***				
	(0.144)				
ln potential female wage		1.250***		1.483***	0.699
		(0.363)		(0.386)	(0.452)
ln potential male wage			0.905**	-0.560	0.781**
			(0.362)	(0.345)	(0.396)
Partial R-Squared	0.113	0.020	0.023		
Obs	2964	2964	2964	2964	2964

First-stage regression where dependent variable is log of actual wage (relative, female, and male respectively) and the main independent variable is a proxy for potential wage. Each observation is a CZ-by-year. Controls include CZ and year fixed effects, CZ-by-year demographic controls such as average years of education by gender, share of black and hispanic population and sexratio.

Table 2. Main results - child penalty in weekly employment

	Child penalty	Child penalty	Child penalty
Panel A: Relative only			
ln potential rel. wage	-0.166***	-0.162***	-0.161**
	(0.060)	(0.062)	(0.063)
CZ race composition	No	Yes	Yes
CZ education by sex	No	Yes	Yes
CZ sexratio	No	No	Yes
CZ fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Obs	2964	2964	2964
Panel B: Relative controlling for average			
ln potential rel. wage	-0.151***	-0.158***	-0.158***
	(0.055)	(0.053)	(0.053)
CZ race composition	No	Yes	Yes
CZ education by sex	No	Yes	Yes
CZ sexratio	No	No	Yes
CZ fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Obs	2964	2964	2964

Reduced-form regression where dependent variable is child penalty in weekly employment and the main independent variable is a proxy for log relative potential wage. Each observation is a CZ-by-year. Controls include CZ and year fixed effects, CZ-by-year demographic controls such as average years of education, share of black and hispanic population and sexratio.

Table 3. Mechanisms - Pre- vs post-birth employeemnt

	Pre-birth employment	Post-birth employment
Panel A: Relative only		
ln potential rel wage	-0.018	0.070***
	(0.013)	(0.017)
CZ race composition	Yes	Yes
CZ education by sex	Yes	Yes
CZ sexratio	Yes	Yes
CZ fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Obs	2964	2964
Panel B: Relative controlling for average		
ln potential rel wage	0.024*	0.079***
	(0.013)	(0.020)
CZ race composition	Yes	Yes
CZ education by sex	Yes	Yes
CZ sexratio	Yes	Yes
CZ fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Obs	2964	2964

Reduced-form regression where dependent variable is child penalty in weekly employment and the main independent variable is a proxy for log relative potential wage. Each obs. is a CZ-by-year. Controls include CZ and year fixed effects, CZ-by-year demographic controls such as average years of education, share of black and hispanic population and sexratio.

Table 4. Mechanisms - Spouses educational gap

	Δ spouse education	Δ spouse education
Panel A: Relative only		
ln potential rel wage	0.092*	
	(0.051)	
Panel B: Relative controlling for average		
ln potential rel wage		0.194***
		(0.062)
CZ race composition	Yes	Yes
CZ education by sex	Yes	Yes
CZ sexratio	Yes	Yes
CZ fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Obs	2964	2964

Reduced-form regression where dependent variable is employment for mothers (column 1) and employment for non-mothers (column 2) and the main independent variable is a proxy for log relative potential wage. Each observation is a CZ-by-year. Controls include CZ and year fixed effects, CZ-by-year demographic controls such as average years of education, share of black and hispanic population and sexratio.

Table 5. Mechanisms - Age at birth

	Age at birth	Age at birth
Panel A: Relative only		
ln potential rel wage	1.249***	
	(0.179)	
Panel B: Relative controlling for average		
ln potential rel wage		1.390***
		(0.172)
CZ race composition	Yes	Yes
CZ education by sex	Yes	Yes
CZ sexratio	Yes	Yes
CZ fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Obs	2964	2964

Reduced-form regression where dependent variable is employment for mothers (column 1) and employment for non-mothers (column 2) and the main independent variable is a proxy for log relative potential wage. Each observation is a CZ-by-year. Controls include CZ and year fixed effects, CZ-by-year demographic controls such as average years of education, share of black and hispanic population and sexratio

Table 6. Mechanisms - Elicited Gender Norms

	Gender Bias State 1980-2010	**Gender Bias State 1990-2010**	**Gender Bias CZ 1990-2010**
Panel A: Relative only			
ln potential rel wage	-0.097*	-0.032	-0.057
	(0.050)	(0.047)	(0.064)
State/CZ race composition	Yes	Yes	Yes
State/CZ education by sex	Yes	Yes	Yes
State/CZ sex ratio	Yes	Yes	Yes
State/CZ fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Obs	154	118	357
Panel B: Relative controlling for average			
ln potential rel wage	-0.108**	-0.051	-0.023
	(0.049)	(0.057)	(0.074)
State/CZ race composition	Yes	Yes	Yes
State/CZ education by sex	Yes	Yes	Yes
State/CZ sex ratio	Yes	Yes	Yes
State/CZ fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Obs	154	118	357

Reduced-form regression where dependent variable is a measure of gender bias at the State/CZ-by-year level obtained from GSS data. the main independent variable is a proxy for log relative potential wage. Each obs. is a CZ-by-year. Controls include CZ and year fixed effects, CZ-by-year demographic controls such as average years of education, share of black and hispanic population and sexratio.

Appendix

A Theoretical framework

Since the optimal solution to the utility maximization problem is expected to exhaust available time and resources (in other words, constraints are binding), I replace the inequality constraints with equality constraints, ensuring that both time and budget are fully allocated at the optimum. Hence, the problem can be re-written as:

$$U = \log(c_f^W) + \log(t_f^L) + \log(c_m^W) + \log(t_m^L) + \log(c^H) - \gamma(t_f^W - t^*)^2$$

Subject to the following constraints:

$$c_f^W + c_m^W = (1-\tau_f)wt_f^W + (1-\tau_m)wt_m^W \tag{i}$$

$$c^H = At_f^H \tag{ii}$$

$$t_f^H + t_f^W + t_f^L = 1 \tag{iii}$$

$$t_m^W + t_m^L = 1 \tag{iv}$$

At this point, I can plug constraints (ii) and (iv) into the objective function. Also note that, since $\log(At_f^H)$ can be rewritten as $\log(A)+\log(t_f^H)$, the term $\log(A)$ can be ignored since it is a constant and will not affect the solution to the optimization. Therefore, I rewrite the problem in the following way:

$$U = \log(c_f^W) + \log(t_f^L) + \log(c_m^W) + \log(1 - t_m^W) + \log(t_f^H) - \gamma(t_f^W - t^*)^2$$

Subject to the following constraints:

$$c_f^W + c_m^W = (1-\tau_f)wt_f^W + (1-\tau_m)wt_m^W \tag{i}$$

$$t_f^H + t_f^W + t_f^L = 1 \tag{iii}$$

Now, in order to further simplify the problem, I focus on a smaller subset of this optimization. I fix t_f^W, and solve the following problem:

$$\max_{t_f^L, t_f^H} \quad \log(t_f^L) + \log(t_f^H)$$

Subject to:

$$t_f^L + t_f^H = 1 - t_f^W$$

Once more, I plug the constraint into the objective function and derive a unique first order condition:

$$\max_{t_f^L} \quad \log(t_f^L) + \log(1 - t_f^L - t_f^W)$$

$$t_f^L : \frac{1}{t_f^L} - \frac{1}{1 - t_f^L - t_f^W} = 0 \tag{FOC}$$

From this FOC I obtain the following optimum solutions:

$$t_f^L = \frac{1 - t_f^W}{2}$$
$$t_f^H = \frac{1 - t_f^W}{2}$$

This means that, for each value of t_f^W, the wife will split the remaining available time equally between house-work and leisure. Moreover, the maximum utility that the household obtains is:

$$\begin{aligned} &\log(t_f^L) + \log(t_f^H) \\ = \log(&\frac{1 - t_f^W}{2}) + \log(\frac{1 - t_f^W}{2}) \\ &= 2\log(\frac{1 - t_f^W}{2}) \end{aligned}$$

Hence, I can simplify our original problem by rewriting $\log(t_f^L) + \log(t_f^H)$ in this way:

$$U = \log(c_f^W) + \log(t_f^L) + \log(c_m^W) + \log(1 - t_m^W) + \log(t_f^H) - \gamma(t_f^W - t^*)^2$$

$$U = \log(c_f^W) + 2\log(\frac{1 - t_f^W}{2}) + \log(c_m^W) + \log(1 - t_m^W) + -\gamma(t_f^W - t^*)^2$$

subject to:

$$c_f^W + c_m^W = (1 - \tau_f)wt_f^W + (1 - \tau_m)wt_m^W \tag{i}$$

Now, I apply the same trick to the consumption variable. Hence, I fix t_f^W and t_m^W, and I focus on the following problem:

$$\max_{c_f^W, c_m^W} \quad \log(c_f^W) + \log(c_m^W)$$

Subject to:

$$c_f^W + c_m^W = (1-\tau_f)wt_f^W + (1-\tau_m)wt_m^W \tag{i}$$

Notice that c_f^W and c_m^W enter the objective function with the same functional forms and weights, hence the solution will take the following form:

$$c_f^W = c_f^M = \frac{(1-\tau_f)wt_f^W + (1-\tau_m)wt_m^W}{2}$$

This expression shows that the level of consumption depends, among others, on the wage rate w. Next, I rewrite the problem in the following way:

$$\max_{t_f^W, t_m^W} \quad 2\log(\frac{(1-\tau_f)wt_f^W + (1-\tau_m)wt_m^W}{2}) + log(1-t_m^W) + 2\log(\frac{1-t_f^W}{2}) - \gamma(t_f^W - t^*)^2$$

At this point I can take the derivative of the utility function with respect to τ_f and τ_m and obtain the following two first order conditions:

$$t_f^W: \quad \frac{2(1-\tau_f)w}{(1-\tau_f)w \cdot t_f + (1-\tau_m)w \cdot t_m} - \frac{2}{1-t_f} - 2\gamma(t_f - t^*) = 0 \tag{FOC1}$$

$$t_m^W: \quad \frac{2(1-\tau_m)w}{(1-\tau_f)w \cdot t_f + (1-\tau_m)w \cdot t_m} - \frac{1}{1-t_m} = 0 \tag{FOC2}$$

Note that the wage rate w does not affect the optimality conditions, since it cancels out in the numerator and denominator. At this point, I can apply the Implicit Function Theorem, which takes the following form:

$$\begin{pmatrix} \frac{dt_f^W}{d\tau_f} & \frac{dt_f^W}{d\tau_m} \\ \frac{dt_m^W}{d\tau_f} & \frac{dt_m^W}{d\tau_m} \end{pmatrix} = -\underbrace{\begin{pmatrix} \frac{\partial \text{FOC1}}{\partial t_f^W} & \frac{\partial \text{FOC1}}{\partial t_m^W} \\ \frac{\partial \text{FOC2}}{\partial t_f^W} & \frac{\partial \text{FOC2}}{\partial t_m^W} \end{pmatrix}^{-1}}_{M_1} \underbrace{\begin{pmatrix} \frac{\partial \text{FOC1}}{\partial \tau_f} & \frac{\partial \text{FOC1}}{\partial \tau_m} \\ \frac{\partial \text{FOC2}}{\partial \tau_f} & \frac{\partial \text{FOC2}}{\partial \tau_m} \end{pmatrix}}_{M_2}$$

Where:

$$M_1 = \begin{pmatrix} -2\gamma - \frac{w^2(1-\tau_f)(2-2\tau_f)}{(t_f w(1-\tau_f)+t_m w(1-\tau_m))^2} - \frac{2}{(1-t_f)^2} & -\frac{w^2(1-\tau_m)(2-2\tau_f)}{(t_f w(1-\tau_f)+t_m w(1-\tau_m))^2} \\ -\frac{w^2(1-\tau_f)(2-2\tau_m)}{(t_f w(1-\tau_f)+t_m w(1-\tau_m))^2} & -\frac{w^2(1-\tau_m)(2-2\tau_m)}{(t_f w(1-\tau_f)+t_m w(1-\tau_m))^2} - \frac{1}{(1-t_m)^2} \end{pmatrix}$$

$$M_2 = \begin{pmatrix} \frac{t_f w^2(2-2\tau_f)}{(t_f w(1-\tau_f)+t_m w(1-\tau_m))^2} - \frac{2w}{t_f w(1-\tau_f)+t_m w(1-\tau_m)} & \frac{t_m w^2(2-2\tau_f)}{(t_f w(1-\tau_f)+t_m w(1-\tau_m))^2} \\ \frac{t_f w^2(2-2\tau_m)}{(t_f w(1-\tau_f)+t_m w(1-\tau_m))^2} & \frac{t_m w^2(2-2\tau_m)}{(t_f w(1-\tau_f)+t_m w(1-\tau_m))^2} - \frac{2w}{t_f w(1-\tau_f)+t_m w(1-\tau_m)} \end{pmatrix}$$

From here, we can obtain expressions for the closed form solutions. Since solutions are quite complex (mostly due to the non-linearity in the penalty term), I evaluate them numerically for a range of possible values and obtain the following:

- $\frac{dt_f^W}{d\tau_f} \leq 0$
- $\frac{dt_f^W}{d\tau_m} \geq 0$
- $\frac{dt_m^W}{d\tau_f} \leq 0$
- $\frac{dt_m^W}{d\tau_m} \geq 0$

Finally, from the FOCs, it can easily be derived that:

- $\frac{dt_f^W}{d\gamma} \leq 0$

- $\frac{dt_m^W}{d\gamma} = 0$

B Data

B.1 Wages

To construct a measure of actual and potential wages in a CZ-by-year I closely follow Autor et al. (2008). I only keep individuals aged 18-64 who are employed and work for a wage. I exclude individuals living in group-quarters or who report being in school. Furthermore, I only keep individuals who work full-year full-time (at least 40 weeks a year and at least 35 hours per week). I multiply top-coded values by 1.5 and I drop observations with wage lower than half the minimum wage in that state and year. I compute *ln_hourly_wage* for each individual in my sample. Finally, to create an average for each CZ-by-year, I collapse individual wages using Census statistical weights (*perwt*) multiplied by the number of hours worked and the allocation factor indicating the probability that each individual in a given *county-group/puma* belongs to a CZ.

B.2 Industry shares

Industry shares are computed from 1% 1970 Decennial Census. I create 15 broad industry categories (see Table B.2). Figure B.1 shows the national distribution of industry by gender in 1970. Men are disproportionately employed in the manufacturing sector (29.8%), while women are over-represented in the professional services sector (29.6%).

Table B.1. Industry classification

1	Agriculture, forestry and fisheries
2	Mining
3	Construction
4	Manufacturing
5	Transportation
6	Communication
7	Utilities
8	Wholesale trade
9	Retail trade
10	Finance
11	Business and repair services
12	Personal services
13	Entertainment and recreation
14	Professional Services
15	Public Administration

Source: 1% Census Microdata 1970

Figure B.1. Proportion employed in each industry by gender - 1970 Census

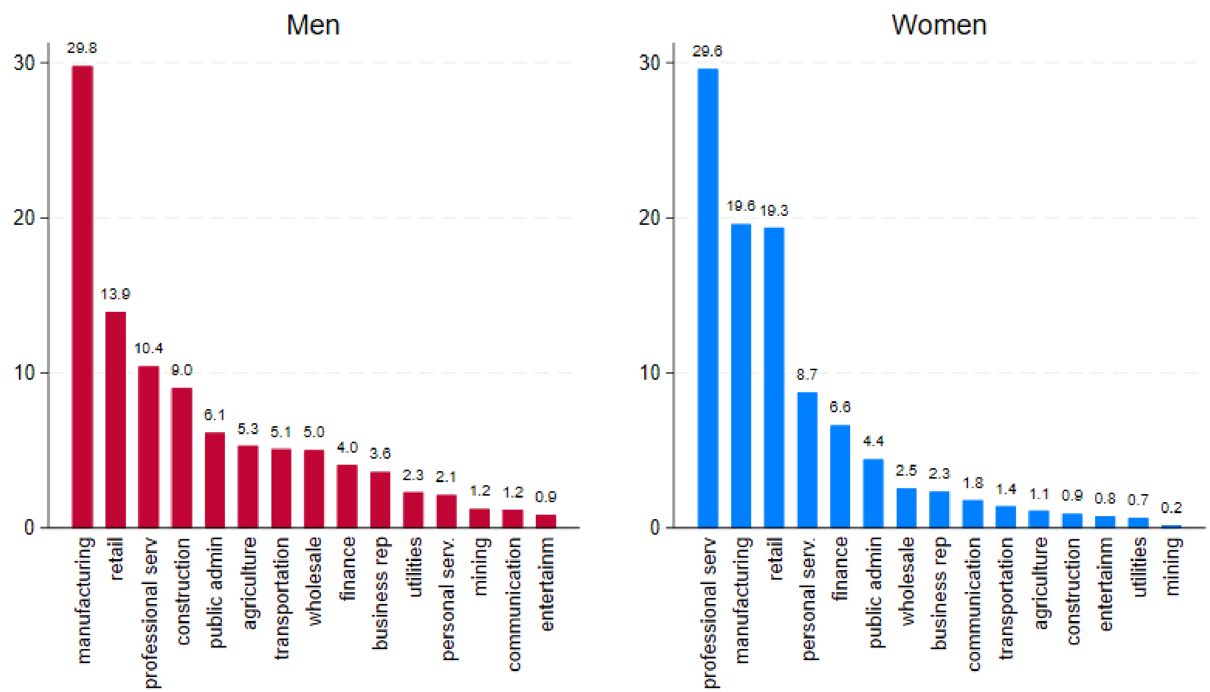

Source: 1% Census Microdata 1970

B.3 Occupation shares

Occupational codes are not consistent across Census years. Hence, I first create a balanced panel of consistent occupations over the period 1970-2010 using the crosswalk provided by Autor and Dorn (2013). Then, I group the consistent occupational codes into 28 categories (see Table B.3 for a list of the occupations. Figure B.2 shows the proportion of each gender employed in each occupation. Men are disproportionately employed in managerial jobs, while women are disproportionately represented in administrative support jobs.

Table B.2. Occupation classification

1	Executive, administrative and managerial occupations
2	Engineers, architects and surveyors
3	Other technicians
4	Physician/nurses
5	Health assistants
6	Teachers and therapists
7	Lawyers and judges
8	Writers, artists, entertainment, and athletes
9	Sales occupations
10	Administrative support
11	Cleaning services
12	Other personal services
13	Protective services
14	Food services
15	Farm and forestry workers
16	Mechanical and electronic repair
17	Construction trades
18	Mining extraction
19	Production supervisor
20	Precision metal and wood-working
21	Precision textile, food and assorted materials
22	Plant operators
23	Metal and plastic machine operators
24	Textile machine operators
25	Miscellaneous machime operator
26	Fabricators, assemblers, inspectors and testers
27	Transportation occupations
28	Movers

Source: 1% Census Microdata 1970

Figure B.2. Proportion employed in each occupation by gender - 1970 Census

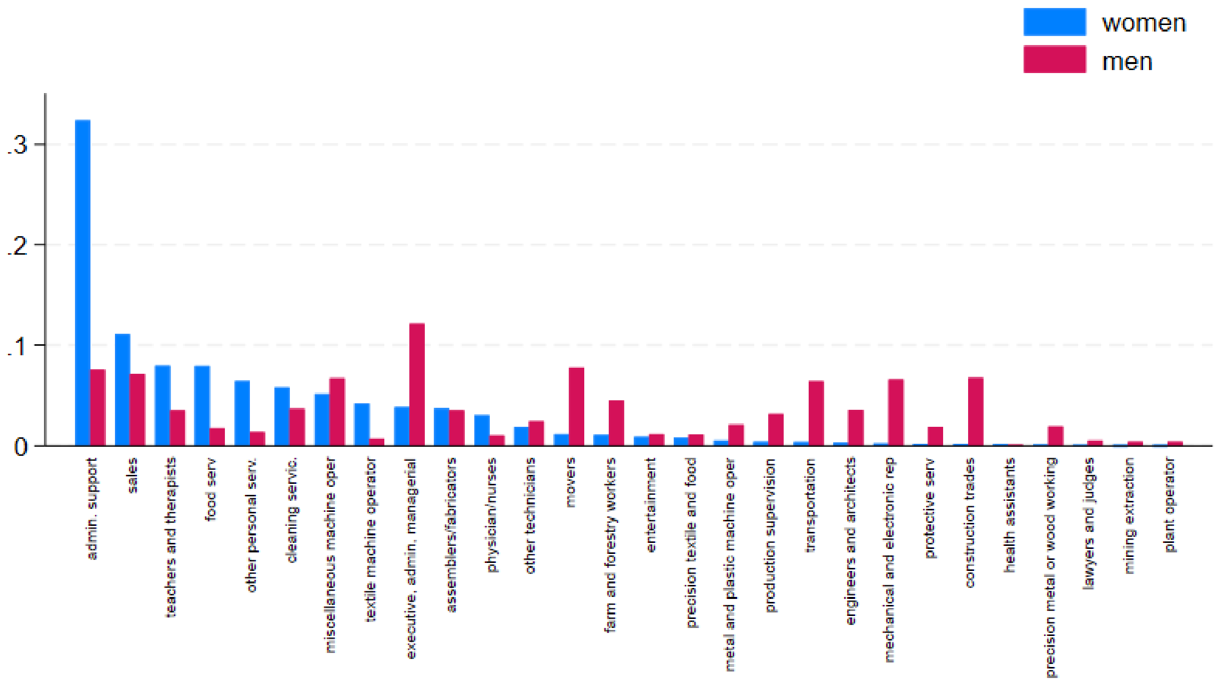

Source: 1% Microdata Census 1970