STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
5350 Master's Thesis in Economics
Academic Year 2018–2019

Female Labor Supply and Earnings Inequality Under Skill-Biased Technological Change

Jörn Onken (41211)

Abstract

The gender education gap has reversed, women today account for the majority of college graduates. At the same time, skill-biased technological change strongly rewards highly-educated workers through the increased skill premium. In this thesis, I analyze the implications of combining these two facts through a heterogeneous agent model using U.S. data from 1964 to 2018, in which agents differ in terms of gender and education. Extending the model of Caselli and Coleman (2006), I first investigate how firms shift their production towards more skilled-labor-efficient technologies over time. I find that a decrease of the relative price of skilled labor productivity over time is needed to generate a skill premium growth. Next, I show that the model fits the data on labor supply and both intra- and inter-household earnings inequality very well. Then, four potential explanations for the sharp increase of female labor supply in the last century are tested: Decreasing home production hours, assortative matching, rising female education and skill-biased technological change. The results show that the great changes in time use, enabled in larger part by more effective home production technology, are the most important factor for female labor supply and intra-household inequality. Lastly, a counterfactual analysis is conducted to determine whether these model results differ between married and cohabiting households, which they do as cohabiting women work more and have higher relative earnings compared to married women.

Keywords: Female Labor Supply, Skill-Biased Technological Change, Earnings Inequality, Family Economics, Aggregate Productivity, Heterogeneous Agents
JEL: E23, J21, J22, J24

Supervisor: Kelly Ragan
Date submitted: 13 May 2019
Date examined: 27 May 2019
Discussant: Johannes Matt
Examiner: Mark Sanctuary

Contents

List of Figures

List of Tables

1 Introduction

The increasing educational attainment of women in industrialized countries is one of the most striking phenomena anyone analyzing microdata encounters. Within just 50 years, women have overtaken men in terms of college graduation rates (Flood, King, Rodgers, Ruggles, & Warren 2018), described as the reversal of the gender education gap. Furthermore, female labor supply has experienced significant growth on both the extensive and the intensive margin since the 1960s. Another trend evident in the data and widely discussed in public is the increasing income inequality within developed countries. One important cause of this inequality in the U.S., as noted by Katz and Autor (1999), is skill-biased technological change. This term describes the increasing relative productivity of skilled compared to unskilled workers, and can explain why the relative wage for highly-educated workers, called the skill premium, has risen tremendously over the last few decades. Learning about these important trends regarding gender, education, inequality, and skill-biased technological change motivated me to connect them in a quantitative economic model in order to better understand how they are related.

In particular, the scope of this thesis is to examine different candidates that could explain the observed time trends in female labor supply and earnings inequality, both within and between households. Hence, the research question I will answer is two-fold: What are the drivers behind the large increases of female labor supply in the U.S., and how do these changes affect earnings inequality between and within households? In order to investigate these questions, I set up a structural partial equilibrium model, consisting of heterogeneous households and a representative firm. My model enables me to test how well each of the four different candidates – the reduction in home production time, changes in spousal matching, increase in female education and skill-biased technological change – can replicate the empirical trends.

One of the main features of my model is to allow for household heterogeneity. While this complicates the quantitative analysis considerably compared to the standard assumption of the representative household underlying most macroeconomic models, it is necessary for answering my research question. This argument is supported by Doepke and Tertilt (2016) who note that abstracting from household heterogeneity becomes particularly problematic when the differences between households are subject to changes over time. As I will present in the following chapters, there have been major changes to female labor supply and educational attainment in the U.S. over the last six decades (Flood et al. 2018), which strongly suggests that modeling multiple-member households is vital for analyzing aggregate labor supply and consumption inequality.

My structural model is based on Caselli and Coleman (2006), one of the most noteworthy works published on skill-biased technological change. One of the main contributions of this thesis is to embed the production function from Caselli and Coleman (2006), who only model the production side of the economy, into a model with heterogeneous agents, while adding gender as a second dimension of differentiation besides skill level. Therefore, I extend the Caselli and Coleman model by incorporating both households and gender, which to my knowledge constitutes the first attempt at doing so while maintaining skill heterogeneity. In addition, instead of analyzing firms' production technology for one year across countries as Caselli and Coleman (2006) do, I consider the U.S. over a period of 54 years, using data from the IPUMS-CPS (2018) and the IPUMS-AHTUS (2018), for the purpose of replicating the stylized facts regarding female labor supply and earnings inequality in the U.S.

The further contributions of this thesis consist of determining the decrease of the relative price of skilled labor productivity over time for the U.S., as well as examining different candidates for explaining female labor supply and determining their effect on earnings inequality. My main finding is that the reduced time spent by women working in the household compared to previous decades, enabled for example through more efficient technologies and changing social norms incentivizing men to participate in household work, is by far the most important driving force behind the changes in female labor supply and earnings inequality. Thus, my thesis supports the findings of Greenwood, Seshadri, and Yorukoglu (2005) and underlines the great importance of freeing up time for women in increasing their labor market attachment.

This thesis is organized as follows: Chapter 2 provides an overview of the extensive literature on female labor supply and skill-biased technological change. Chapter 3 illustrates the main facts motivating the quantitative analysis, which will begin in chapter 4 with the introduction of a structural heterogeneous agent model. Chapter 5 describes how the model is parameterized and solved in MATLAB. Chapter 6 presents the model results, including a discussion on how to quantitatively incorporate skill-biased technological change into the model. Chapter 7 provides counterfactual analyses in order to examine different candidates for explaining female labor supply changes, as well as to determine the model results and fit for the subsample of unmarried, cohabiting households. In the end, I will conclude my thesis with a brief summary of the results and a discussion of the limitations of my work as well as possible extensions of the model to overcome these.

2 Literature Review

This thesis is combining two features: The changes in female labor market and educational attachment on the one hand and skill-biased technological change on the other, each of them affecting earnings inequality. Therefore, I will provide an overview of both strands of literature and subsequently discuss how they are connected. While I cannot claim that examining the relationship between these two fields is a novel approach, the research on this topic is still fairly young. Notable work includes Greenwood, Guner, Kocharkov, and Santos (2016), Cerina, Moro, and Rendall (2017), Rendall (2017), Ngai and Petrongolo (2017) and Cortes, Jaimovich, and Siu (2018), which I will accentuate in detail after presenting both research areas individually. This chapter is concluded with summarizing evidence on the role of female labor supply and skill-biased technological change for earnings inequality in the U.S.

To begin with, there exists an extensive literature on the properties of female labor supply in the US and the driving forces behind its shifts over time. In order to provide a structured overview of the multifaceted work done on female labor force participation (LFP) over the last decades, I present the most decisive factors for its development in the following order: The changes in the division of household labor, the compatibility of motherhood and working for women, the increasing returns to labor market experience, child care and parental leave policies, the taxation of married couples, cultural and social norms, intra-household bargaining, the effects of changing household composition in the U.S. (especially assortative matching), the increased educational attainment of women, and the effects of the increasing skill-premium caused by skill-biased technological change and sectoral transformation.

One of the most notable changes over the past decades is how considerably different the day of an average American woman in 2019 looks from one in the 1960s. Despite being only two generations apart, the everyday life of a housewife in the 1960s, dominated by hard manual housework, seems extremely outdated today. However, since then shifts in the division of labor within the household and technological innovations have freed up time for women to work in the labor market, as a vast literature on time use in the U.S. documents. Aguiar and Hurst (2007) study time use by gender in the U.S. between 1965 and 2003 and find that men substituted part of their reduced market work hours with home production, whereas women used some of the time gained by reduced household work to participate in the labor market.

Analyzing the American Heritage Time Use Study (AHTUS), Fisher, Egerton, Gershuny, and Robinson (2007) support these results by showing that leisure increased for both genders, while work hours of men and women converged over time. Furthermore, Ramey and Francis (2009) provide evidence that the rise in female compensated the decline in male working hours almost perfectly, such that the average prime age individual today spends as much time working in the labor market as in 1900.

These far-reaching shifts in time use over the last 100 years raise more questions than they answer. While the overall patterns in the data are well-documented, identifying the drivers behind these changes proves to be much more difficult. One possible explanation is put forward by Greenwood et al. (2005) who argue that advances in home appliances increased female LFP by reducing the time married women had to devote to home production. Their main result is that the model reproduces both the increasing female LFP and the declining household work time by women found in the data. According to Greenwood et al., the narrowing of the gender pay gap on the other hand was not as important for increasing female LFP, since women could not have allocated enough time to working in the market without more efficient household appliances granting them more time to allocate outside of the household.

Contradicting these results, Jones, Manuelli, and McGrattan (2015) find that household technology improvements cannot explain the rise of married female LFP after World War II. They attribute their deviating results to Greenwood et al. (2005) using a Leontief production function for the home good and assuming female labor to be indivisible, meaning Greenwood et al. do not consider part-time work despite it constituting a significant part of the female work force (Jones et al. 2015). Instead, Jones et al. come to the conclusion that even small decreases in the gender pay gap can explain both the increase of married and the stagnation of single female LFP after World War II in the U.S.

Besides household production, the other important reason for women spending more time on non-market work than men, and one that cannot be divided between partners as easily as doing the dishes or buying groceries, is having children. From pregnancy over child birth to raising children, having children exerts a significant time (and subsequently monetary) cost that is still mostly born by women. This is in part due to unchangeable biological reasons, but also to social norms and expectations as well as preferences. The degree of (in)compatibility between having children and participating in the labor market is therefore particularly important for female labor supply on both the extensive and the intensive margin.

Accentuating the significance of fertility for female LFP, Erosa, Fuster, and Restuccia (2016) show that the impact of children on female labor supply explains about half of the divergence in hourly wages over the life-cycle between men and women. This huge effect is caused mainly by mothers' lower expected employment rate and working hours due to having children. The remainder of the gender pay gap can be attributed to women working around 10% less than men when employed, regardless of whether they have children (Erosa et al. 2016). Bertrand, Goldin, and Katz (2010) study the gender pay gap for MBA graduates and find that having children comes with a huge earnings penalty, especially after the first year after giving birth. Interestingly, this effect is reversed for fathers, who earn more than childless men (Bertrand et al. 2010). In addition, work hour differences, which are also mainly due to having children, can explain a large part of the gender difference in pay (Bertrand et al. 2010).

While less debated, in addition to the gender wage gap, there is also a wage gap between mothers and childless women called the family gap. Waldfogel (1998) finds that a woman's average wage decreases by 4.6% for the first child, while Erosa, Fuster, and Restuccia (2010) estimate that the average wage of a women with children is only 89% of that of a women without children.

Concerned with the link between education and fertility, Hazan and Zoabi (2015) illustrate that fertility and education are related in a U-shape, meaning that low- and high-educated women have the most children. While they state that fertility is usually decreasing in education and income within a country, which can explain the high fertility of low-educated women, the reason for fertility increasing again for high-educated women is unclear. Their explanation for this phenomenon is that highly educated women purchase market substitutes for home production such as child care and housekeeping that allow them to combine high fertility and market work. Mothers with low education (and low income) on the other hand do not marketize household production as much and engage in most of the child-raising themselves. In addition, Hazan and Zoabi also find that highly educated women invest more in their children, for instance in the quality of their education, which increases the growth of economic inequality over time.

Moreover, similar to household labor, technological progress has affected the time requirement of raising children. Albanesi and Olivetti (2016) provide an example for this by showing that the increasing diffusion of infant formula between the 1930s and 1960s allowed mothers to combine market work and child care more easily. Together with advances in maternal health, this allowed female LFP to increase alongside a fertility rise. In addition, medical progress has fundamentally reduced the uncertainty fertility imposes on female education and labor supply decisions. Goldin and Katz (2002) present evidence that oral contraceptives introduced in the 1970s have increased the returns to female education, as they reduced the costs of investing into a long-term professional career and made planning a career easier. This in turn increased their labor market participation later on when married.

Another significant channel of how a woman's fertility decision affects her labor market outcome, besides the direct time cost, is through the immense importance of work experience in today's world and the subsequent loss of it which working mothers inevitably incur. The earnings benefit of having accumulated job market experience is being referred to as the returns to experience and can be estimated empirically. A noteworthy study doing exactly this has been conducted by Olivetti (2006), who finds that the returns to experience overall have risen between the 1970s and 1990s, but that this increase was more distinct for women than for men. This in turn raised women's opportunity costs of temporarily dropping out of the labor force during motherhood and of working less hours during child-rearing, contributing to their LFP increase (Olivetti 2006).

Moreover, using abortion legislation as an instrument for fertility, Bloom, Canning, Fink, and Finlay (2009) estimate the reduction of a mother's lifetime labor supply associated with having one child to be almost two years. Combining this with the increasing returns to experience for women found by Olivetti, the detrimental and long-lasting effect of dropping out of the labor force for two years due to motherhood on women's careers are evident.

As the significance of returns to experience demonstrates, reducing the length of the career interruption associated with fertility is crucial in ensuring that women can combine motherhood with working in the labor market. Two important instruments to facilitate this are child care and maternal leave policies, and the effects of both have been examined extensively by economists.

Erosa et al. (2010) analyze the impact of parental leave policies on fertility and labor supply using a search and matching model with human capital accumulation on the job. Their main result is that voluntary parental leave policies have no effect or can even reduce the time women spend on child-rearing, whereas mandatory parental leave policies lead to increases in both the fertility rate and female participation. Domeij and Klein (2012) examine whether child care can raise welfare by allowing mothers to work and find that half-subsidizing day care almost doubles LFP among mothers with small children. Likewise, Bick (2016) also finds large positive effects of subsidized child care on mother's labor supply using German data, in particular for part-time-working mothers of children aged zero to two. Ragan (2013) considers cross-country data and notes that day care subsidies are crucial in explaining the high levels of female labor market participation in Sweden, as day care increases market hours while reducing hours spent on non-market work at home. Guner, Kaygusuz, and Ventura (2014) underline the importance of child care subsidies for female LFP, particularly for less-educated women who on average have more children than college graduates. They estimate that introducing a fully subsidized, universal child care program in the U.S. would raise female participation rates as well as hours worked.

The effectiveness of child care and maternal leave policies in raising female LFP depends in large part on their accessibility. Martínez and Iza (2004) relate the increased availability and affordability of child care services in the U.S. to the rising demand for high-skilled labor induced by skill-biased technological change. They argue that as the demand and therefore the wages for skilled relative to unskilled labor (the skill premium) grow, the relative cost of child care, which does not require highly-educated workers, decreases. This in turn made purchasing market child care more affordable for women, allowing them to allocate more time to working in the labor market. Hazan and Zoabi (2015) also find that the costs of child care relative to female wages have decreased, opening it up to women who were previously not able to afford it to the same extent. However, unlike Martínez and Iza, Hazan and Zoabi note that this relative cost decrease only applies to highly educated women, and that child care has in fact become more expensive for less-educated women.

Besides child care and parental leave subsidies, the taxation of labor income for married couples also plays an important role for female LFP. This is proven by Bick and Fuchs-Schündeln (2017), who investigate differences in aggregate labor supplies of married couples across the U.S. and selected European countries. Their results show that the cross-country differences in female labor supply within Europe can be explained by variations in tax systems. Namely, joint taxation of married couples, as practiced in the U.S. and Germany, increases the marginal tax rate faced by the second earner (mostly the wife) and hence disincentivizes her from working (the so-called marriage penalty), whereas separate taxation systems like in Sweden do not exert a negative effect on second-earner employment. The result of Apps and Rees (2004) that individual taxation of couples increases both female LFP and fertility supports this.

In contrast to these quantifiable policy instruments, much more difficult to measure, yet nonetheless crucial factors for female labor supply are social norms and values. Noteworthy studies on this challenging matter include Fernández, Fogli, and Olivetti (2004), who argue that the cultural transmission of social attitudes towards working women is passed on across generations. Notably, men obtaining their preference for marrying either a working or a stay-at-home wife from their mother's behavior leads to an increase in female LFP over time. Similar to this study, Fernández (2013) assumes that attitudes towards working married women are passed on to the next generation and finds that especially the perceived long-run costs of working for women are vital in this process. In the model of Fogli and Veldkamp (2011), women learn about the effects of fertility on female LFP through the number of working women in society. This results in accumulating effects, as more women working implies a larger growth of female LFP.

Closely related to social norms and values, intra-household bargaining models also serve as prominent explanations for the increase of female LFP in the literature. For instance, Eckstein and Lifshitz (2015) assume that households can be one of two types: In classical households, the husband sets his own optimal labor supply, while the wife treats his decision as given (non-cooperative bargaining). On the contrary, spouses in modern households play a symmetric Nash game (cooperative bargaining). Eckstein and Lifshitz find that in modern households, women work about 10% more, while male labor supply is not affected. Therefore, if one assumes a shift in social norms over the last decades that led to more households being of the modern type, the model of Eckstein and Lifshitz provides a good explanation for both the rise in female and stagnation in male labor supply.

Besides through household bargaining, the fact that married women's labor supply is influenced by spousal decision making is further reflected in the fact that civil status is an important factor for explaining female LFP. For instance, Browning, Chiappori, and Weiss

(2014) note that the pattern of hours and wages between married and single individuals is reversed between genders: While married men both work more hours and earn higher wages than single men, single women work and earn more than married women. According to Browning et al., two effects can explain this pattern: The division of household labor within married couples, where women typically do the majority of chores, and a marriage selection bias caused by mostly high-wage men and low-wage women being able and willing to enter marriage. However, Browning et al. state that these wage differences by marital status are less pronounced for recent cohorts, since married women today participate more in the labor market which makes their wages converge to those of males.

Further, one of the most influential factors shaping both male and female labor supply is assortative matching, meaning that people tend to choose partners with characteristics similar to their own, in particular regarding education. Notably, according to Browning et al., the share of couples with equal schooling is relatively constant at about 50% of married couples in the U.S. On the contrary, the share of couples in which the wife is higher educated has surpassed that of couples with a higher educated husband for couples with husbands born in 1960 and later (Browning et al. 2014). This marks the reversal of the educational gap between men and women, and assessing its implications for labor supply and income inequality will be part of my quantitative analysis. Nonetheless, Browning et al. also demonstrate that, while there is a strong correlation of 0.7 between spouses' school years, the correlation of their log wages is considerably lower with a correlation coefficient of only 0.3 (although it has been increasing from 0.2 in the 1970s). This indicates that the convergence of female to male school years does not perfectly translate to wage convergence.

A significant channel of how this changing household composition matters for labor supply is through household insurance against labor market shocks. Choi and Valladares-Esteban (2016) study the role of household insurance through spousal income against labor market shocks under a system of publicly provided unemployment insurance (UI). By examining the effects of these shocks on single and married households, they demonstrate large effects of self insurance. Notably, in their model UI does not improve the welfare of married households, who are self-insured through the pooling of household earnings, joint savings and the added worker effect. In consistence with this result, Blundell, Pistaferri, and Saporta-Eksten (2016) report strong evidence for the importance of family insurance within two-earner households in smoothing wage shocks, especially for poor families that do not hold assets. Furthermore, Shore (2015) studies couples' income dynamics and finds that the volatility of spousal incomes is positively correlated, which could be explained through assortative matching.

As Browning et al. (2014) note, the remarkable rise of the share of couples with a more educated wife cannot be fully explained by changing preferences in spousal matching.

Rather, the increased educational attainment of women is a core driver behind the shifts in household composition. The fact that women today are obtaining more education than men is referred to as the reversal of the gender education gap, and it is particularly relevant for female labor supply and earnings because of the growing skill premium, which describes that the monetary returns (in terms of future labor income) of being more educated have been starkly increasing in recent decades. For instance, Albanesi and Prados (2017) argue that the skill premium rise during the 1990s led to huge increases of the top incomes for (college-educated) men, which subsequently caused a decline in the LFP of their wives through a household income effect. According to Albanesi and Prados, this effect was amplified by the increase of assortative matching and can explain why female LFP flattened in the 1990s.

Mulligan and Rubinstein (2008) attribute the reduction of the gender wage gap in the U.S. since the 1970s to unobservable changes in the composition of the female workforce. They describe that the increased demand for skilled labor caused women to invest more into their education. Subsequently, it became more expensive for these well-educated women to remain out of the labor force due to the large opportunity costs, which increased female LFP particularly for high-skilled women. In line with these findings, Chiappori, Iyigun, and Weiss (2009) also show that the overall returns to schooling have increased more for women than for men. They state that through positive assortative matching, investing in their education not only generates a labor market, but also a marriage market return for women. The reason for this is that being more educated grants women a larger share of the martial surplus, both because their outside option improves and because they can generate a larger household income upon marriage. Chiappori et al. argue that the lower educational investment of women compared to men in the past was mostly due to household work, which decreased the returns from education after marriage. Nowadays, while the market return to schooling has increased for both genders, according to Chiappori et al., women invest more in education than men because the reduction in household hours due to technological progress has increased their marriage return to education.

The main reason for the steep increase of the returns to education (the skill premium) is skill-biased technological change, meaning that the relative productivity of skilled compared to unskilled labor has increased rapidly since the 1970s. Early important work in this field includes Katz and Murphy (1992) who show the increase of the college wage premium in the U.S., and Katz and Autor (1999) who also find that the wage difference between college and high school graduates has increased sharply after an initial drop in the 1970s. In addition, according to Acemoglu (1998) firms choose their production technology endogenously depending on the relative supplies of skilled and unskilled labor, which he calls the "directed technology effect". However, as technologies are fixed in the short-run, an increased availability of skilled labor initially reduces the skill premium due to the supply effect, which can explain the college wage premium decrease in the U.S.

during the late 1960s and 1970s (Acemoglu 1998).

Caselli and Coleman (2006) examine cross-country differences in production with skilled and unskilled labor being imperfect substitutes in a model where firms choose an optimal production technology based on the availability of both types of labor. Their results indicate that rich countries use skilled labor more efficiently than poor countries because firms in these countries choose technologies favoring skilled labor, as it is available in abundance. Moreover, Caselli and Coleman find that the ratio of skilled to unskilled labor productivity is larger in more developed countries, and grows over time, generating skill-biased technological change. Interestingly, they show that the variation in labor productivities accounts for 40% of cross-country variation in income.

Combining these findings with the shifts in female labor supply, matching and education documented above, it becomes evident that skill-biased technological change affects female labor supply and earnings in particular. One of the main reasons for this is that it implies a sectoral shift to services where women have a comparative advantage over men, as many authors have illustrated. In the following, I will therefore present noteworthy work highlighting how female labor supply and skill-biased technological change are related.

An early paper on this is written by Galor and Weil (1996), who study potential causes of the growing demand for female labor. They show that technological innovations increased the returns to skilled-labor occupations, in which women have a biological comparative advantage over men in contrast to physically demanding manual labor jobs. Further, they find that this factor, together with the increasing labor market participation of women, explains why the gender gap in wages has been decreasing over time. Consistent with these findings, Black and Juhn (2000) show that the rising demand for skilled labor contributed to the rise in skilled female LFP. According to them, the additional demand for skilled labor was primarily served by women, since most skilled men were already participating in the labor market. These well-educated women subsequently mostly switched from non-participation or low-skill jobs to positions in need of high-skilled labor, increasing female LFP along both margins.

However, Blau and Kahn (2007) argue that skill-biased technological change, while increasing the demand for skilled female labor, would not boost female LFP any further, but rather increase women's wages relative to men. They base this prognosis on an estimated drop of female wage elasticity by 50-56%, meaning that higher wages would not imply further female labor supply growth. Herrmann and Machado (2012) measure the selection of men and women into jobs using high school test scores which serve as a measure of cognitive ability before entering the labor market. Their results show that the gender wage gap cannot be explained by the job selection of women into either high- or low-paying job. Herrmann and Machado find no relationship between ability (as measured by the test score) and full-time work participation, proving that working women are not self-selected exclusively from either end of the skill distribution.

In addition, Greenwood et al. (2016) examine two exogenous trends, the technological progress in household production and the wage changes over time due to the increasing skill premium and decreasing gender wage gap. They show that these two can explain the patterns of marriage and divorce and female LFP present in the data. According to Greenwood et al., technological improvements in household production constitute the most important factor for the increase of married female LFP, followed by the narrowing of the gender wage gap. The rise of income inequality between households on the other hand can be attributed in large part to the growing college skill premium, which is strengthened by the prevalence of assortative matching implying that high-earning men are matched with high-earning women. However, Greenwood et al. note that the effect of assortative matching on inequality is only effective when women are working, as otherwise the inequality between households is the same as between men.

Furthermore, Cerina et al. (2017) examine employment polarization in the U.S., which describes the increase of LFP rates (the extensive margin of labor supply) at the bottom and top of the skill distribution, while employment in the center declines. They show that that this pattern is mostly caused by women, whose employment polarization is driven by skill-biased technological change after 1980, causing high-skilled women to substitute household production hours with market work. In order to spend less time in the household, these high-skilled women purchase market substitutes for home production (for instance helpers, cleaners and nannies), and these substitutes are mostly provided by low-skilled women. Subsequently, the LFP rates of women at the top and bottom ends of the skill distribution increase, while women at the center of the distribution are less likely to work.

Moreover, Rendall (2017) studies women's comparative advantage under two labor demand shifts, a standard skill-biased one and a brain-biased one where brain inputs become more productive than brawn inputs in both high- and low-skilled jobs. Most notably, her model offers a good explanation for the narrowing of the gender wage and education gaps in the U.S. since the 1980s. According to Rendall, these effects can mostly be explained through a brain-biased labor demand shift, which can also be interpreted as a gender-biased shift due to the comparative advantage of women in brain-biased jobs. Interestingly, she finds that this brain demand shift was especially pronounced in the unskilled sector.

Likewise, Ngai and Petrongolo (2017) examine the effects of the growing service sector share on the gender gaps in hours and wages, assuming that women have a comparative advantage in producing services. Their results show that the growing dominance of the service sector increases women's relative wages and hours worked. They provide two explanations for why women may have a comparative advantage in services. The first reason states that women have a clear disadvantage against men in brawn-intensive occupations, while the second argues that women have superior communication and interpersonal skills compared to men, and these skills are irreplaceable by automation. Ngai and Petrongolo argue that this female advantage in the tertiary sector is reflected in wages, since the rise of the service sector share is associated with growing female relative wages.

The last paper I present regarding skill-biased change and gender is that of Cortes et al. (2018), who investigate why the share of college-educated men working in cognitive, high-wage occupations has been decreasing, whereas the share of women in these jobs has grown. Similar to Ngai and Petrongolo, they bring up the argument that the demand for female-oriented (social) skills in high-wage occupations is increasing relative to other occupations, which introduces a female bias. According to Cortes et al., this growing demand for social skills is the main reason for the observed gender reversal in these occupations, which cannot be explained by the increasing female share among college graduates alone.

To conclude this chapter, I summarize relevant studies on how the increased female labor market attachment has affected pay inequality. For example, Acemoglu, Autor, and Lyle (2004) estimate the effect of female LFP on the wage structure in the U.S. utilizing the state-level variation in military mobilization during World War II, since women worked more after the war in states with high mobilization due to the greater number of fallen men. They find that higher female LFP lowered the wages of both genders, although the stronger decrease of female wages indicates that male and female labor are imperfect substitutes. Furthermore, states with greater female LFP saw increased earnings inequality between high- and low-skilled men, which Acemoglu et al. attribute to the low level of female education at the time, meaning that women competed mostly with low-skilled men, in turn increasing the returns to education for men.

Furthermore, positive assortative matching reinforces inequality, as several studies have demonstrated. Hyslop (2001) finds that female labor supply can explain much of the rise in family and in female earnings inequality in the early 1980s. Moreover, he estimates that positive assortative matching can explain about 25% of the increase in permanent family earnings inequality. In support of this, Fernández and Rogerson (2001) as well as Greenwood et al. (2016) also show that increased assortative matching raises household income inequality. Gihleb and Lifshitz (2016) indicate that the share of couples in which the woman is more educated is increasing, and that women in these couples are more likely to work than women who married a more educated husband. They argue that this furthers inequality between households and is much more important for inter-household inequality than positive assortative matching.

In addition, Card and Hyslop (2018) study why earnings inequality among women has declined since the 1960s. They report a decreasing relevance of family factors like the number of children and spousal income for female labor supply, which contributed a major part to female earnings inequality in the past. Heathcote, Storesletten, and Violante (2010) investigate the effects of increasing wage inequality on overall welfare and find that today's higher relative wages for college graduates and for women have led to more schooling and a more even division of household work between genders. Moreover, while gender- and especially skill-biased technological change increased average welfare, the household type consisting of two low-skilled (high school graduate) partners incurred a welfare loss caused by the rising skill premium.

To summarize, female labor supply, education, wages and earnings in the U.S. have all seen tremendous changes over the past decades, which in turn has affected household composition and inequality. The next chapter is therefore dedicated to illustrating these trends using empirical U.S. data.

3 The Facts

Before introducing my model, I present empirical trends and patterns regarding female labor supply, time use, spousal matching, education, wages, skill-biased technological change and earnings inequality in the U.S. These stylized facts provide evidence of huge shifts in many key variables over the last decades, such as the increasing returns to education and stronger attachment of women to the labor market. Since these findings motivate my analysis of female labor supply under skill-biased technological change and set the stage for the quantitative part of this thesis, it is important to know of the main trends in the data before commencing any quantitative study.

For my analysis, I use U.S. microdata obtained from IPUMS for the Current Population Survey (CPS) from 1964 to 2018 (Flood et al. 2018) and the American Heritage Time Use Study (AHTUS) from 1965 to 2012 (Fisher, Gershuny, Flood, Roman, & Hofferth 2018). The main reason for using the U.S. Current Population Survey over the decennial U.S. census data (IPUMS-USA), which is also available on IPUMS and widely used in the literature to show long-term trends of female labor supply, is its annual availability. Furthermore, the early time use surveys contained in the AHTUS were conducted in 1965, 1975, 1985 and 1995, whereas the census data was gathered every ten years from 1790. Hence, for better comparability of household survey and time use data, I use the CPS.

Following Mulligan and Rubinstein (2008), who also utilize the CPS to study female labor supply and wages, I impose an age restriction of 25 to 55 years. Restricting on the minimum age of 25 ensures that most individuals will have completed their undergraduate college education, which is imported since education enters my model exogenously. The maximum age is set to 55 to mitigate the effects of retirement on labor supply and time use. Additionally, like Mulligan and Rubinstein (2008) and Herrmann and Machado

(2012), I restrict the sample to non-Hispanic whites to avoid changing demographics over time being the driver behind patterns in the data. One example for this would be that black women, especially when young, work more in the labor market than white women (Fernández 2013), meaning that an increasing share of African Americans would increase female labor supply ceteris paribus, without any changes in individual behavior. In line with Mulligan and Rubinstein (2008) and Herrmann and Machado (2012), I also drop all military personnel, as they do not report hours worked.

Moreover, since households in my model are composed of opposite-sex couples, I restrict the sample to men and women living with or married to a partner of the other gender. In addition, since one of the assumptions of my model is that the husband is always working, I only consider households where the male partner is working for pay for any positive amount of hours. This also allows for better comparability between households, as the household production time of both partners would differ when the husband is unemployed. Therefore, unless stated otherwise, all trends presented in the following referring to CPS and AHTUS data are for this sample of heterosexual couples in which the male partner is working and receives non-zero earnings. A full description of the data, the number of observations, the sample restrictions I imposed and how I constructed the key variables can be found in the Data Appendix A.

Regarding educational attainment, I follow the skill-level distinctions by Cortes et al. (2018) and Greenwood et al. (2016) and consider an individual with four or more years of college education (the typical duration of most undergraduate degree programs in the U.S.) as high-skilled, and everyone with less than four years of college education as low-skilled. Hence, the sample consists of four different household types: Two in which both partners are equally educated, and two in which either the male or female partner is the only one with a college degree.

This chapter is organized in five sections presenting time trends of the key variables most important for my quantitative analysis. First, I will outline the increase of female labor force participation, especially for skilled women, as well as the tremendous decrease in time spent on household production by women. Subsequently, I show patterns of spousal matching based on education and of the share of cohabiting (non-married) couples. Following this, I present the reversal of the gender education gap caused by the increasing educational attainment of women. Since skill-biased technological change is one of the core forces in my model, I will then demonstrate the growing skill premium through plotting the development of real wages over time by gender. I conclude this chapter by showing the growing earnings inequality between U.S. households as well as the declining earnings inequality within households.

3.1 Trends in Labor Supply and Time Use

First of all, female labor force participation (LFP) has increased greatly since the 1960s, especially for married/cohabiting women. This trend is evident on both the extensive margin (the percentage of women in the labor force, called the LFP rate) and the intensive margin, typically measured by annual hours worked. Figure 3.1 shows the immense increase of labor market participation on both margins of married/cohabiting women over time. Both the female LFP rate and the average annual hours worked were computed using the individual-level survey weight *ASECWT* over all women in the sample, including both full-time and part-time working women.

Figure 3.1: Labor Force Participation Rate and Average Annual Hours Worked of Women

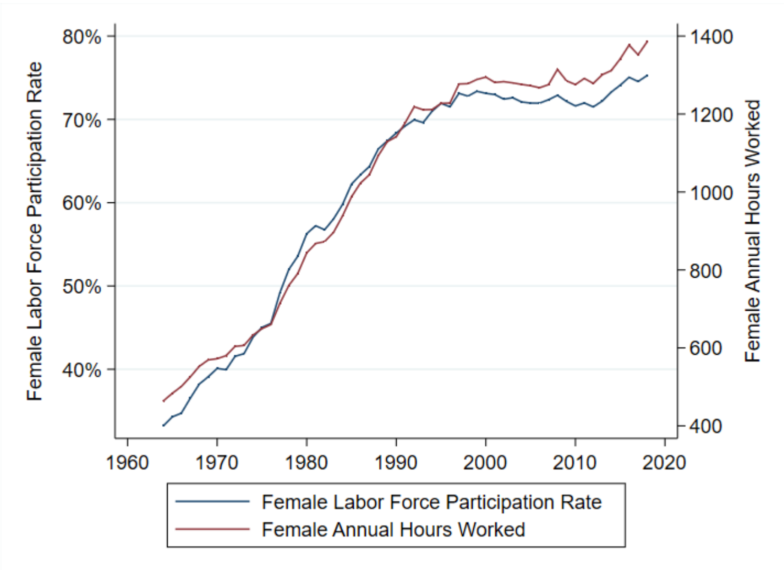

Note: Aged 25 to 55, non-Hispanic whites, living with a working partner
Source: Author's rendering of IPUMS-CPS data (2018)

As can be seen in figure 3.1, both measures have increased starkly since the 1960s. The female LFP rate has more than doubled from 33% in 1964 to 75% in 2018, meaning three in four married/cohabiting women are working today. Even more distinctively, annual hours worked have almost tripled from 464 to 1,385. Interestingly, an almost identical growth pattern can be observed in both variables. After a steep increase over the late 1960s, the 1970s, and 1980s, both measures begin to stagnate in the mid-1990s, and this stagnation has lasted until today. Both rates display a modest rise starting in 2012, with the relative growth being larger for annual hours worked. From 1995 to 2018, annual hours grew by 13% in total, while the LFP rate increased by only 5%. This indicates that, in addition to more women entering the labor market, the already working women have increased their working hours over the last 20 years.

While the increase of female labor market attachment (especially in annual work hours) is tremendous, it is important to note that women are still working significantly less than men, whose average annual hours worked have stagnated at around 2300 hours since the 1960s. The main reason for this big gap between male and female hours worked is that, despite the consistently high female LFP rate of around 75% since the 1990s, most women do not work full-time. Some of the reasons for this, such as child care and differences in home production between genders, have been discussed in the literature review.

Figure 3.2 depicts the survey weight-adjusted shares of married/cohabiting women who are working full-time full-year (FTFY, defined as having worked for pay at least 50 weeks last year and at least 35 hours last week), part-time (all other working women), and not at all. Within each year, these three fractions sum up to 100%, denoting all women in the sample.

Figure 3.2: Fractions of Women by Work Status

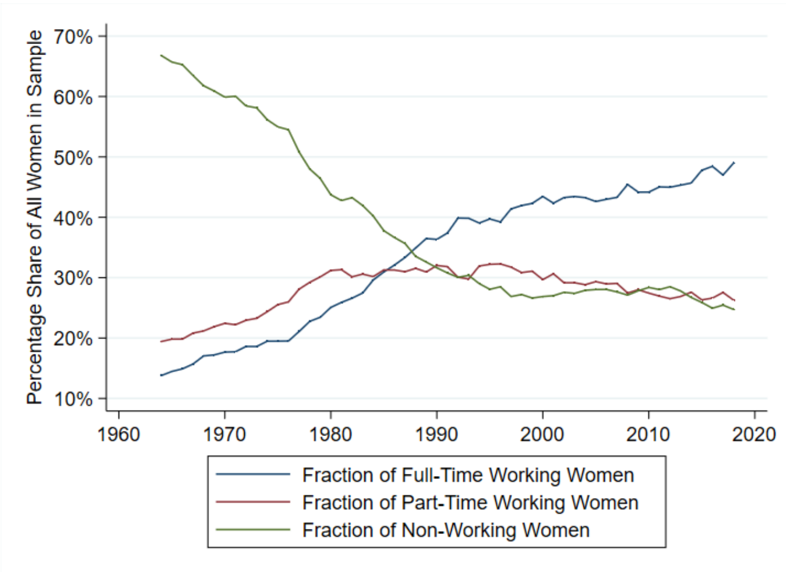

Note: Aged 25 to 55, non-Hispanic whites, living with a working partner
Source: Author's rendering of IPUMS-CPS data (2018)

It is evident that the main shifts over time are the sharp decrease of the share of women living with a working partner but not working themselves and the rising fraction of FTFY working women, which has reached an all-time maximum of 49% in 2018. The share of part-time working women on the other hand has only increased moderately since the 1960s and, after plateauing in the 1980s and 1990s, has been slightly declining since 2000.

A well-known, interesting fact is that female LFP differs greatly by educational attainment, and has always been higher for more educated women. As figure 3.3 illustrates, female LFP rates for college and non-college graduated women comove almost perfectly over time, with the exception of the current decade, which has seen a steep increase of LFP for college graduates coupled with stagnating LFP for lower educated women. The gap between the two rates has been relatively consistent at around ten percentage points until the end of the 20th century. This mirrors in part the higher expected wages for college-educated women, but also captures endogenous factors such as different fertility rates, preferences, and social attitudes between the two groups. Interestingly, figure 3.3 proves that college-educated women constitute the driving force behind the recent increase in overall female LFP depicted in 3.1.

Figure 3.3: Labor Force Participation Rate of Women by Education

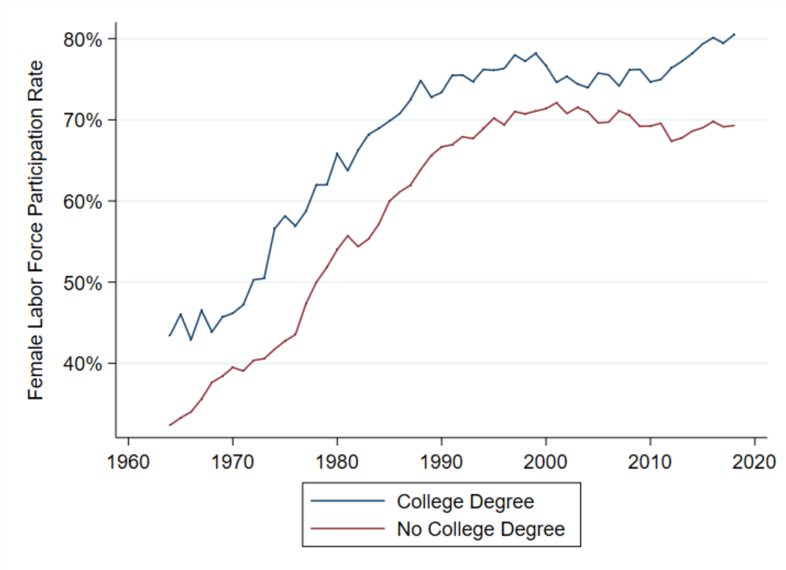

Note: Aged 25 to 55, non-Hispanic whites, living with a working partner
Source: Author's rendering of IPUMS-CPS data (2018)

At the same time, a married or cohabiting women's decision of whether to work is not only influenced by her own education, but also by her partner's education and resulting labor market opportunities. As mentioned in the beginning of this chapter, the sample of married or cohabiting couples consists of four different household types depending on each partner's education: Non-college & non-college, college & non-college, non-college & college, and college & college. Figure 3.4 shows that female LFP has increased in a similar pattern over time regardless of household type: First, female LFP grew throughout the 1960s, 1970s, and 1980s, which was then followed by almost 30 years of plateauing since around 1990. The only exception to this are college-educated women living with an also college-educated partner, whose LFP rate has risen by almost 10% since 2010. Hence, the increase of college-educated women's LFP illustrated in figure 3.3 can be attributed to women living with an equally educated partner, since the LFP rate of college-educated women living with less-educated men has been quite stable for the past 25 years.

Figure 3.4: Labor Force Participation Rate of Women by Household Type

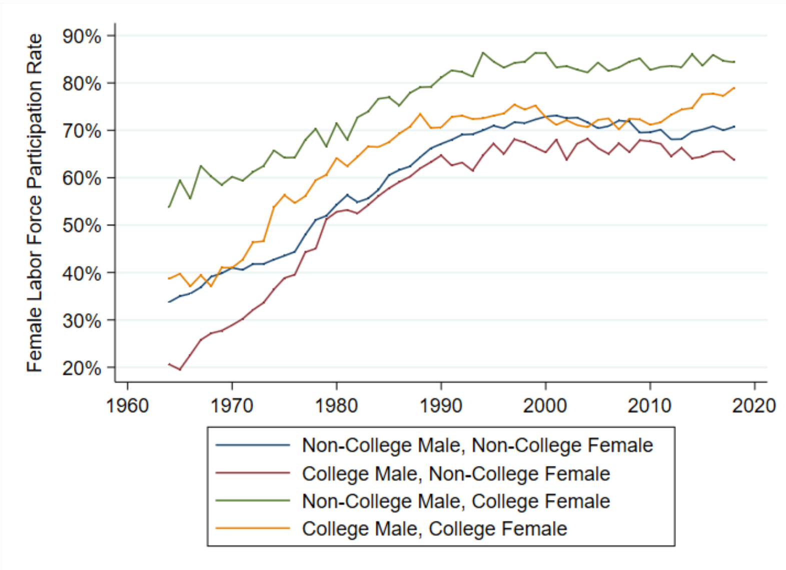

Note: Aged 25 to 55, non-Hispanic whites, living with a working partner
Source: Author's rendering of IPUMS-CPS data (2018)

Women living with a more educated partner have always shown the lowest labor force participation, while women more educated than their partner have been most likely to work at all times. This is not only in line with the findings of Gihleb and Lifshitz (2016), but also follows basic economic intuition. A rational division of labor in the household would predict the more educated (and potentially higher-earning) partner to participate in the labor market, while the less educated partner takes over household duties like chores or childcare. Interestingly, female LFP is not the same for the two types of assortatively matched couples with equally educated partners, but has been larger in most years for couples in which both partners hold a college degree. At first glance, this seems to contradict the basic economic theory of leisure as a normal good, which would predict that women in non-college & non-college households should work more due to the lower household earnings exerting a positive income effect on hours worked. However, one potential explanation for this paradox is the marketization of home production. More educated couples, through their higher household income, can afford to substitute home production with market services such as household helpers and day care, making it easier for the woman in these households to work (Cerina et al. 2017). In addition, as noted above, differences in women's preferences and fertility decisions between these two groups also have to be considered.

In conclusion, as the evidence for the sample of married/cohabiting women shows, a greater number of these women are working today than ever before in the past 50 years, and the women who participate in the labor market are working more hours than ever before. A core driving force behind these tremendous changes are the large shifts in the division of labor within the household that have been occurring since the 1960s.

Combining the CPS data with the American Heritage Time Use Study, the most extensive and harmonized set of U.S. time use data available from 1965, allows for comparing both time worked in the labor market and in the household over time. I define household work as the sum of unpaid domestic work, which includes all chores, child care, and adult care. I do not consider time spent on personal care and sleep to be part of home production, as they are individual activities mostly to ones own benefit, unlike the core activities necessary for maintaining the household that are considered as production of a public good. As the AHTUS only records the respondent's, but not their partner's education level, there is no way to distinguish time worked in the household by household type. In addition, the AHTUS is not available annually like the CPS. However, comparing labor market and household time use over the last decades for men and women by the two education levels (college and non-college) already yields highly interesting results, as can be seen in figure 3.5.

Figure 3.5: Hours per Week Spent Working in the Labor Market and in the Household by Gender and Skill

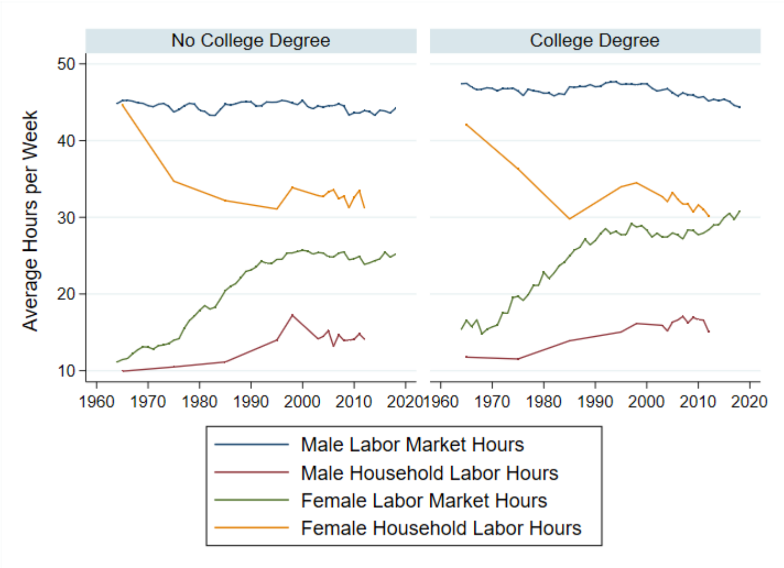

Note: Aged 25 to 55, non-Hispanic whites, couples with a working male

Source: Author's rendering of IPUMS-CPS data (2018) and IPUMS-AHTUS data (2018)

The patterns of male and female time use are very similar between the two educational groups. Women have substituted a significant share of their household hours with working in the labor market. Whereas in 1965, a woman without a college degree was on average almost spending as much time on household production as the equally educated man spent working in the labor market, female market and non-market work have converged over the years, in particular for college-educated women. Figure 3.5 indicates that for these women hours worked have recently started to exceed hours spent on home production. However, since the AHTUS data ends in 2012, this cannot be backed by the data until the AHTUS gets updated.

Conversely, the labor market hours of men are relatively stable at around 45 hours per week over the past 50 years. This stagnation persists almost regardless of education; only college-educated men display a decrease in average work hours of almost 3 hours since the turn of the century. On the contrary, men of both educational groups have increased their participation in household work by about 50% over the last decades, which together with the constant work hours indicates that men have sacrificed leisure time to support their partners in doing household work. While this seemingly contradicts the findings of Aguiar and Hurst (2007) that male market work hours have decreased significantly since the 1960s such that they enjoy more leisure today, this difference can be attributed to different sample restrictions, as Aguiar and Hurst consider all men, married and single. Nonetheless, the increase of male household labor of about 5 additional hours per week today compared to 1965 does not compensate for the large reduction in women's household production hours, which indicates that less time needs to be allocated to maintaining a household today than in the 1960s, e.g. due to the improvements in technology mentioned by Greenwood et al. (2005). Moreover, the increase of male household labor hours came to a halt at the end of the 20th century. For the latest set of time use data ranging from 2003 to 2012, men's household hours are fluctuating around their 2003 values of about 15 hours per week.

To summarize, this section has not only demonstrated the tremendous changes in female labor supply and time use over the past 54 years, but also the important role of education, both of the woman herself and of her partner, in shaping female labor decision. Thus, the next two sections are dedicated to the changing composition of households and the remarkable growth of female educational attainment.

3.2 Matching and Marriage

As indicated by figure 3.4, female LFP differs by household type, which underlines the importance of examining changes in household composition for understanding female labor supply. Besides the well-known facts of decreasing marriage and increasing divorce rates and a rising share of single households in the U.S. (Doepke & Tertilt 2016), there have also been shifts within couples over the past decades. In order to analyze these, I again use my sample of married/cohabiting couples with a working male. All descriptive statistics in this section have been computed using the ASEC household weight *ASECWTH*.

Firstly, considering the four household types in the sample based on education, it is evident that the share of couples consisting of two non-college graduates has declined continuously since the 1960s, as depicted in figure 3.6. This is mainly driven by the increasing educational attainment of both genders, which will be presented in the following section 3.3 on education. The share of couples with a more educated male is relatively stable over time at about 10%. On the other hand, a slow rise of couples where the woman has a higher education than her partner can be observed. This is due to the reversal of the gender education gap, which also will be shown in the next section. One of the most striking facts is the steep growth of the fraction of strictly college graduate couples from 5% in 1964 to 38% in 2018, making it the predominant household type for the first time in 2018.

Figure 3.6: Fractions of the Four Different Household Types by Education

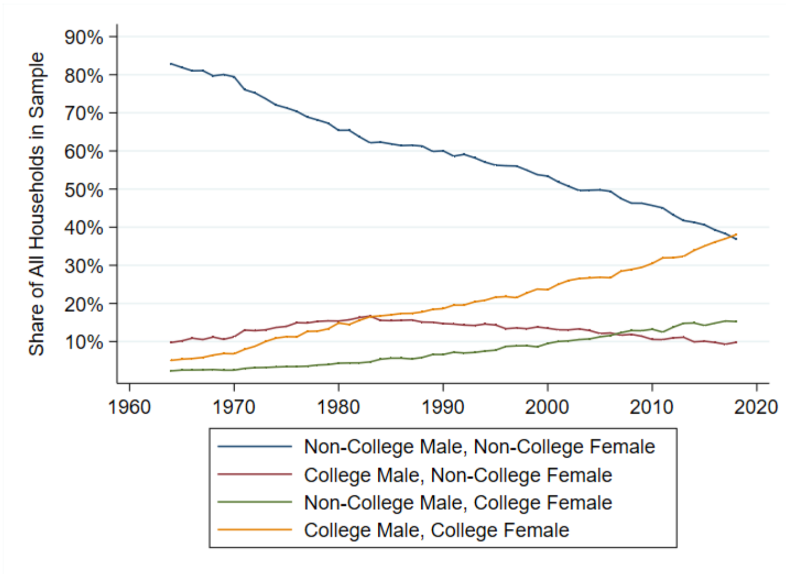

Note: Aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

In addition, figure 3.6 illustrates that in terms of education, the share of women "marrying up" by selecting a more educated partner has fallen below the share of women "marrying down" in 2007, again because of the now overall higher education of women compared to men. As the fractions of exclusively non-college and college graduate couples evolved in opposing directions, it is not evident how the share of equally matched couples with partners of equivalent education progressed over time. A closer look at the data reveals that the share of couples with similar levels of schooling (below college or college education) has decreased slightly from approximately 88% in 1964 to 75% in 2018. This does however not indicate a decrease in assortative matching over time, as it is driven by the great decline of strictly non-college graduate couples due to the general increase in education. The growth of the share of college-educated women selecting an equally educated partner has in fact far exceeded that of college-educated women living with a non-college man.

Further, alongside the decreasing share of married-couple households in the whole U.S. population documented by Doepke and Tertilt (2016), the civil status of couples has also been changing steadily over the past decades. Luckily, the CPS provides information on cohabiting couples through the category "unmarried partner" in the variable *RELATE*, albeit only from 1995 onwards. Comparing the share of legally married and unmarried cohabiting couples in my sample in figure 3.7 reveals that the share of unmarried couples has been slowly increasing during the past 20 years and is now exceeding 10% of all couples.

Figure 3.7: Fractions of Households by Civil Status

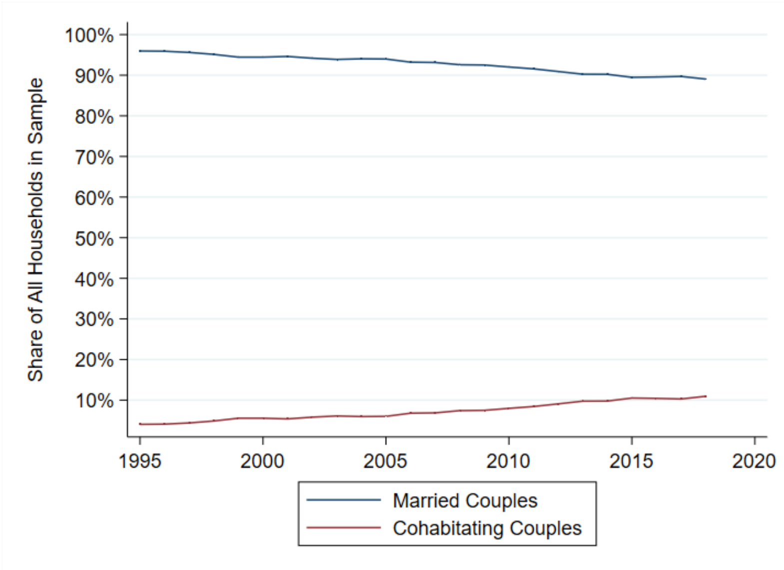

Note: Aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

Nonetheless, it has to be noted that this sample excludes subgroups of the population like younger couples (due to the age restriction of 25 to 55) and student couples (as I only consider couples with employed men) for which living together unmarried is more common. Notwithstanding, an ever-increasing share of cohabiting couples, even of only 10%, raises the question of why most research focused on household labor supply only distinguishes between single individuals and married couples. I deviate from this by including unmarried couples into my sample where possible (from 1995 to 2018), and later on conduct a counterfactual analysis using my model by restricting the sample on cohabiting and on married couples to determine whether their characteristics are different enough to considerably alter my results.

3.3 The Reversal of the Gender Education Gap

As pointed out in the previous section, the increased educational attainment of women has disrupted the composition of household types in the U.S. and consequently impacted both male and female labor supply. Before turning to the increasing monetary payoff of being well-educated in terms of earnings, I present the main trend in educational investment in the U.S. of the last decades: The reversal of the gender education gap, which describes that women have overtaken men in terms of obtained education.

In order to show the reversal of the gender education gap in my sample, I first plot the shares of college graduates by gender relative to all men or women, respectively in figure

3.8a. It can be seen that, while the college-educated shares grew tremendously for both genders over time, the increase is far steeper for women. Whereas initially the gender education gap was more than 7 percentage points in favor of men in 1964 and remained relatively constant at first, it started to narrow at the end of the 1980s. Between 2006 and 2007, both lines intersect, and since then the share of college-educated women has exceeded that of men in the sample. Moreover, it has been widening since its reversal and stands at more than 5 percentage points in the benefit of women today.

Browning et al. (2014) offer two explanations for why women may have acquired less schooling than men in the past: First, they received a lower expected return from investing into their education in the labor market due to discrimination. This was particularly the case for the 1960s, 1970s and 1980s and can explain why the gender education gap persisted during these years. Second, married women back then have also had a significantly lower expected return from schooling due to their larger role in child care, both because of social norms and for biological reasons, creating large breaks in women's employment histories. With an increasing contribution of men to child care and household work, as shown in figure 3.5, and changing social norms, this effect has weakened over time and incentivized more women to obtain tertiary education.

Figure 3.8: Fractions of College Graduates Relative to All Men or Women, Respectively, for Different Age Groups

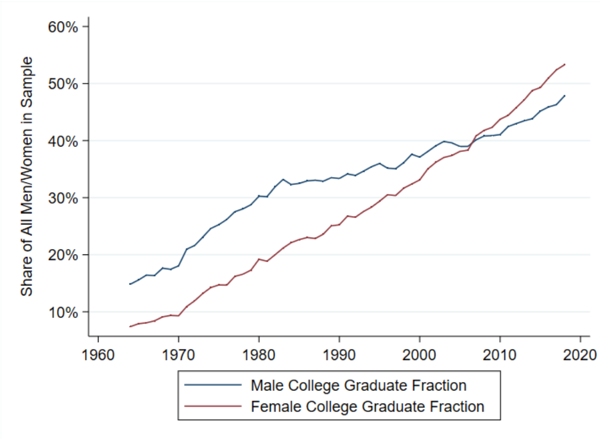

(a) Aged 25 to 55

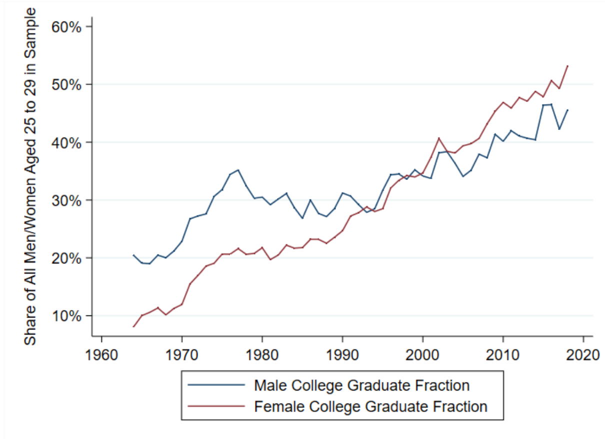

(b) Aged 25 to 29

Note: Non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

However, one has to take into consideration that the actual reversal of the gender education gap occurred prior to the intersection point in figure 3.8a. The reason for this is that considering the whole sample will only provide a lagged representation of actual college graduation trends, as it also includes preceding cohorts up to the maximum restricted age of 55. Thus, changes in educational attainment take several years to materialize on the level of the whole population. One way to overcome this, as done by Heathcote et al.

(2010), is to focus on a sample of 25-29-year-olds. They find that the gender education gap reversed in the late 1980s among 25-29-year-olds Heathcote et al. (2010), whereas in my sample the reversal is recognizable more than a decade later in the early 2000s as depicted in figure 3.8b. This difference is mostly due to my sample having the minimum age of 25 also imposed on women, while Heathcote et al. use a CPS sample of married households in which only the husband need to be 25 to 59 years old. By only restricting the husband's age, the sample of Heathcote et al. includes a large number of women below the age of 25 due to men typically marrying younger women. This leads to an underrepresentation of young women in my sample compared to theirs, causing the reversal of the gender education gap to occur delayed.

Nonetheless, figure 3.8b shows that for younger people, the college-educated share of women has been comparable to that of men throughout the 1990s. Most importantly, the persistent female education advantage in the 25-29-year-olds sample since the 2000s has important implications for the complete sample. Namely, as older cohorts with a larger male college share drop out, the education gap continues to widen, which is exactly what can be seen in figure 3.8a over the last years.

The ever-increasing education of women directly affects their labor supply. As women invest more time and money into their human capital accumulation, it becomes more expensive for highly-educated women to stay out of the labor force, for example in order to raise children. This is due to increased opportunity costs of not working faced by well-educated and potentially high-earning women, and has been shown empirically by Olivetti (2006). Hence, one would expect a larger share of college-educated women among FTFY compared to part-time and non-working women, and figure 3.9 confirms this.

Figure 3.9: Fractions of Female College Graduates by Work Status, Respectively

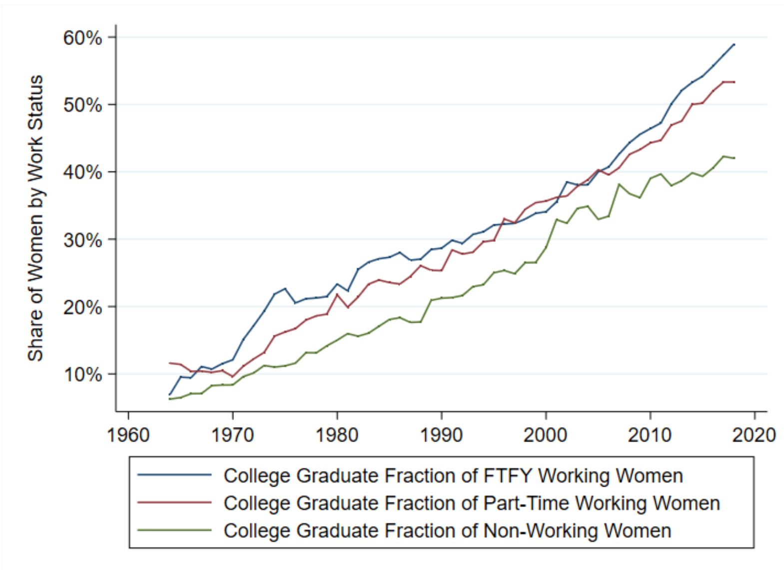

Note: Aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

To conclude the section on education and as a follow-up on the increasing share of co-habiting couples shown in figure 3.7, figure 3.10 provides insight into the composition of married and unmarried couples in 2018 by both partners' education.

Figure 3.10: Fractions of the Four Different Household Type by Education over Civil Status in 2018

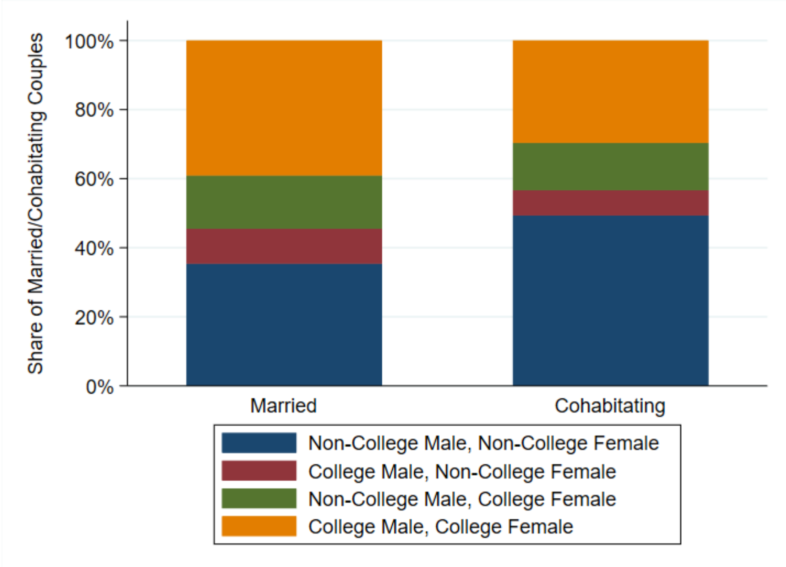

Note: Aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

It is evident that the skill distribution within couples differs greatly by their civil status. While the shares of unequally-skilled couples are similar, strictly non-college educated couples are far more prevalent in the cohabiting sample, whereas the college-college household type fraction is greater among married couples. One possible reason for this could be a preference towards marriage among the more educated. Furthermore, my sample excludes a large number of young cohabiting college & college couples in which the male is not working, for example due to attending graduate school, or in which the woman has a college degree but is younger than 25 years old. Nevertheless, the different household composition by civil status combined with an increasing share of unmarried couples is a strong argument for taking these trends into consideration when analyzing household behavior, as I will demonstrate later using my model.

3.4 Wages and Skill-Biased Technological Change

Before examining the wage differences between non-college and college graduates, it is important to note the difference in wages between genders. One way to illustrate this is by plotting the female relative wage, which is defined as the ratio of the average female to average male hourly wage. Since the share of part-time workers is significantly higher among women, I restrict the analysis to full-time full-year (FTFY) workers, again defined as working more than 50 weeks a year for at least 35 hours per week. Otherwise the part-time penalty, meaning that not working full-time results in a lower hourly wage, would considerably lower the female-to-male wage ratio. Figure 3.11 presents how female relative wages of FTFY workers evolved over time and shows that female relative wages have increased immensely over time.

Figure 3.11: Relative Female Wage (Female to Male Average Hourly Wage Ratio) for FTFY Workers

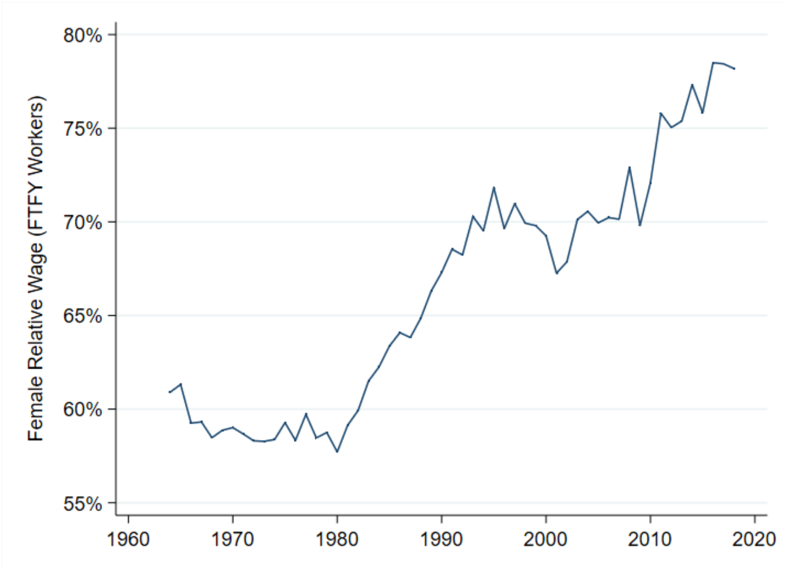

Note: Aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

Most of the growth can be attributed to the 1980s and the past decade. Overall, female wages increased from 61% in 1964 to more than 78% of male wages in 2018. The reasons for today's gender wage gap of about 22% in this sample are multifarious and include lower labor market experience due to childbirth breaks, different occupational choices and working less hours.

Furthermore, the previous section has shown that women are obtaining more education than men, and this education pays off increasingly well due to skill-biased technological change, which describes that the productivity of high-skilled labor is increasing overproportionally compared to that of unskilled labor. One way this can be observed in the data is through the growing wage ratio of skilled to unskilled workers called the skill premium. Most commonly used in this context is the ratio of college to non-college graduates, the college wage premium. Heathcote et al. (2010) for example find that the male college wage premium declined during the 1970s, but has been growing since then. Katz and Autor (1999) come to the same result, namely that the wage difference between college and high school graduates has increased sharply after an initial drop in the 1970s. Acemoglu (1998) explains the decreasing skill premium in the 1970s with technologies being fixed in the short-run, while the greater supply of skilled labor lowered the college wage premium.

Liu (2017) offers a concrete explanation for the larger supply of skilled workers in the 1970s: The U.S. government raising financial aid for college students through the Basic Education Opportunity Grant Program. In his dissertation, Liu extends the model of Caselli and Coleman (2006) with agents' choosing whether to obtain a college education or not. One of the findings of his model is a negative relationship between the (financial) availability of college and the skill premium. As financial support for students was raised in the 1970s, more people could afford higher education in the U.S. The huge scale of this grant program is reflected in it constituting 29% of total tuition fee revenue in the U.S. in 1980 (McPherson & Schapiro 1991). This increased the supply of skilled labor and reduced the skill-premium in the short-run, until skill-biased technological change increased the demand for skilled labor during the 1980s sufficiently to raise the skill premium again (Liu 2017).

My sample shows the same patterns, as presented in figure 3.12. In addition, displaying the overall and the female college wage premium reveals that the peak of the total premium in the late 2000s was driven by males. Interestingly, the female college wage premium exceeds the male college premium in 2018 and stands at an all-time high of 1.68. This implies that today, obtaining a college degree offers on average a stronger relative wage increase for women than it does for men.

Figure 3.12: College Premium (College to Non-College Average Hourly Wage Ratio) for FTFY Workers, Overall and by Gender

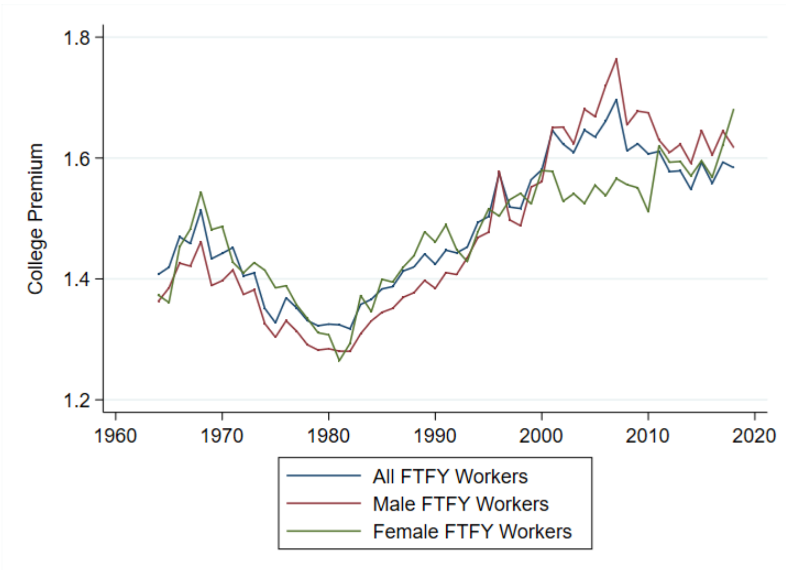

Note: Aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

3.5 The Rise of Earnings Inequality

Concluding the chapter on empirical facts, I show how wage and earnings inequality between different subgroups of the population developed over time. This draws on the trends in female labor supply, matching, education, gender pay differences, and college wage premium presented in the previous sections, since they all affect inequality between and within households. All monetary variables depicted here have been adjusted for inflation using CPI data (United States Bureau of Labor Statistics 2019) and are expressed in constant 2010 US-Dollars. Figure 3.11 has already demonstrated the differences in hourly wages between genders, yet there are also large wage differences within genders. Figure 3.13a represents the development of wage inequality among FTFY working men and women, measured as the log ratio of the 90th percentile to the 10th percentile hourly wage, called the 90-10 ratio. Taking the log of the 90-10 ratio means that it can be interpreted as the difference between the 90th and 10th percentile log wage. Another popular measure of inequality is the log 90-50 ratio comparing the 90th percentile log wage to the median log wage, presented in figure 3.13b.

Figure 3.13: Hourly Wage Inequality of FTFY Workers by Gender

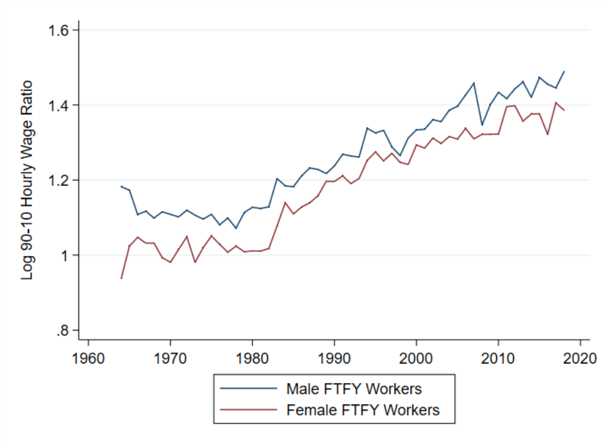

(a) Log 90-10 Ratio

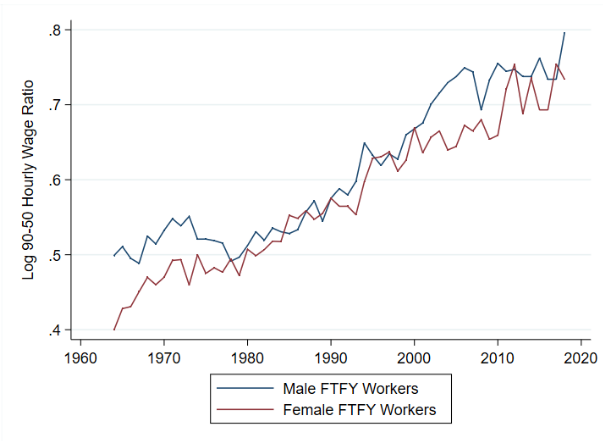

(b) Log 90-50 Ratio

Note: Aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

While the 90-10 ratio is obviously greater than the 90-50 ratio, the pattern of both is quite comparable. Notably, the male wage dispersion has been higher than the female one in most years, especially when comparing the 90th to the 10th percentile log wage. One of the reasons for this is the great gender disparity at top-paying jobs, as supported by Albanesi and Olivetti (2009) who find that the gender wage gap is greatest within management and sales occupations. This causes the male 90th percentile wage to be considerably higher than the female one at all times, while the gender gap is smaller at the bottom of the wage distribution. In addition, the big gaps between the male and female ratios in the 2000s can be attributed to the much higher college wage premium for men in this period depicted in figure 3.12.

Besides hourly wages, another important factor for earnings inequality is the amount of hours worked, as mentioned by Heathcote et al. (2010). Figure 3.14 illustrates the variance of log annual hours worked within both genders and shows that male annual hours variance has been fairly constant over the decades. This is in support of the stagnating weekly market work hours of men found in 3.5, as they have always been mostly working FTFY, regardless of their education and spousal matching. On the contrary, the variance of log female annual hours worked has decreased significantly, which can be linked to the changing composition of the female work force presented in 3.2. As the share of FTFY working women typically working 40 hours a week grows, the variance of their annual work hours converges to that of men.

Figure 3.14: Variance of Log Annual Hours Worked by Gender

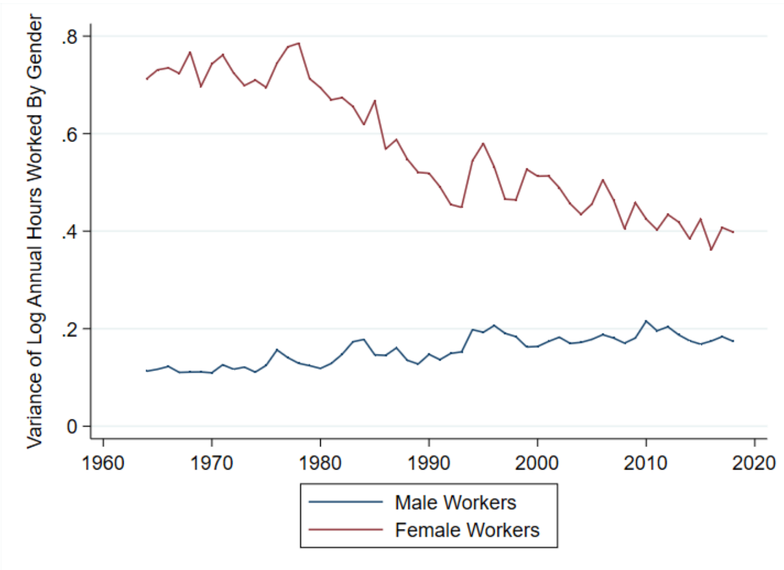

Note: FTFY and part-time workers, aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

Combining the trends in wage and hours dispersion, there are two different levels of earnings inequality to be considered in my sample: Intra-household inequality between partners and inter-household inequality between the different types of households. Firstly, within-household inequality can be depicted through the ratio of a woman's earnings to her partner's. As one would expect, this ratio varies greatly depending on both partners' education levels. Figure 3.15 shows the development of the average ratio over time for the four household types in my sample, with a ratio of one meaning that both partners have the same labor income. First, the ratio is greater than one for couples in which the woman is more educated, proving that the higher education compared to her partner reflects in larger relative earnings. The enormous fluctuations of this curve in the 1960s and 1970s can be explained by the small sample size, as non-college & college households constituted less than 5% of the sample in these years (see figure 3.6). In contrast, the relative earnings ratio is lowest for households with a more educated male. One interesting fact to note is that the distance of both non-college & college and college & non-college relative earnings from one are practically the same in 2018, indicating that education relative to the partner is almost perfectly mirrored in earnings today and documenting the narrowing of the gender pay gap (at least for the sample of couples living together).

Figure 3.15: Average Female to Male Annual Earnings Ratio by Household Type

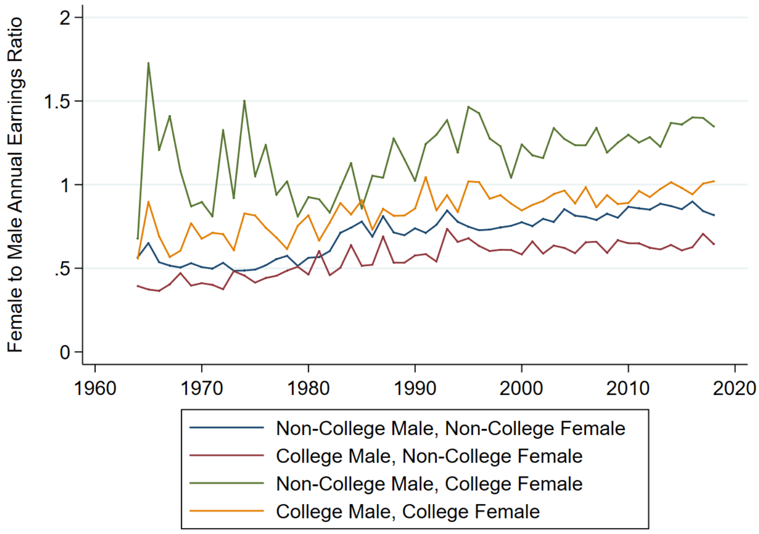

Note: FTFY and part-time workers, aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

For couples with equally educated partners, a ratio of one would indicate perfect intra-household earnings equality, and both of these household types (strictly non-college and strictly college) have been converging towards that. The ratio for couples of two college graduates has even slightly surpassed one in 2017 and 2018, proving that college-educated women earn more on average than their also college-educated partner. These patterns therefore support an overall reduction of intra-household inequality over time, with potential causes being the narrowing gender wage gap (figure 3.11) as well as the increasing female market hours (figure 3.5) and changing work status of women towards FTFY jobs (3.2).

In contrast, calculating inter-household inequality reveals a tremendous dispersion of earnings between households. Figure 3.16 considers household earnings as the sum of both partners' annual earnings for the different household types. It can be seen that household earnings were relatively equal in the 1960s for the three household types in which at least one partner has a college degree. Over time, household earnings dispersed, in particular for the college & college household type, whose average annual earnings were almost double of the non-college & non-college type in the 2000s.

Figure 3.16: Average Annual Household Earnings by Household Type in Constant 2010 USD

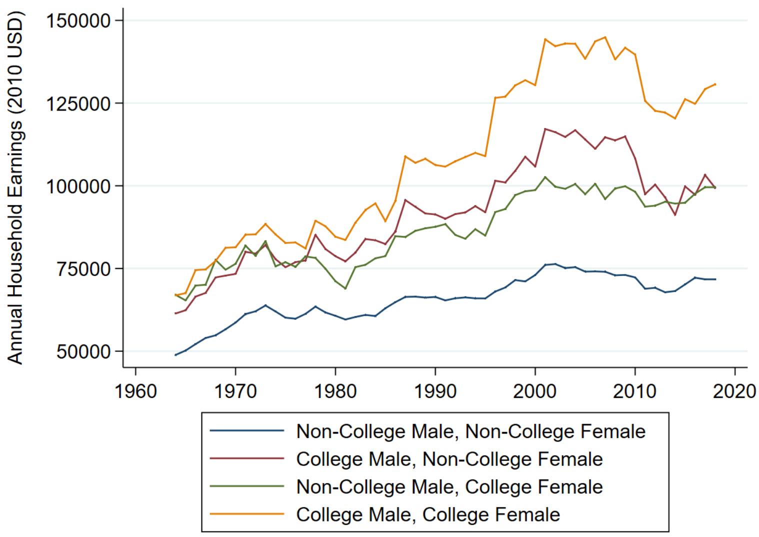

Note: FTFY and part-time workers, aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

This pattern translates into an increasing Gini coefficient over time, as presented in figure 3.17a. Notably, the 1980s and late 1990s saw the largest inter-household inequality increase. Quite surprisingly, inequality has been reduced in my sample over the last 10 years. This is indicated by a decrease of the Gini coefficient following the Great Recession in the late 2000s, but also shows in declining average household earnings of households with a college-educated male in figure 3.16. Furthermore, the Gini coefficient of individual earnings depicted displays the same pattern.

Figure 3.17: Gini Coefficients over Time

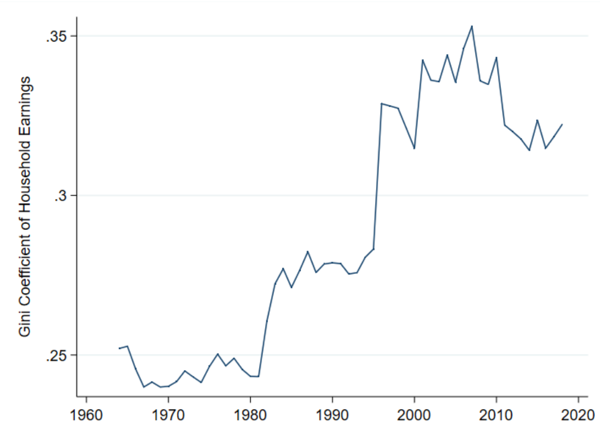

(a) Gini Coefficient of Average Annual Household Earnings by Household Type

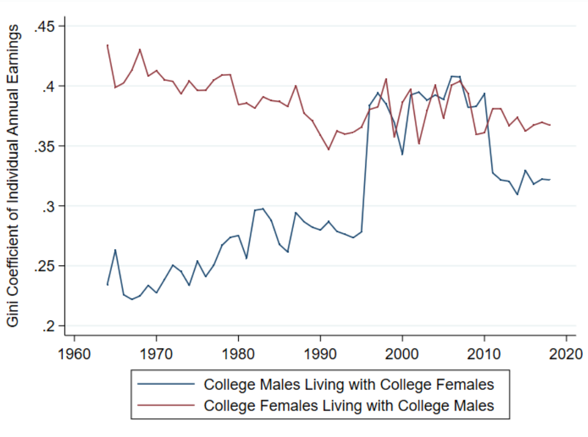

(b) Gini Coefficient of Average Annual Individual Earnings for Men and Women in Households of Two College Graduates

Note: FTFY and part-time workers, aged 25 to 55, non-Hispanic whites, couples with a working male
Source: Author's rendering of IPUMS-CPS data (2018)

To identify the driver behind this drop in inequality, I plot the Gini coefficients of the average individual earnings of men and women in the college & college household type in figure 3.17b. The result demonstrates that men's Gini coefficient dropped from 0.41 in 2006 to 0.31 in 2014, while earnings inequality among women in this household type have been decreasing over the past decades, as described by Card and Hyslop (2018). This indicates that the Great Depression caused men's earnings to become more equal in my sample of working men, in particular affecting college-educate men who have both reduced their weekly work hours (figure 3.5) and faced a lower wage premium over non-college educated men (figure 3.12) in recent years.

In conclusion, this chapter has illustrated an overview of the great shifts in female labor supply, household composition, education, wages and earnings over the past 54 years in the U.S. Taking all these changes into account, I will now introduce a heterogeneous household model in order to quantitatively examine how they have affected female labor supply and earnings inequality.

4 Introducing a Structural Model

4.1 Model Setup

In order to determine the effects of skill-biased technological change on consumption inequality, I introduce a static partial equilibrium model of a production economy with endogenous wages and heterogeneous two-person households. In this model, households maximize the unitary CES utility function from Heathcote et al. (2010), extended with time spent on household production for each partner. This type of utility function assumes that partners have the same preferences and jointly maximize household utility, with both partners having individual labor supply and wages but consuming goods as a household.

On the contrary, in a collective model both partners maximize their own utility, and household utility is then formed as a weighted sum of the individual spousal utilities. The main reason for why I set up a unitary household model instead of a collective one is that it does not require the specification of how partners' utilities are weighted. This simplifies the analysis, since there is no household bargaining in my model such that the bargaining power of each partner (and hence their utility weight) is ambiguous. In any case, simply assuming an equal distribution of power within the household and weighting both utilities with 0.5 would mean that the collective model effectively becomes unitary. Nonetheless, it has to be noted that partners having the same preferences is a significant abstraction from reality, where partners most likely have individual preferences (Chiappori 1992).

Another approach employed in the literature (Bick 2016) is to abstract from male labor supply and just solve the female partner's optimization problem. This however is not an option for my analysis, since gender differences in labor supply are a crucial part of this thesis. Moreover, incorporating individual labor supplies while leaving consumption on the household level is sufficient for the purpose of this thesis, where the focus is on labor supply and income inequalities. Analyzing consumption inequality would require a more complex collective model with public and private consumption goods and some kind of sharing rule.

Individuals are heterogeneous in terms of their gender and skill level, which is assigned exogenously. In the model economy, the number of male and female individuals is equal and they are all living together with one individual of the other gender, meaning there are no singles. Agents can be of two types of skill-level, high- and low-skilled, which correspond to high school graduates (including those with some college education) and college graduates holding at least an undergraduate degree. While Caselli and Coleman (2006) note that there are other possible skill level thresholds, like for instance completed primary education, using completed college education is the most appropriate for my analysis as it is motivated by the increase of female college graduation rates observed in the data (figure 3.8). Hence, there are four different types of individuals in my model: skilled and unskilled males and skilled and unskilled females.

The household utility function is a version of the one presented in Heathcote et al. (2010), extended with household production time:

$$U(c_{i,j}, l_{i,j,m}, l_{i,j,f}) = \frac{c_{i,j}^{1-\eta}}{1-\eta} + \psi \frac{(1 - l_{i,j,m} - \overline{h_{i,j,m}})^{1-\sigma}}{1-\sigma} + \psi \frac{(1 - l_{i,j,f} - \overline{h_{i,j,f}})^{1-\sigma}}{1-\sigma} \tag{4.1}$$

In the utility function, $i = \{u, s\}$ and $j = \{u, s\}$ denote the skill levels of the male and female partner (unskilled or skilled). $c_{i,j}$ is the joint consumption of the household, while $l_{i,j,m}$ and $l_{i,j,f}$ denote male and female labor supply respectively. Hence, both consumption and individual labor supplies depend on the skill level of both partners. In particular, it is necessary that individual labor supply not only depends on one's own, but also on the partner's skill level in order to more accurately depict the labor decisions within the household. Time spent on home production is given by $\overline{h_{i,j,m}}$ and $\overline{h_{i,j,f}}$ for each partner and is treated as a gender- and skill-specific fixed time cost parameter rather than a choice variable in order to simplify the analysis. This means that home production is part of the utility function and time constraint, but cannot be varied by households. Making this assumption captures the fact that some amount of time has to be dedicated to necessary tasks like doing chores, shopping or taking care of children. In the following analysis, household work hours will be reduced over time according to the data in order to model increased household productivity. Finally, the parameters η, σ and ψ denote the coefficient of relative risk aversion (CRRA), the substitution parameter for leisure (the elasticity of substitution between male and female leisure is therefore $\frac{1}{\sigma}$) and the utility weight of leisure.

Every individual is endowed with one unit of time which – after working the fixed amount of time $\overline{h_{i,j,m/f}}$ in the household – they can allocate between working in the market ($l_{i,j,m/f}$) and enjoying leisure ($d_{i,j,m/f}$). All three time uses depend on the skill distribution within the household. Hence, the household time constraint is given by:

$$d_{i,j,m} + d_{i,j,f} = 2 - l_{i,j,m} - l_{i,j,f} - \overline{h_{i,j,m}} - \overline{h_{i,j,f}} \tag{4.2}$$

The household's consumption expenditure $c_{i,j}$, with the price of the consumption good normalized to 1, is financed through the individual labor incomes from allocating time to market work, which gives the following household budget constraint:

$$c_{i,j} = w_i l_{i,j,m} + w_j l_{i,j,f} \tag{4.3}$$

The household's consumption is therefore determined by how much labor it supplies. Note the absence of the gender subscript on the wages, which resembles the assumption that firms do not differentiate between genders and only hire with regards to skill level. This implies that there is no gender discrimination in my model, meaning that equally skilled people get paid identical wages. While this is a broad simplification of reality in which employers might have a gender bias in hiring, especially in the earliest time periods of the model in the 1960s, the focus of this thesis on skill-bias and the reversal of the gender education gap justifies this abstraction. Hence, w_i denotes the male and w_j the female wage.

Since individuals are heterogeneous across two dimensions (gender and skill), the economy consists of four different household types – skilled-skilled, skilled-unskilled, unskilled-skilled, and unskilled-unskilled, the first skill level referring to the male – with varying labor supplies and consumption. In order to abstract from modeling a marriage market, I impose the matching of couples and their educational attainment exogenously based on the empirical distribution of household types, following the definition used in section 3.2.

Another important part of my analysis is the distinction between married couples and couples living in cohabitation. In the literature on household labor supply, households are mostly defined as married couples. However, as shown by Doepke and Tertilt (2016) and illustrated in figure 3.7, declining marriage and increasing divorce rates together with a growing share of unmarried couples in the U.S. over recent decades call for an updated definition that better reflects the reality of how couples in the U.S. are living today. Alongside following the standard approach of focusing on married couples only, I will therefore also extend this definition to include unmarried couples living in cohabitation, which has become an increasingly widespread household type in the U.S. (figure 3.7). In my analysis, I am comparing the results obtained by using both the strictly married as well as the married or living in cohabitation definition and discuss the differences.

On the production side of the economy, there is a continuum of competitive firms of measure 1 that demand the labor supplied by the households, turning the model into a production economy in partial equilibrium with endogenous wages. The production function of the representative firm is taken from Caselli and Coleman (2006) and extended with gender:

$$y = k^{\alpha}\left[(A_u L_u)^{\theta} + (A_s L_s)^{\theta}\right]^{\frac{1-\alpha}{\theta}}, \tag{4.4}$$

$$\text{where } L_u = \phi_{u,u}((l_{u,u,m} + l_{u,u,f}) + \phi_{u,s} l_{u,s,m} + \phi_{s,u} l_{s,u,f} \tag{4.5}$$

$$\text{and } L_s = \phi_{s,s}(l_{s,s,m} + l_{s,s,f}) + \phi_{u,s,f} l_{u,s,f} + \phi_{s,u} l_{s,u,m} \tag{4.6}$$

As explained above, the key assumption is that firms do not regard gender in their hiring decisions, only skill level. Therefore, while unskilled labor L_u and skilled labor L_s are composed of weighted male and female labor supplies by household type, the firm only hires labor with respect to L_u and L_s and does not observe their composition. An interpretation of this assumption would be that firms are receiving anonymized applications that only reveal the skill level (the education) of the applicant. L_u and L_s represent the link between the labor supplied by the households and the labor employed by the firm. As I will explain in the next chapter, the weights $\phi_{i,j}$ are obtained from the data.

Furthermore, capital k is exogenous, but is not important for the equilibrium as shown in the next section. The ratio $\frac{A_s}{A_u}$ models the relative productivity of skilled to unskilled labor, where an increase of this ratio over time indicates skill-biased technological change. The elasticity of substitution between high- and low-skilled labor is $\frac{1}{1-\theta}$, so $\theta \to 0$ resembles the Cobb-Douglas case. α denotes the weight of capital in production.

Like in the model of Caselli and Coleman, firms determine the level of skilled-labor-intensitivity of their production technology by choosing A_s and A_u from a set of available technologies, which is defined as:

$$(A_s)^\omega + \gamma(A_u)^\omega \leq B \tag{4.7}$$

ω, γ and B are all strictly positive, exogenous parameters the firm takes as given. ω and γ determine the trade-off between the two labor productivities, while B fixes the upper ceiling of all possible technology sets.

4.2 Equilibrium

The model equilibrium is characterized by labor market clearance, households maximizing their utility and firms maximizing profits. Solving the model therefore requires deriving the optimality conditions for both the households and the firm. Starting with the first, households maximize their utility (4.1) subject to the budget (4.3) and time (4.2) constraints by choosing consumption and labor supplies, which yields:

$$\frac{w_i}{w_j} = \left(\frac{1 - l_{i,j,m} - \overline{h_{i,j,m}}}{1 - l_{i,j,f} - \overline{h_{i,j,f}}}\right)^{-\sigma} \tag{4.8}$$

The detailed derivation of the optimality condition (4.8) is explained in the Mathematical Appendix B.1. As the optimality condition shows, the ratio of the male wage w_i to the female wage w_j within each household, both depending on the respective skill level, is a function of both partners' labor supplies $l_{i,j,m/f}$, household production hours $\overline{h_{i,j,m/f}}$ and the substitution parameter between male and female leisure σ.

Turning to the other agent, the representative firm maximizes profits subject to the production function (4.4) and the technology constraint (4.7). Its choice variables consist of the productivities A_u and A_s, the capital stock K, and the labor inputs L_u and L_s.

Hence, wages are endogenously determined. The firm has three optimality conditions: One relating labor supplies, one relating productivities, and one for capital. As in Caselli and Coleman (2006), the first-order condition for capital is irrevelant for my analysis and will not be considered. Therefore, I will only derive the other two optimality conditions.

Starting with labor supply, combining the firm's first-order conditions for L_s and L_u yields the ratio of the skilled to the unskilled wage, the skill premium:

$$\frac{w_s}{w_u} = \left(\frac{A_s}{A_u}\right)^{\theta} \left(\frac{L_s}{L_u}\right)^{\theta-1} \tag{4.9}$$

This is the first firm optimality condition. The second, capturing the firm's trade-off between the productivity levels A_u and A_s, is:

$$\left(\frac{A_s}{A_u}\right)^{\omega-\theta} = \gamma \left(\frac{L_s}{L_u}\right)^{\theta} \tag{4.10}$$

Again, the detailed derivations of both conditions are documented in Appendix B.2. Given equations (4.9) and (4.10), solving the model is then simply a matter of combining the household and firm optimality conditions by aggregating the household labor supplies and assuming labor market clearance.

First, as my model is aimed at predicting female labor supplies, equation (4.8) needs to be solved for female labor supply $l_{i,j,f}$ for each of the four different household types (skilled-skilled, skilled-unskilled, unskilled-skilled, and unskilled-unskilled). As shown in Appendix B.1, this yields:

$$l_{i,j,f} = 1 - \overline{h_{i,j,f}} - (1 - l_{i,j,m} - \overline{h_{i,j,m}}) \left(\frac{w_i}{w_j}\right)^{\frac{1}{\sigma}} \tag{4.11}$$

Since I assume that wages only depend on skill level, solving this for the two equally skilled households s-s and u-u is straight-forward. This is because in these households, the partners' wage ratio $\frac{w_i}{w_j}$ is equal to 1, as both can expect to be paid the same hourly wage in the market. Hence, equation (4.11) simplifies to:

$$l_{u,u,f} = l_{u,u,m} + \overline{h_{u,u,m}} - \overline{h_{u,u,f}} \tag{4.12}$$

for the unskilled-unskilled and

$$l_{s,s,f} = l_{s,s,m} + \overline{h_{s,s,m}} - \overline{h_{s,s,f}} \tag{4.13}$$

for the skilled-skilled household. Due to the identical wages, both partners will enjoy the same amount of leisure in these household types, meaning female labor market hours will be pinned down by male labor supply as well both partners' home production hours.

For households with heterogeneously skilled partners earning different wages, female labor supply will also depend on the wage ratio (the skill premium). For households with a higher skilled male, this gives female labor supply as:

$$l_{s,u,f} = 1 - \overline{h_{s,u,f}} - (1 - l_{s,u,m} - \overline{h_{s,u,m}}) \left(\frac{w_s}{w_u}\right)^{\frac{1}{\sigma}} \tag{4.14}$$

For couples consisting of a skilled female living with an unskilled male, female labor supply is determined by:

$$l_{u,s,f} = 1 - \overline{h_{u,s,f}} - (1 - l_{u,s,m} - \overline{h_{u,s,m}}) \left(\frac{w_u}{w_s}\right)^{\frac{1}{\sigma}} \tag{4.15}$$

The next chapter is dedicated to illustrating how the model is solved in MATLAB, using the same CPS (Flood et al. 2018) and AHTUS (Fisher et al. 2018) data presented in chapter 3. I will explain in detail the aggregation of individual labor supplies, the parameterization of the model, the wage-setting mechanism, and how the representative firm chooses its optimal production technology.

5 Solving the Model

5.1 Prerequisites

The previous chapter outlined the model and showed the derivation of its equilibrium. Now, the next step in order to solve the model is to link the household and the firm problem. For this, the eight individual labor supplies (male and female for each of the four household types) need to be aggregated and expressed as functions of only the wage ratio $\frac{w_s}{w_u}$. The goal is to obtain $L_s(\frac{w_s}{w_u})$ and $L_u(\frac{w_s}{w_u})$, which can then be combined with the firm's optimality condition (5.1).

To do so, I use the equations (4.12) to (4.15) derived in the previous chapter. I then obtain the male labor supplies, male household hours and female household hours for each household type from the CPS and AHTUS data by computing the average hours per week within each household type per year. I assume each individual to have 100 hours of productive time per week that can divided between working in the market, working in the household, and enjoying leisure. Hence, I divide each of these values by 100 to obtain $l_{i,j,m}$, $\overline{h_{i,j,m}}$ and $\overline{h_{i,j,f}}$ for each household type. Unfortunately, as the AHTUS data only covers the years 1965, 1975, 1985, 1995, 1998 and 2003 to 2012, my analysis is limited to these years, since the AHTUS data provides the time spent on household production. A detailed explanation of the sample selection and all data manipulations can be found in the Data Appendix A.

As equations (4.12) and (4.13) show, this already explicitly determines two of the four female labor supplies, $l_{u,u,f}$ and $l_{s,s,f}$, while the remaining two can both be expressed as functions of the skill premium, $l_{s,u,f}(\frac{w_s}{w_u})$ and $l_{u,s,f}(\frac{w_s}{w_u})$ (the latter by taking the inverse of the wage ratio $\frac{w_u}{w_s}$). The fractions of the four household types – $\phi_{u,u}$, $\phi_{u,s}$, $\phi_{s,u}$ and $\phi_{s,s}$ – are also taken from the data for each year, and are the same as depicted in figure 3.6.

Thus, following equations (4.5) and (4.6), I obtain the two aggregates of unskilled labor $L_u(\frac{w_s}{w_u})$ and skilled labor $L_s(\frac{w_s}{w_u})$ as functions of the skill premium. These can be plugged into the firm's optimality condition:

$$\frac{w_s}{w_u} = \left(\frac{A_s}{A_u}\right)^{\theta} \left(\frac{L_s(\frac{w_s}{w_u})}{L_u(\frac{w_s}{w_u})}\right)^{\theta-1} \tag{5.1}$$

5.2 Parameterization

The next step is to set the four parameters required for solving the model: The substitution parameter between male and female leisure in the utility function σ, the substitution parameter between skilled and unskilled labor in the production function θ, and ω and γ determining the trade-off between A_s and A_u that the firm faces when choosing a production technology. As I will explain in the following, I obtain all four parameters from the literature.

Regarding σ, I follow Heathcote et al. (2010) who set it equal to 3 and argue that this yields reasonable Frisch elasticities of 0.48 for men and 1.46 for women. Using $\sigma = 3$ implies an elasticity of substitution between male and female leisure of $\frac{1}{3}$. This is consistent with Attanasio and Weber (1995), who find elasticities of 0.34 and 0.48 depending on the sample. Other researchers find larger estimates such as 0.66 (Domeij & Flodén 2006), 0.7 (Pistaferri 2003) and 0.67 (Attanasio, Low, & Sánchez-Marcos 2008), which imply a lower CRRA coefficient. However, changing σ exerts only a very marginal effect on the model results, hence I maintain the value of 3 from Heathcote et al. (2010).

Yet, the model is much more sensitive to changes in θ. Here I take the value from Caselli and Coleman (2006), as my model is an extended version of theirs. Caselli and Coleman set $\frac{1}{1-\theta}$, the elasticity of substitution between low- and high-skilled labor, to 1.4, implying $\theta = 0.286$. In doing so, they approximate the value set by Katz and Murphy (1992) of $\frac{1}{1-\theta} = 1.41$, which is also used by Heathcote et al. (2010).

For ω, I also follow Caselli and Coleman and set $\omega = 0.41$. They state that $\omega > \theta/(1-\theta)$ is a necessary condition for a symmetric equilibrium, meaning that the representative firm will never set $A_u = 0$ or $A_s = 0$ and employ only one of the two types of labor. Rearranging this inequality condition, one obtains that $\omega > \theta$, which is fulfilled by setting $\theta = 0.286$ and $\omega = 0.41$.

Caselli and Coleman (2006) obtain the final parameter γ by taking the natural logarithm of the second firm optimality condition (4.10) and then regressing $log(A_s/A_u)$ on $log(L_s/L_u)$. However, this requires knowledge of A_s/A_u. Caselli and Coleman obtain this ratio through the firm's optimality condition with respect to labor supplies:

$$\frac{w_s}{w_u} = \left(\frac{A_s}{A_u}\right)^{\theta} \left(\frac{L_s}{L_u}\right)^{\theta-1} \tag{5.2}$$

By taking both labor supplies and the wage ratio from the data, they determine A_s/A_s through this equation. Nonetheless, this is not applicable to my model if I want to maintain endogenous wages. Therefore, I take γ from Liu (2017), who extends the same model by Caselli and Coleman with an endogenous education choice and normalizes γ to 1. Table 5.1 summarizes the external parameter values.

Table 5.1: External Parameter Values

Parameter	Description	Source	Value
σ	Substitution between male and female leisure	Heathcote et al. 2010	3
θ	Substitution between skilled and unskilled labor	Caselli and Coleman 2006	0.286
ω	Technology trade-off parameter	Caselli and Coleman 2006	0.41
γ	Technology trade-off parameter	Liu 2017	1

5.3 Solution Method

To solve the model in MATLAB, I first normalize both the wage w_u and the productivity A_u of low-skilled workers to 1. This leaves two variables for the firm to set in order to maximize its profit in each period: A_s and w_s. The main difference between the model of Caselli and Coleman and my extended version is that there are no households in their model, meaning both the labor supplies and the wage ratio are derived from the data. In my model on the other hand, the firm practically solves two problems by both deciding how much labor to employ at what wages and which production technology to choose.

The way I model this is as a sequential decision process solved through a guess & iterate algorithm in MATLAB. First, the firm (or rather I) takes an initial guess of the production technology ratio A_s/A_u, which due to normalizing A_u consists of setting some A_s. Given this production technology, the firm then solves its first optimality condition:

$$\frac{w_s}{w_u} = \left(\frac{A_s}{A_u}\right)^{\theta} \left(\frac{L_s(\frac{w_s}{w_u})}{L_u(\frac{w_s}{w_u})}\right)^{\theta-1} \tag{5.3}$$

By plugging in the initial guess of A_s, the parameter θ, and the aggregated labor supplies derived as described in section 5.1, the only remaining unknown is the skilled-labor wage w_s. Thus, the firm will choose the optimal w_s through this condition. For the two heterogenous household types (skilled-unskilled and unskilled-skilled), female labor supply depends on the wage ratio. Therefore, women in these two households will observe the wage ratio w_s/w_u offered by the firm and adjust their labor supply accordingly through equations (4.14) and (4.15).

Subsequently, given the now explicit aggregate labor supplies L_u and L_s, the firm chooses a new optimal A_s through its second optimality condition:

$$\left(\frac{A_s}{A_u}\right)^{\omega-\theta} = \gamma \left(\frac{L_s}{L_u}\right)^{\theta} \tag{5.4}$$

After obtaining A_s, the firm again determines w_s via equation (5.3), which yields new labor supplies by the households, leading to a new A_s. This process is repeated until the optimal A_s set by the firm is within close distance of the previous iteration's value. This final optimal A_s will be the equilibrium value and pin down the equilibrium skilled wage w_s and labor supplies L_s and L_u, as well as the individual female labor supplies.

6 Results

6.1 Basic Model

Simulating the model as described, taking the data inputs from the CPS and AHTUS microdata and the parameters as noted in table 5.1, the model outputs the skilled labor productivity A_s, the skilled wages w_s and the aggregate labor supplies L_u and L_s for each model period. The results are plotted in figure 6.1.

Figure 6.1: Results of the Basic Model with Fixed $\gamma = 1$

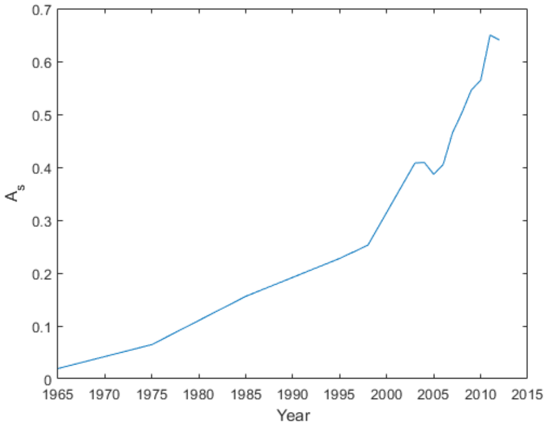

(a) Productivity Ratio A_s/A_u

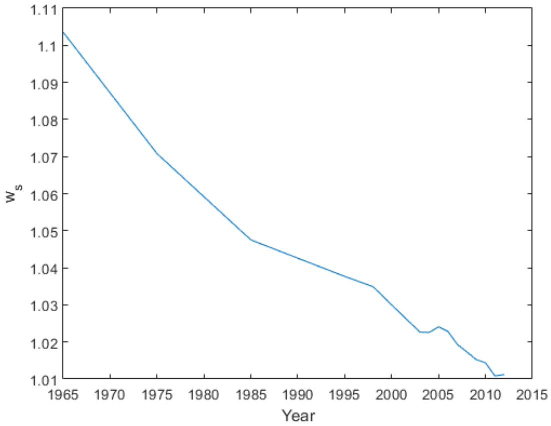

(b) Skill Premium w_s/w_u

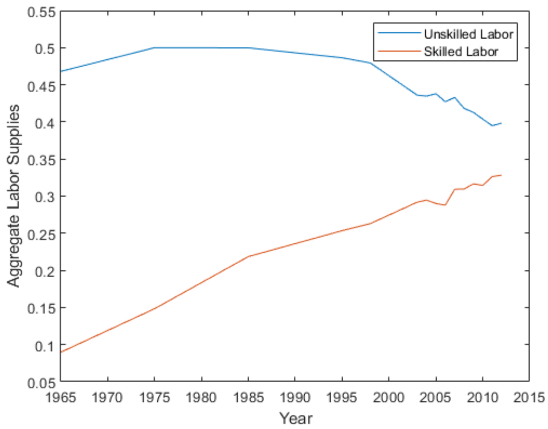

(c) Aggregate Labor Supplies L_s and L_u

These first results indicate that the model is able to generate an endogenous growth of the productivity ratio A_s/A_u over time (figure 6.1a), which is particularly pronounced after the turn of the century. This is caused by the availabilities of aggregated skilled and unskilled labor (L_s and L_u) in each period (figure 6.1c). As figure 3.6 showed, more than 80% of the CPS sample consisted of unskilled-unskilled households in the 1960s. Hence, firms during this time choose a production technology that more efficiently utilizes the vastly available unskilled labor, because the technology constraint (4.7) imposes a trade-off between the two productivities. This results in a tremendously low ratio of skilled to unskilled productivity (A_s/A_u) close to zero during the early model periods.

However, as more skilled labor becomes available due to the increasing share of college graduates (especially women) and the decreasing time spent in the household freeing up hours for working in the labor market, firms set a more skilled-labor-intensive technology, increasing A_s/A_u. A valuable insight is that the immense growth of A_s/A_u in the 2000s and the current decade cannot be explained by changes in time use, as figure 3.5 showed that male labor supply and both genders' home production hours have remained fairly constant since 2000. Rather, the huge rise of highly-educated women, which is reflected in the increasing shares of skilled-skilled and unskilled-skilled households shown in figure 3.6, boosts the supply of skilled labor. Firms then adapt to this by choosing a more skilled-labor efficient production technology.

The skill premium (equivalent to the college wage premium due to my skill level definition) depicted in figure 6.1b on the other hand does not develop as expected. While figure 3.12 illustrated a sharp increase of the college premium since the end of the 1970s from 1.3 to 1.7 today, the skill premium in my model decreases over time, approaching unity. This can again be explained by the relative availability of skilled and unskilled labor shown in figure 6.1c: Since skilled labor is extremely scarce during the early period, firms are willing to pay a premium to skilled workers by setting a relatively high skilled wage w_s despite producing with a more unskilled-labor-efficient technology. Over time, as more skilled labor becomes available, firms no longer have to pay this premium in order to attract skilled workers, which is reflected in the skill premium decrease in figure 6.1b.

Interpreting these results, it is evident that the endogenous growth of the productivity ratio A_s/A_u is not strong enough to generate even a modest skill premium rise. Even changing the parameter values σ, θ and ω does not change this anomaly. While reducing ω flattens the skill premium decrease, it is bounded by $\omega > \theta/(1-\theta)$, as Caselli and Coleman (2006) note, in order to rule out asymmetric equilibria, implying $\omega > 0.4$. In any case, I could not generate a skill premium increase for any combination of parameters while keeping γ fixed to one. In order to gain insight into what growth pattern of the productivity ratio A_s/A_u would be required to replicate the skill premium growth found in the data, I next compute the implied productivity level set by the firm given wages and labor supplies. This abstracts from the double-decision problem of the firm and is exactly what Caselli and Coleman (2006) do by plugging the ratio of wages and labor supplies from the data into equation (4.9).

6.2 On Skill-Biased Technological Change

In order to estimate the A_s/A_u growth consistent with the skill premium increase in the data, I first need to obtain the college wage premium for each model period. To do so, I take the average skilled and unskilled CPI-adjusted hourly wage of all individuals reporting positive hours worked in the week before the survey across gender for each year from the CPS (Flood et al. 2018). It is crucial to compute the wage of all workers, not just full-time full-year workers, because most women in the model will be working part-time. In addition, despite the labor supply decision only being modeled for women, I take the average wage across the whole sample due to my assumption of no gender wage discrimination. Thus, there are only two wages, w_s and w_u, irrespective of gender.

As my model includes households unlike the model of Caselli and Coleman, it is sufficient to only obtain the college wage premium w_s/w_u from the data, since this will endogenously determine the individual and aggregate labor supplies through the households' utility maximization. Hence, I compute the firm's optimal A_s given wages by solving the following equation for each model period:

$$\frac{w_s}{w_u} = \left(\frac{A_s}{A_u}\right)^{\theta} \left(\frac{L_s(\frac{w_s}{w_u})}{L_u(\frac{w_s}{w_u})}\right)^{\theta-1} \tag{6.1}$$

It has to be noted that this method of course merely yields the labor productivity ratios that perfectly reproduce the skill premium growth. If one were to subsequently use the obtained A_s, this would be equivalent to not modeling a firm at all and exogenizing wages by substituting w_s and w_u into the female labor supply conditions (4.12) to (4.15). Hence, this analysis only serves the illustrative purpose of comparing the required productivity growth to the one obtained endogenously in figure 6.1a. The resulting skilled labor productivity growth path is presented in figure 6.2.

Figure 6.2: Productivity Ratio A_s/A_u Required to Replicate Skill Premium Growth of the Data

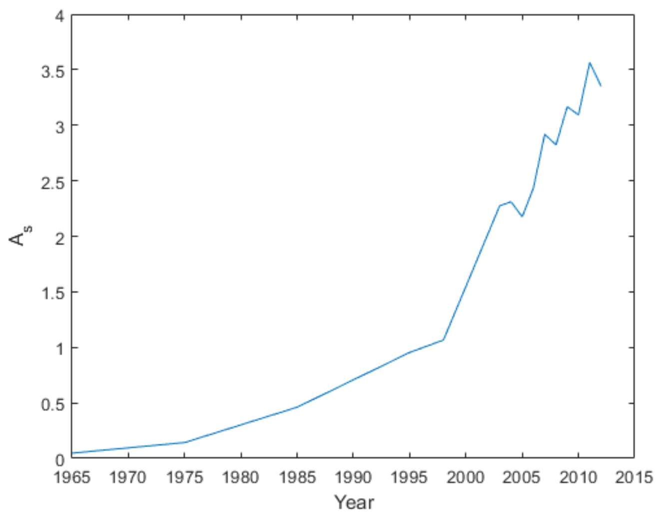

It can be seen that the overall growth pattern is strikingly similar to the endogenous one depicted in figure 6.1a. After a modest increase during the 20th century, the growth of A_s/A_u accelerates from 1998. Subsequently, following a marginal decline in 2005, the growth continues until 2012, before the ratio decreases again in 2012. However, the main difference between the endogenous and exogenous growth of relative skilled-labor productivity is that the latter is far steeper. While the initial ratio in 1965 is relatively comparable (0.02 in the endogenous and 0.05 in the exogenous case), the growth rate is considerably higher for the exogenous ratio, which reaches a maximum of 3.6 in 2011.

This implies that firms in reality choose production technologies that are much more geared towards skilled labor than they do in my model. Examining a potential cause for why my model fails to reproduce this result requires a reconsideration of constraint (5.1), which defines the set of available production technologies:

$$(A_s)^\omega + \gamma(A_u)^\omega \leq B \tag{6.2}$$

Here, γ can be interpreted as the relative price of the productivity of unskilled workers (A_u) compared to that of skilled workers (A_s), as noted by Liu (2017). Furthermore, Liu explains that a decrease in γ over time would represent unskill-biased technological change, and provides the example of the invention of the assembly line making it cheaper for firms to increase the productivity of unskilled workers (A_u). Liu models this by slightly altering the technology choice constraint to $\delta(A_s)^\omega + \gamma(A_u)^\omega \leq B$ and defining $\delta = 0.668/\gamma$. In this notation, a decrease of δ would indicate skill-biased technological change. In contrast, in my (or Caselli and Coleman's) notation without δ, skill-biased technological change would be represented by an increase of γ.

As noted in section 5.2, I obtained γ from Liu who fixes it at 1. This is because Caselli and Coleman (2006) do not report their values for γ, only stating that γ differs between countries, while ω is constant across countries. In their static model, more developed countries have already moved along the technology frontier, which describes the set of available technologies. Firms in these countries have shifted their production towards more skilled-labor-efficient technologies than firms in less-developed countries, where the technology set incurs a larger relative cost on A_s for the firms. This greater cost of skilled labor productivity is modeled via a lower γ for more-developed countries in Caselli and Coleman, which means that the relative price of unskilled labor productivity A_u is lower.

Applying this relationship to my model is equivalent to letting γ vary between the years, as the U.S. moves along the technology frontier through skill-biased technological change. Hence, while skill-biased technological change is reflected in an increasing skilled labor productivity ratio A_s/A_u in my model, its root lies in a growing γ over time. This makes it less costly for firms to increase the relative productivity of skilled labor (A_s). A possible reason for this, as noted by Liu (2017), is the advent of computers, which have made it much easier for the firm to enhance the productivity of their top-educated workers.

My initial result in figure 6.1 therefore indicates that when facing a constant γ every year, firms do not choose a high enough skill-biased productivity growth – given the available labor supplies – to increase the skill premium over time. Thus, the next step is to estimate how γ would have to evolve over time such that my model yields the greater growth of A_s

presented in figure 6.2. This is achieved through the second firm optimality condition:

$$\left(\frac{A_s}{A_u}\right)^{\omega-\theta} = \gamma \left(\frac{L_s}{L_u}\right)^{\theta} \tag{6.3}$$

Solving this for γ, taking the wage ratio w_s/w_u (the skill premium) from the data and the productivity ratio A_s/A_u obtained as described above through equation (6.1), yields the growth path of γ depicted in figure 6.3:

Figure 6.3: γ Required to Replicate Skill Premium Growth of the Data

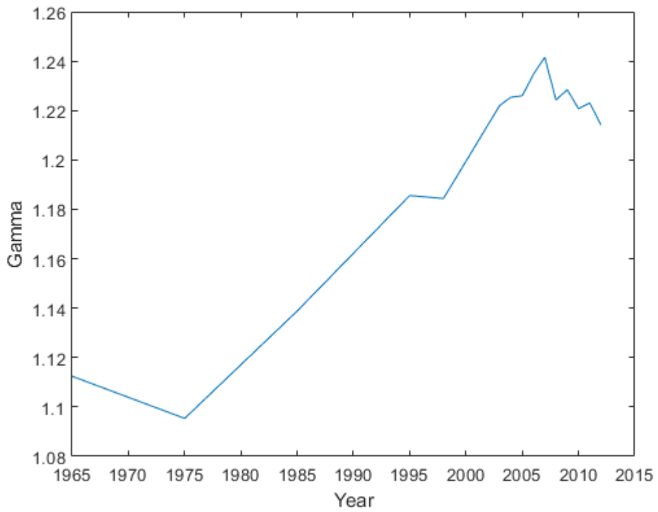

Additionally assuming this γ growth effect therefore enhances the modest growth of A_s/A_u generated solely by the increased availability of skilled labor L_s over time in the initial model (figure 6.1a). Both effects, the exogenous γ growth and the increased availability of skilled labor, combined are needed in order to generate the growth path of A_s/A_u necessary to replicate the skill premium development found in the data. This relation is also found by Liu (2017), who states that the skill premium is greater when the cost of the skill-biased technology (A_s in my model) is relatively cheaper, meaning when γ is larger. This is because a greater γ, resembling a higher relative cost of unskilled labor productivity, incentivizes firms to gear their production technology towards skilled labor, resulting in a skill premium rise and a distinct increase of A_s/A_u over time.

To summarize, the results presented here suggest that firms utilize a more skill-biased technology over time (reflected in a higher A_s/A_u ratio) as the supply of skilled labor increases, which is amplified by the introduction of new technologies driving down the relative price of skilled workers' productivity. These combined effects together cause firms to select more skill-biased technologies over time. Skill-biased technological change in my model is therefore endogenous in the sense that firms choose to shift production towards more skilled-labor-efficient technologies as γ increases over time.

In the following, I will analyze female labor supply and earnigns inequality while treating the determined γ growth path as exogenous. As noted above, this is equivalent to omitting the firm and exogenizing wages, since estimating γ or A_s/A_u myself will always yield the ratio that perfectly reproduces the skill premium. However, assuming an exogenously growing γ is a necessary step in order to ensure the model fit, given the lack of data on A_s/A_u for the U.S. over time. Hence, while I keep the firm in my model, by imposing the γ development shown in figure 6.3, it will always select just the right A_s/A_u ratios each year to reproduce the skill premium from the data, which de facto exogenizes wages.

6.3 Updated Model

Now I simulate the same model as in section 6.1, again using the guess & iterate algorithm described in section 5.3. The only difference is the implementation of the time-invariant γ estimated as described.

Figure 6.4: Results of the Updated Model with Time-Varying γ

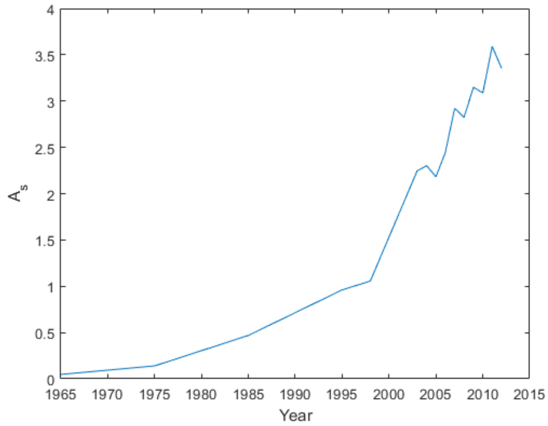

(a) Productivity Ratio A_s/A_u

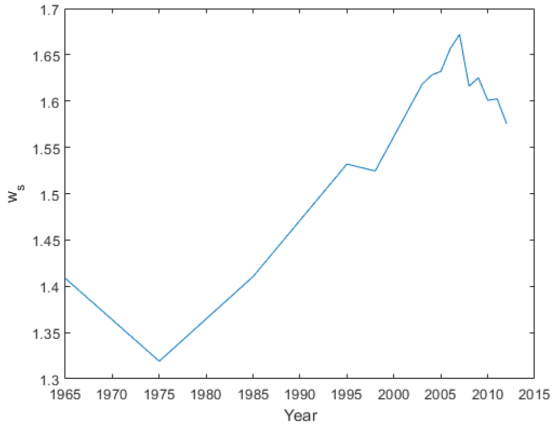

(b) Skill Premium w_s/w_u

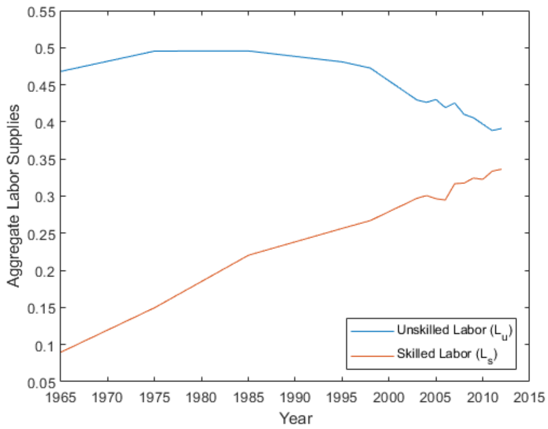

(c) Aggregate Labor Supplies L_s and L_u

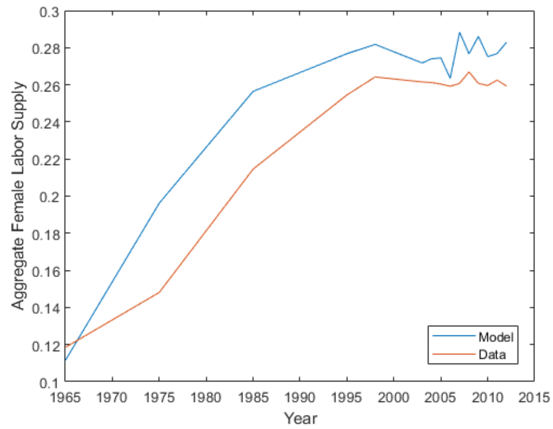

(d) Aggregate Female Labor Supply

Source: Data from IPUMS-CPS (2018) and IPUMS-AHTUS (2018)

Comparing the results of the updated model depicted in figure 6.4 to those of the initial model in figure 6.1, it is evident that adding the exogenously growing γ has altered the growth paths of the productivity ratio (6.4a) and the skill premium (6.4b). As expected, the productivity ratio develops exactly as in figure 6.2, since γ was calibrated to replicate the skill premium growth, just as A_s was in the beginning of chapter 6.2. Likewise, the skill premium replicates the data perfectly due to exogenizing γ, showing a steep increase from 1975 followed by a decline since the mid 2000s, consistent with figure 3.12.

As figure 6.4c demonstates, having γ grow over time does not change aggregate labor supplies relative to the basic model. Rather, it functions as a scale factor on the productivity ratio A_s/A_u, which affects the skill premium through equation (6.1). It can be seen that the total amount of unskilled labor employed decreases slightly, while aggregated skilled labor grows almost continuously. This reflects the increased demand for skilled labor over time caused by the shift in production technology used by the firms. As both

γ and the share of educated workers in the sample grow, firms adapt their technology to utilize skilled labor more efficiently and employ more skilled workers. The fact that the rise of L_s dominates the decrease of L_u implies that employment has increased overall. Interestingly, this illustrates the increasing female labor force participation, since hours worked by men have been stagnating and since all men in the sample are working. Hence, the aggregate labor supply increase can be attributed to women both entering the labor force and increasing their work hours.

In support of this finding, figure 6.4d depicts the aggregate female labor supply, both as predicted by the model and as found in the CPS data (average weekly hours worked for women in each year). It can be seen that female labor supply in the model displays the same pattern as the data, namely a stark increase that flattens during the 1990s. While the model slightly overpredicts female labor supply compared to the data, the overall model fit regarding female labor supply is satisfactory. In order to interpret figure 6.4d, it has to be noted that aggregate female labor supply in the model is derived as the weighted sum of the individual female labor supplies for each household type:

$$L_f = \phi_{u,u} l_{u,u,f} + \phi_{u,s} l_{u,s,f} + \phi_{s,u} l_{s,u,f} + \phi_{s,s} l_{s,s,f} \tag{6.4}$$

Hence, the interpretation of figure 6.4c is straight-forward: Aggregate female labor supply L_f is the weekly labor supply of an average woman, where a value of for instance 0.2 denotes 20 hours of market work per week. Therefore, the average work hours of women in my model have increased from just over 10 to almost 30 hours per week, whereas the actual increase in the data is slightly less pronounced, only reaching close to 26 hours in 2012. One potential reason for this is that women in my model receive the same wage as men of the identical skill level, making them earn more than in practice due my abstraction from the gender wage gap. Furthermore, omitting the time cost of fertility in my model also leads to overpredicting average female hours worked.

The aggregate labor supplies by skill level, L_s and L_u, presented in figure 6.4c can however not be interpreted the same way, as the population weights in equations (4.5) and (4.6) do not add up to one. This is because L_s and L_u not only depend on the average amount of hours worked of each (skilled or unskilled) individual, but also on the distribution of skill levels within the sample. The share of women in the sample on the other hand is constant at 0.5 in each year, allowing for the direct interpretation of L_f.

Besides female labor supply, the other economic variable I will analyze using this model is earnings inequality. In particular, I distinguish between intra-household inequality, measured by the average ratio of women's weekly earnings relative to their male partners, and inter-household inequality, denoted by the Gini coefficient of weekly household earnings (the sum of male and female earnings within each household). The results are presented in figure 6.5, compared to the data equivalents obtained from the CPS. Note that the values for the data curves have been calculated for the complete sample in Stata and have only been imported into MATLAB in order to illustrate them side by side with the model estimates.

Figure 6.5: Intra- and Inter-Household Earnings Inequality

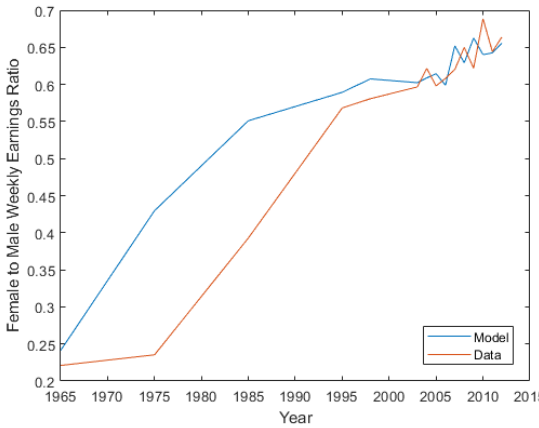

(a) Average Female to Male Weekly Earnings Ratio Within Households

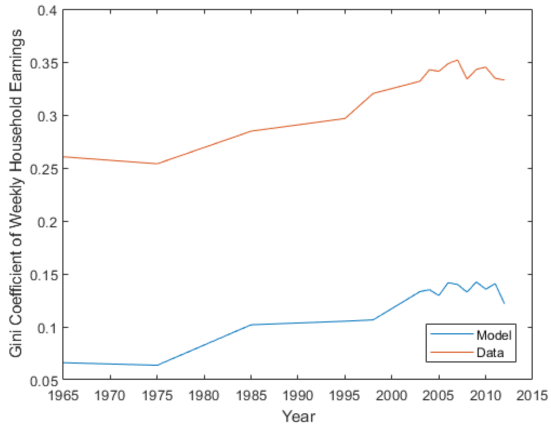

(b) Gini Coefficient of Weekly Household Earnings

Source: Data from IPUMS-CPS (2018)

As can be seen, the model fits the intra-household earnings inequality very well, in particular in recent years, while there exists a persistent gap between the two during the 1970s and 1980s in which the model predicts relative female earnings to be greater than in the data. As male working hours in the model are taken from the data, and as female working hours are very close to the actual values (figure 6.4d), this indicates that wages as the only other determinant of weekly earnings are responsible for the poorer model fit in the early model periods. Noting that one of the limitations of my model is its lack of the gender wage gap, this finding reflects that the gender wage gap was more pronounced during there years. Hence, in the model where the (un)skilled wage is the average wage earned by (un)skilled male and female workers, women earn a larger wage than they actually do in practice. This effect is reversed for men, whose hourly wage in the model is lower compared to the data, as it is driven down by women.

For the inter-household earnings inequality, the model results share the same pattern as the data, but are considerably lower. This can be explained by the lack of earnings dispersion in the model. As there are only four distinct household types, there are consequently only four different household earnings values in each model period. Keeping in mind that the Gini coefficient measures the area below the Lorentz curve, the lack of data points (distinct household earnings) in the model leads to a lower Gini coefficient. This is the case in particular at the right tail end of the earnings distribution, where a large share of total earnings in the economy is concentrated. The model does not consider this, as it only outputs one value per household type.

Due to its good fit for female labor supply and reasonable replication of both the patterns in both measures earnings inequality, this updated version of the model with exogenous γ growth will function as the baseline for the ensuing analysis, which all other results will be compared against. In the next chapter, I will use the model to test different candidates that could explain the rapid increase of female labor supply over the last decades of the 20th century shown in figure 6.4d.

7 Counterfactual Analyses

This chapter is concerned with conducting two sets of counterfactual analyses. First, I test four different empirical trends for their impact on female labor supply and earnings inequality. Second, I investigate whether the model results differ notably between married and cohabiting households and how well the model fits the data on cohabiting households.

7.1 Female Labor Supply Candidates

Using the model set up as adjusted in chapters 4 to 6, I test four candidates for explaining the increase of female labor supply from the 1960s to the 1970s: The decreasing time women spend on home production, the rise of assortative matching, the rising female educational attainment and skill-biased technological change increasing the skill premium. The goal is determine which of these factors caused women to participate more in the labor market, both on the extensive and the intensive margin.

To test the importance of these four candidates, there are two possible approaches: Either one could take model with the values for 1965 and change one of the four candidates over time according to the data, leaving everything else constant at the 1965 level, or one could change all values except for the variable of interest over time. I will utilize the latter method, as this allows for the interaction of the remaining variables changing over time. As the purpose of this thesis is to investigate female labor supply and earnings inequality under skill-biased technological change, I plot the productivity ratio A_s/A_u chosen by the firms, aggregate female labor supply, intra-household inequality and inter-household inequality for each candidate, and compare them to the data. It is important to note that close resemblance of the model estimates to the data therefore occurs when assuming a candidate to not change from its 1965 value barely affects the result. In such a case, this provides evidence for a candidate being rather insignificant in explaining the trends in the data.

7.1.1 Home Production

Starting with home production, I simulate the model as in section 6.3, with the only exception being that both female and male weekly household hours remain on their 1965 levels of 44 hours on average for women and 10 hours for men. This represents the case where home production technologies have not become more efficient, meaning everyone is forced to keep using the laundry machines, stoves, vacuum cleaners etc. that were available in 1965. Furthermore, assume that men do not increase their participation in household work and child care over time, but continue to only spend a fourth of their wife's time in the household.

Figure 7.1: Counterfactual Analysis with Constant 1965 Household Hours

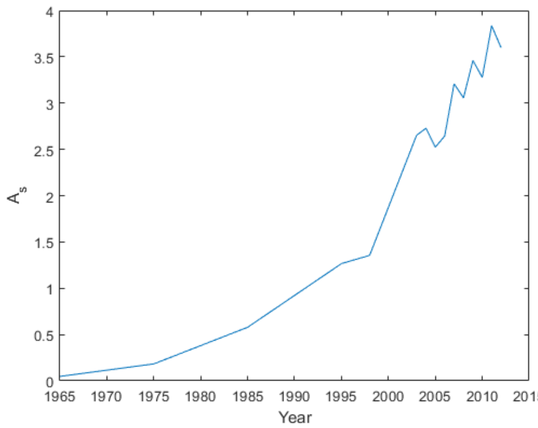

(a) Productivity Ratio A_s/A_u

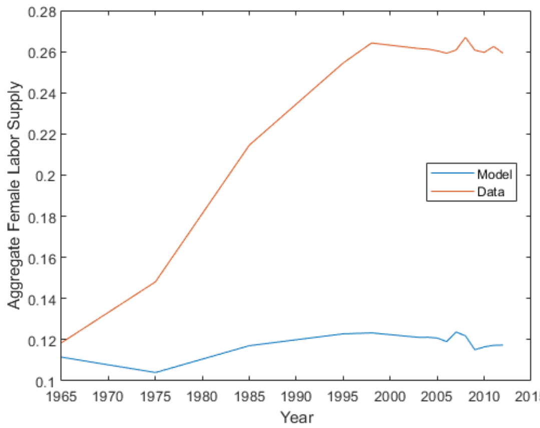

(b) Aggregate Female Labor Supply

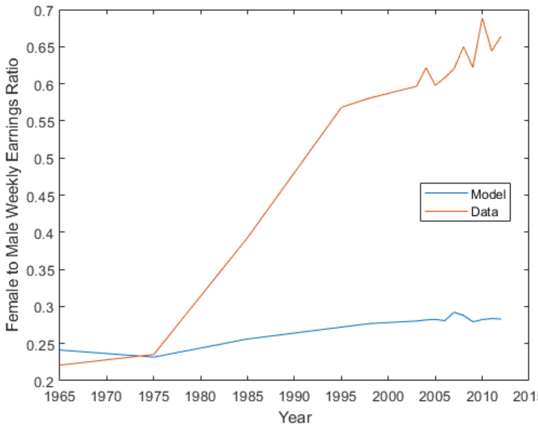

(c) Average Female to Male Weekly Earnings Ratio Within Households

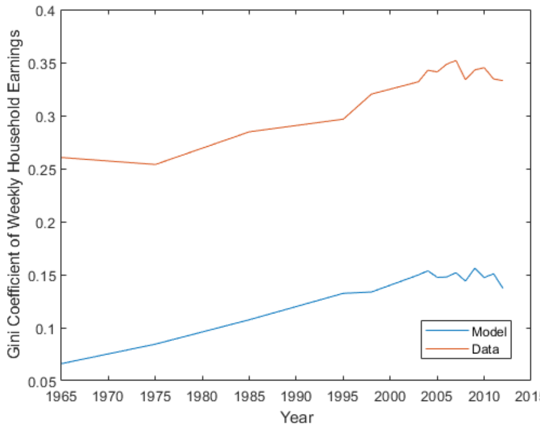

(d) Gini Coefficient of Weekly Household Earnings

Source: Data from IPUMS-CPS (2018) and IPUMS-AHTUS (2018)

The results of this model economy are presented in figure 7.1. First, as one would expect, confining women to their role of a 1965 housewife absolutely suppresses their labor supply. Figure 7.1b shows that the aggregate female labor supply under constant household hours stagnates and never exceeds 13 hours a week. Hence, women today would be as scarce on the labor market as in the 1960s if they still had to spend an average of 44 hours in home production as full-time housewives.

Turning to intra-household, women's relative earnings also show almost no increase over time, despite the growing skill premium. This is due to their constant labor supply over time, which keeps their earnings low, while men's earnings are fixed in the model due to obtaining male working hours from the data. Therefore, when spending the same amount of hours in the household as in 1965, women's weekly earnings continue to be around a quarter of men's.

Lastly, the inter-household inequality is barely affected by keeping home hours constant. This indicates that the role of women in reducing households earnings inequality is very limited to begin with, as the same pattern persists when women are barely working.

7.1.2 Assortative Matching

In order to assess the effects of different matching patters over time, I change the population weights $\phi_{i,j,m/f}$ in each model period. First, I calculate the shares of both unskilled and skilled women living with an either equally or differently skilled partner in 1965. For instance, of the 92.1% unskilled women in 1965, 89% live with an unskilled and 11% with a skilled partner, while of the 7.9% skilled women in 1965, only 32% live with an unskilled opposed to 68% with a skilled man. I then apply these ratios to the shares of skilled and unskilled women in each year. Doing so leaves the education level of women unchanged and only imposes the matching ratios present in 1965 by effectively keeping male education levels at their 1965 values.

Figure 7.2: Counterfactual Analysis with Constant 1965 Matching

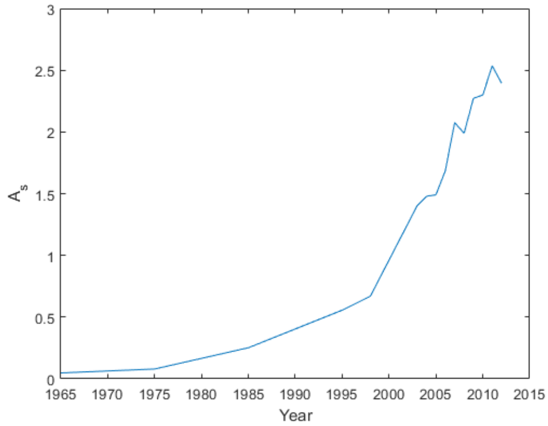

(a) Productivity Ratio A_s/A_u

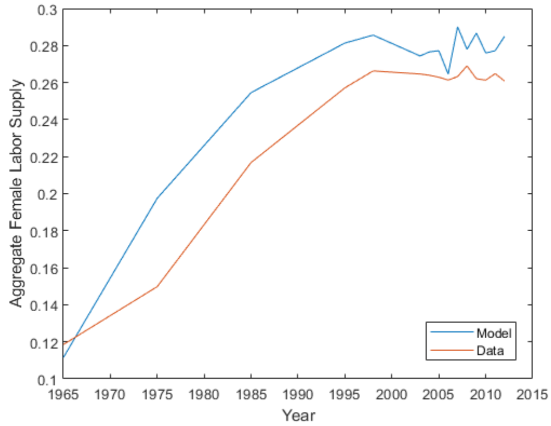

(b) Aggregate Female Labor Supply

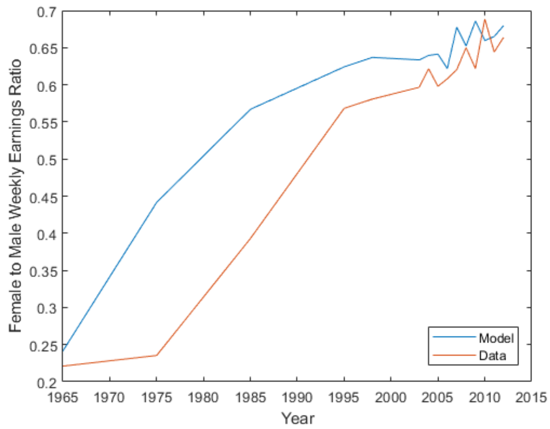

(c) Average Female to Male Weekly Earnings Ratio Within Households

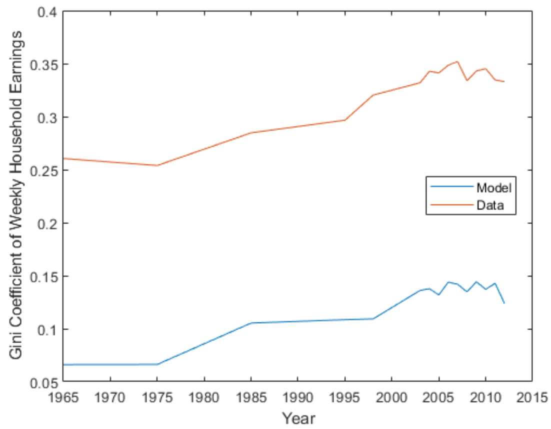

(d) Gini Coefficient of Weekly Household Earnings

Source: Data from IPUMS-CPS (2018) and IPUMS-AHTUS (2018)

The results are illustrated in figure 7.2 and indicate that the productivity growth is less pronounced than in the baseline model, meaning the firm chooses a more unskilled-labor intensive production technology. This is due to the constant education level of men keeping the supply of skilled workers relatively low. Surprisingly, female labor supply and inter-household inequality are unchanged compared to the baseline model. This can be attributed to the fact that while men influence female labor supply through their working and households hours, these do not differ too much between skilled and unskilled men. Furthermore, due to the data limitations of the AHTUS, female household hours do not depend on their partners' education. Intra-household inequality is slightly reduced compared to the baseline model, as expected when keeping male education at the low 1965 level. Overall, these results indicate that changing matching patterns over time cannot explain the trends in the data, as omitting them does not alter the model results.

7.1.3 The Increase of Female Education

Regarding the rising education of women, figure 7.3 depicts the model results for when women would have kept the same education levels over time as they had in 1965. I obtained these by again changing the population weights, now applying the ratio of skilled (92%) to unskilled (8%) women in 1965 to all other model periods.

Figure 7.3: Counterfactual Analysis with Constant 1965 Female Education Level

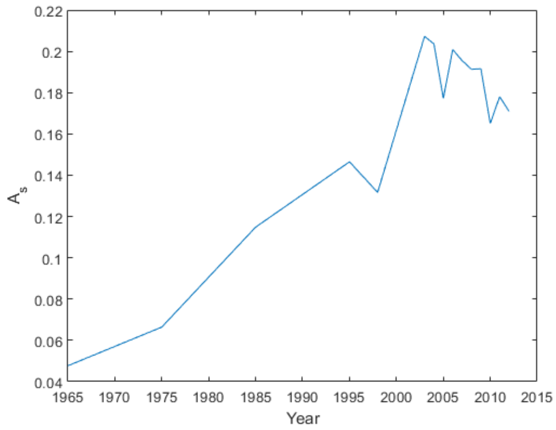

(a) Productivity Ratio A_s/A_u

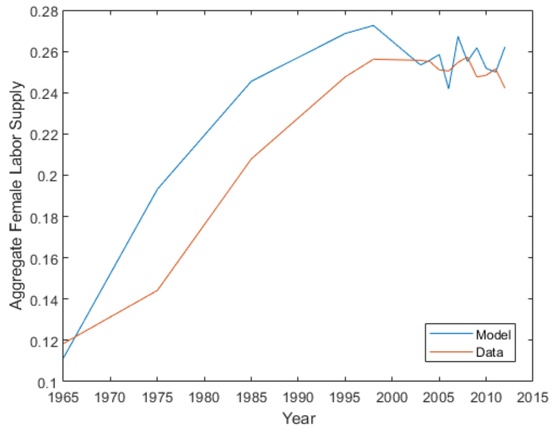

(b) Aggregate Female Labor Supply

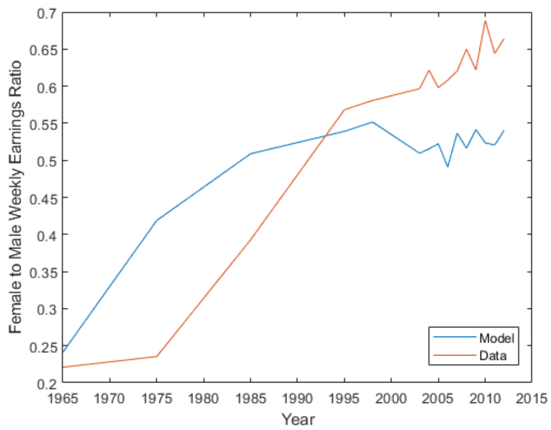

(c) Average Female to Male Weekly Earnings Ratio Within Households

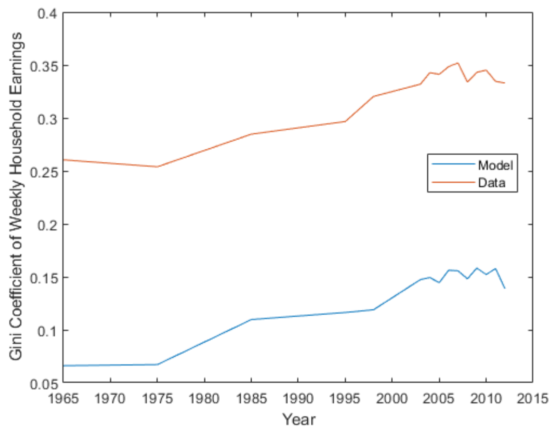

(d) Gini Coefficient of Weekly Household Earnings

Source: Data from IPUMS-CPS (2018) and IPUMS-AHTUS (2018)

Firstly, the productivity ratio changes distinctively. While it still displays a shift towards more skilled-labor-efficient productivities over time, this effect is very weak, as indicated by the low maximum value of 0.2 of A_s. This is caused by the lower supply of skilled labor due to the lack of high-educated women. Therefore, firms have a lower incentive to utilize skilled labor more efficiently. The fact that female labor supply in this case is still growing strongly shows that these overwhelmingly unskilled women are still employed, albeit in unskilled jobs.

This relationship is also evident in the intra-household inequality. In earlier model periods, where firms use more unskilled-labor-efficient production technologies, female relative earnings exceed the values found in the data. However, as firms (although only slightly) shift production towards skilled labor, female relative earnings no longer grow and stagnate below the model ratio. Finally, the lack of female education growth has almost no influence on inter-household inequality compared to the baseline model, which could be explained by the difficulty to accurately estimate household earnings dispersion between households when there are only 4 representative household types with no within-group earnings variance.

7.1.4 Skill-Biased Technological Change

At last, I assume a lack of skilled-labor productivity growth by holding A_s constant at the 1965 equilibrium value of 0.0478.

Figure 7.4: Counterfactual Analysis with Constant 1965 Productivities

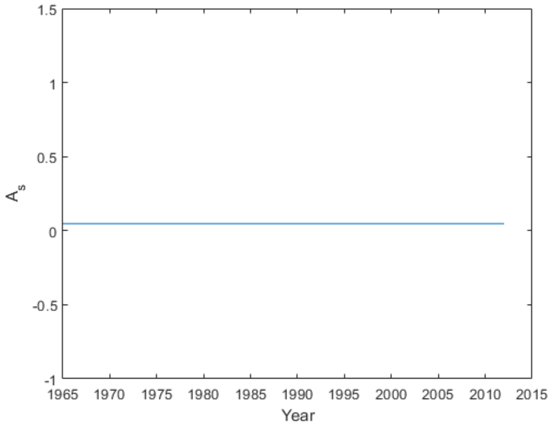

(a) Productivity Ratio A_s/A_u

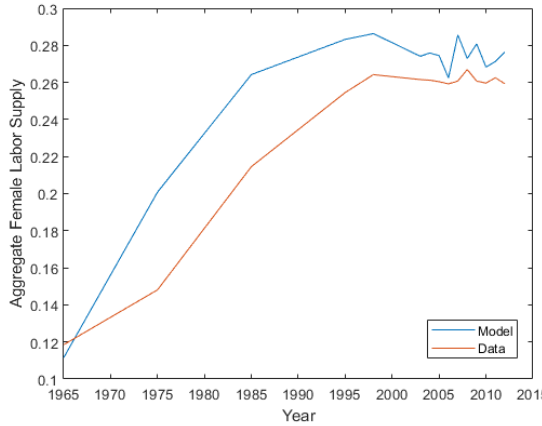

(b) Aggregate Female Labor Supply

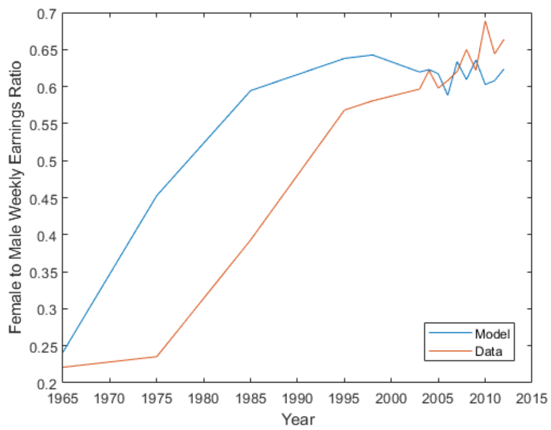

(c) Average Female to Male Weekly Earnings Ratio Within Households

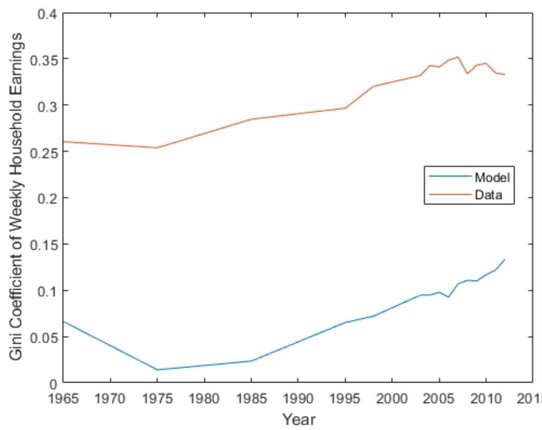

(d) Gini Coefficient of Weekly Household Earnings

Source: Data from IPUMS-CPS (2018) and IPUMS-AHTUS (2018)

A different way to model the lack of skill-biased technological change is to keep γ fixed over time. However, as this has already been discussed extensively when solving the basic model in section 6.1, I maintain the exogenous growth in γ and instead fix A_s. Note that this causes the skill premium w_s/w_u to decrease monotonically over time from 1.4 in 1965 to 0.5 in 2012. This is because, as explained in section 6.2, skill-biased technological change needs to consist of both A_s/A_u and γ increasing over time.

The results, shown in figure 7.4, illustrate that the lack of skill-biased technological change also does not exert a large effect on female labor supply and on intra-household inequality. Only the development of intra-household inequality over time is slightly lower than in the baseline model in the past two decades, as women do not benefit from the reversal of the gender education gap without skill-biased technological change. Instead of the usual pattern of a slight but continuous increase, the Gini coefficient in this case decreases at first, meaning that the overall inequality decreases. This indicates that the increase in between-household inequality from 1975 to 1985 can be attributed to the skill premium growth.

7.2 On Married and Cohabiting Households

The final part of my analysis consists of comparing the model results for cohabiting to those of married households in order to determine whether the civil status of couples changes the model results. As the CPS sample consists of 1,343,031 married and only 41,546 cohabiting individuals, one has to note the issue of the small sample size for cohabiting observations. Furthermore, due to how married couples are defined in the CPS, identifying them is only possible from 1995, as noted in the Data Appendix A.

As the results for strictly married couples look almost identical to the full sample, due to the low share of cohabiting households, one can compare the results for cohabiting households in figure 7.5 to those in figures 6.4 and 6.5. It can be seen that the patterns in female labor supply and earnings inequality for cohabiting couples differ greatly from those for married couples, with the model again displaying a good fit to the data. While the erratic behavior of female labor supply in figure 7.5b is due to the small scale of the y axis, the fluctuations of intra-household earnings inequality in the data graph in figure 7.5c can be attributed to the low sample size. An interesting result is that the average female to male earnings ratio is close to unity, much larger than for the sample of married households. This indicates that women living in cohabitation have earnings more similar to their partners' than married women. Finally, inter-household inequality is fairly constant over time due to the relatively small time frame of 23 years, but is greatly underpredicted by the model.

Figure 7.5: Counterfactual Analysis with Only Cohabiting Households

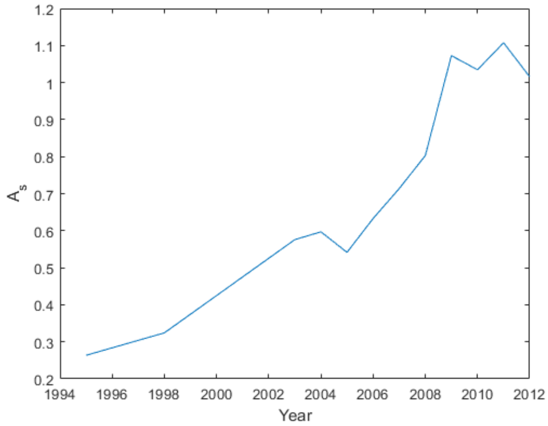

(a) Productivity Ratio A_s/A_u

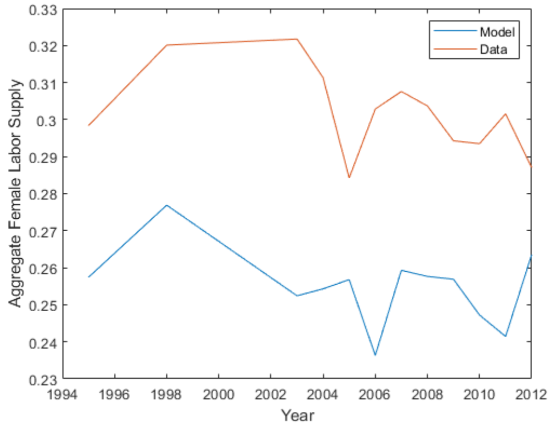

(b) Aggregate Female Labor Supply

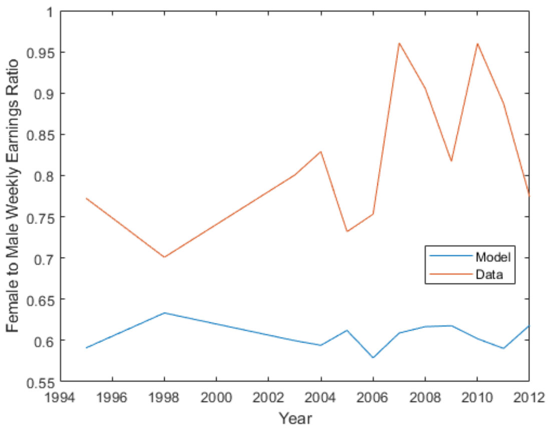

(c) Average Female to Male Weekly Earnings Ratio Within Households

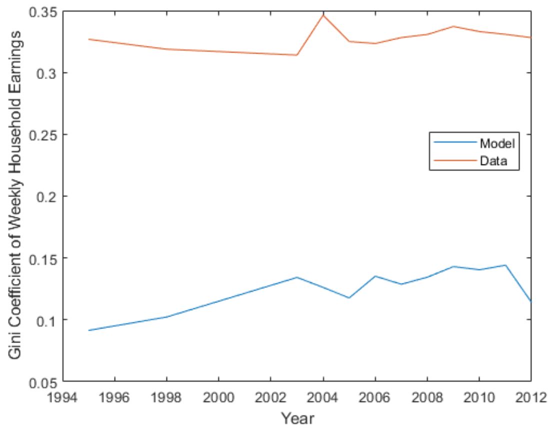

(d) Gini Coefficient of Weekly Household Earnings

Source: Data from IPUMS-CPS (2018) and IPUMS-AHTUS (2018)

8 Conclusion

To summarize, I set up a heterogeneous agent household model in which wages are determined endogenously by including firms which in addition choose their production technology as in Caselli and Coleman (2006). As the first results showed, this model failed to replicate the increase of the skill premium found in the data, because the implied endogenous shift towards more skilled-labor-efficient production technologies was not strong enough. By introducing an external decrease of the relative price of skilled-labor productivity, I ensured that firms shift their production further towards skilled labor, such that skill premium growth from the data is replicated. The resulting model, while de facto having exogenous wages, fits the data on female labor supply as well as both intra- and inter-household inequality very well.

Subsequently, I used this model to analyze four different explanations for the steep growth of female labor supply from the 1960s to the 1990s: The reduction in household time, the rise of assortative matching, the increased educational attainment of women, and the growing skill premium over time. In line with Greenwood et al. (2005), I found that the reduction in household time exerts by far the greatest effect on female labor supply and subsequently on intra-household earnings inequality. Assortative matching, the increase of female education and the rise of the skill-premium, at least in my sample, barely change the results compared to the baseline model. This finding could be due to having exogenized wages. Whereas reduced household hours have a direct effect on the agents' time constraint, leading to more hours worked, the other three effects mostly function through higher wages, which my model cannot depict accurately.

Lastly, I simulated the model using a sample consisting of only cohabiting households. While this sample is rather small and only available for more recent years, the results show that the model again fits the data reasonably well. Comparing the results of the cohabiting to the complete sample, one observes that women in cohabiting couples work more and have earnings closer to their partners'.

It has to be noted that my model is built on many assumptions which serve to simplify the analysis. This naturally creates numerous possible extensions to the model for future research. An obvious one would be endogenizing the agents' education choice, rater than imposing the shares of non-college- and college-educated individuals exogenously like I did. This would require making the model dynamic, for example by setting up an OLG model in which individuals accumulate human capital over time, similar to the groundbreaking model by Ben-Porath (1967). Agents would then decide based on their productivity level and expected wages, whether they would obtain a college degree or not. The results of this analysis would yield valuable insights into causes for the growing female educational attainment in the context of increasing skill-biased technological change.

In addition, another factor not included in my model that heavily influences female labor supply, as explained in chapter 2, is fertility. In my Bachelor's thesis, I have modeled this as another element in the woman's utility function, such that women derive utility from consumption, leisure, and having children. This would allow to analyze the trade-off between working and having children that women often face in reality. A simple way of incorporating the effects of fertility on female labor supply would be to have women bear a fixed time cost of motherhood, a forced reduction of their available time they cannot avoid, similar to Erosa et al. (2010). Moreover, by also modeling that women face a quantity-quality trade-off regarding their children as done by Galor and Weil (2000), one could study intergenerational effects. For instance, Guryan, Hurst, and Kearney (2008) find that higher educated women spend more time on their children. This would would probably translate into an amplifying effect on earnings inequality, as high education would be transmitted to the next generations.

Besides the education and fertility decision, my model also does not consider women's decision of entering the labor market by only focusing on the intensive margin of labor supply. In reality, most of the growth of female labor supply can be attributed to increases along the extensive margin of labor supply (Attanasio et al. 2008), as more and more women entered the labor market. By only considering the average hours worked of women in each household type, my model does not account for this decision.

Furthermore, my simplifying assumption of gender-unbiased wages is debatable, as figure 3.11 has presented a 22% gender gap in hourly wages of FTFY workers. While this measure does not control for work experience, occupational choices, hours worked and many other factors that can explain part of the gender wage gap, there always remains an unexplained gender wage gap even after considering gender differences in these variables. Blau and Kahn (2007) estimate the unexplained gender wage gap to be 9% for the U.S. This difference could for example be due to taste discrimination by employers. One way to incorporate this into my model would be by letting wages not only differ by skill level, but also by gender. For instance, one could introduce a wedge $(1 + \tau)$ between male and female wages, which could be interpreted as a tax firms need to pay for hiring women or as a disutility they obtain from employing women.

The inherent weakness of my model due to its unitary framework is explained by Knowles (2013), who argues that since households are not rational agents, the unitary model overstates how aggregate labor reacts to relative wage changes. As a solution to this, he suggests modeling household bargaining, which would improve the model predictions. However, doing so considerably complicates my model, which is why I have refrained from it. One possible method of implementing this is to incorporate a collective rather than a unitary model by weighting the utilities of both spouses and then varying the weights over time, as done by Fernández and Wong (2014). The main difficulty in doing so is estimating the spousal bargaining powers, which are unobserved. A possible way to do this is to use the relative spousal wages as a proxy for bargaining power, thereby modeling that women's power in household decision making has grown.

A final extension popular in the literature that I have not incorporated is to endogenize assortative matching by modeling a marriage (and divorce) market, similar to Gihleb and Lifshitz (2016) and Greenwood et al. (2016). This could be combined with including household bargaining, where being single provides an outside option and would allow for a better understanding of decisions on the household level.

References

Acemoglu, D. (1998). Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *The Quarterly Journal of Economics*, *113*(4), 1055–1089.

Acemoglu, D., Autor, D., & Lyle, D. (2004). Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury. *Journal of Political Economy*, *112*(3), 497–551.

Aguiar, M., & Hurst, E. (2007). Measuring Trends in Leisure: The Allocation of Time Over Five Decades*. *The Quarterly Journal of Economics*, *122*(3), 969–1006.

Albanesi, S., & Prados, M. J. (2017). *Slowing Women's Labor Force Participation: The Role of Rising Income Inequality*, University of Pittsburgh, Unpublished.

Albanesi, S., & Olivetti, C. (2009). Production, Market Production and the Gender Wage Gap: Incentives and Expectations. *Review of Economic Dynamics*, *12*(1), 80–107.

Albanesi, S., & Olivetti, C. (2016). Gender Roles and Medical Progress. *Journal of Political Economy*, *124*(3), 650–695.

Apps, P., & Rees, R. (2004). Fertility, Taxation and Family Policy. *Scandinavian Journal of Economics*, *106*(4), 745–763.

Attanasio, O., Battistin, E., & Ichimura, H. (2007). What Really Happened to Consumption Inequality in the United States? In *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches* (pp. 515–543). National Bureau of Economic Research, Inc.

Attanasio, O., Low, H., & Sánchez-Marcos, V. (2008). Explaining Changes in Female Labor Supply in a Life-Cycle Model. *American Economic Review*, *98*(4), 1517–52.

Attanasio, O., & Weber, G. (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, *103*(6), 1121–57.

Ben-Porath, Y. (1967). The Production of Human Capital and the Life Cycle of Earnings. *Journal of Political Economy*, *75*(4), 352–365.

Bertrand, M., Goldin, C., & Katz, L. F. (2010). Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors. *American Economic Journal: Applied Economics*, *2*(3), 228–55.

Bick, A. (2016). The Quantitative Role of Child Care for Female Labor Force Participation and Fertility. *Journal of the European Economic Association*, *14*(3), 639–668.

Bick, A., & Fuchs-Schündeln, N. (2017). Taxation and Labour Supply of Married Couples across Countries: A Macroeconomic Analysis. *The Review of Economic Studies*, *85*(3), 1543–1576.

Black, S. E., & Juhn, C. (2000). The Rise of Female Professionals: Are Women Responding to Skill Demand? *American Economic Review*, *90*(2), 450–455.

Blau, F. D., & Kahn, L. M. (2007). The Gender Pay Gap: Have Women Gone as Far as They Can? *Academy of Management Perspectives*, *21*(1), 7–23.

Bloom, D. E., Canning, D., Fink, G., & Finlay, J. E. (2009). Fertility, Female Labor Force Participation, and the Demographic Dividend. *Journal of Economic Growth*, *14*(2), 79–101.

Blundell, R., Pistaferri, L., & Saporta-Eksten, I. (2016). Consumption Inequality and Family Labor Supply. *American Economic Review*, *106*(2), 387–435.

Browning, M., Chiappori, P.-A., & Weiss, Y. (2014). *Economics of the Family*. Cambridge Surveys of Economic Literature.

Card, D., & Hyslop, D. R. (2018). *Female Earnings Inequality: The Changing Role of Family Characteristics on the Extensive and Intensive Margins* (NBER Working Papers No. 25387). National Bureau of Economic Research, Inc.

Caselli, F., & Coleman, W. J. (2006). The World Technology Frontier. *American Economic Review*, *96*(3), 499–522.

Cerina, F., Moro, A., & Rendall, M. (2017). *The Role of Gender in Employment Polarization* (Discussion Papers No. 1704). Centre for Macroeconomics (CFM).

Chiappori, P. (1992). Collective Labor Supply and Welfare. *Journal of Political Economy*, *100*(3), 437–67.

Chiappori, P., Iyigun, M., & Weiss, Y. (2009). Investment in Schooling and the Marriage Market. *American Economic Review*, *99*(5), 1689–1713.

Choi, S., & Valladares-Esteban, A. (2016). *On Households and Unemployment Insurance*. Unpublished.

Cortes, G. M., Jaimovich, N., & Siu, H. E. (2018). *The "End of Men" and Rise of Women in the High-Skilled Labor Market* (Working Paper No. 24274). National Bureau of Economic Research.

Doepke, M., & Tertilt, M. (2016). *Families in Macroeconomics* (NBER Working Papers No. 22068). National Bureau of Economic Research, Inc.

Domeij, D., & Flodén, M. (2006). The labor-supply elasticity and borrowing constraints: Why estimates are biased. *Review of Economic Dynamics*, *9*(2), 242–262.

Domeij, D., & Klein, P. (2012). Should Day Care be Subsidized? *The Review of Economic Studies*, *80*(2), 568–595.

Eckstein, Z., & Lifshitz, O. (2015). Household Interaction and the Labor Supply of Married Women. *International Economic Review*, *56*(2), 427–455.

Erosa, A., Fuster, L., & Restuccia, D. (2010). A General Equilibrium Analysis of Parental Leave Policies. *Review of Economic Dynamics*, *13*(4), 742–758.

Erosa, A., Fuster, L., & Restuccia, D. (2016). A quantitative theory of the gender gap in wages. *European Economic Review*, *85*(100), 165–187.

Fernández, R. (2013). Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century. *American Economic Review*, *103*(1), 472–500.

Fernández, R., Fogli, A., & Olivetti, C. (2004). Mothers and Sons: Preference Formation and Female Labor Force Dynamics. *The Quarterly Journal of Economics*, *119*(4), 1249–1299.

Fernández, R., & Rogerson, R. (2001). Sorting and Long-Run Inequality. *The Quarterly Journal of Economics*, *116*(4), 1305–1341.

Fernández, R., & Wong, J. C. (2014). *Divorce Risk, Wages, and Working Wives: A Quantitative Life-Cycle Analysis of Female Labor Force Participation* (Working Paper No. 19869). National Bureau of Economic Research.

Fisher, K., Egerton, M., Gershuny, J., & Robinson, J. (2007). Gender Convergence in the American Heritage Time use Study (AHTUS). *Social Indicators Research: An International and Interdisciplinary Journal for Quality-of-Life Measurement*, *82*(1), 1–33.

Fisher, K., Gershuny, J., Flood, S. M., Roman, J. G., & Hofferth, S. L. (2018). American Heritage Time Use Study Extract Builder: Version 1.2 [dataset]. Retrieved May 4, 2019, from https://www.ahtusdata.org/ahtus/.

Flood, S., King, M., Rodgers, R., Ruggles, S., & Warren, R. J. (2018). Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [dataset]. Retrieved May 4, 2019, from https://cps.ipums.org/cps/index.shtml.

Fogli, A., & Veldkamp, L. (2011). Nature or Nurture? Learning and the Geography of Female Labor Force Participation. *Econometrica*, *79*(4), 1103–1138.

Galor, O., & Weil, D. N. (1996). The Gender Gap, Fertility, and Growth. *American Economic Review*, *86*(3), 374–387.

Galor, O., & Weil, D. N. (2000). Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and beyond. *The American Economic Review*, *90*(4), 806–828.

Gihleb, R., & Lifshitz, O. (2016). *Dynamic Effects of Educational Assortative Mating on Labor Supply* (IZA Discussion Papers No. 9958). Institute for the Study of Labor (IZA).

Goldin, C., & Katz, L. (2002). The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions. *Journal of Political Economy*, *110*(4), 730–770.

Greenwood, J., Guner, N., Kocharkov, G., & Santos, C. (2016). Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation. *American Economic Journal: Macroeconomics*, *8*(1), 1–41.

Greenwood, J., Seshadri, A., & Yorukoglu, M. (2005). Engines of Liberation. *The Review of Economic Studies*, *72*(1), 109–133.

Guner, N., Kaygusuz, R., & Ventura, G. (2014). *Childcare Subsidies and Household Labor Supply* (IZA Discussion Papers No. 8303). Institute for the Study of Labor (IZA).

Guryan, J., Hurst, E., & Kearney, M. (2008). Parental Education and Parental Time with Children. *Journal of Economic Perspectives*, *22*(3).

Hazan, M., & Zoabi, H. (2015). Do Highly Educated Women Choose Smaller Families? *The Economic Journal*, *125*(587), 1191–1226.

Heathcote, J., Storesletten, K., & Violante, G. L. (2010). The Macroeconomic Implications of Rising Wage Inequality in the United States. *Journal of Political Economy*, *118*(4), 681–722.

Herrmann, M., & Machado, C. (2012). *Patterns of Selection in Labor Market Participation.*

Hyslop, D. R. (2001). Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings. *American Economic Review*, *91*(4), 755–777.

Jones, L., Manuelli, R., & McGrattan, E. (2015). Why Are Married Women Working so much ? *JODE - Journal of Demographic Economics*, *81*(1), 75–114.

Katz, L., & Autor, D. (1999). Changes in the wage structure and earnings inequality. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (1st ed., Chap. 26, Vol. 3, Part A, pp. 1463–1555). Elsevier.

Katz, L., & Murphy, K. M. (1992). Changes in Relative Wages, 1963–1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, *107*(1), 35–78.

Knowles, J. (2013). Why are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply. *Review of Economic Studies*, *80*(3), 1055–1085.

Liu, L. (2017). *Three papers in economic growth and inequality* (Doctoral dissertation, University of Southampton).

Martínez, D. F., & Iza, A. (2004). Skill premium effects on fertility and female labor force supply. *Journal of Population Economics*, *17*(1), 1–16.

McPherson, M. S., & Schapiro, M. O. (1991). Does Student Aid Affect College Enrollment? New Evidence on a Persistent Controversy. *The American Economic Review*, *81*(1), 309–318.

Mulligan, C. B., & Rubinstein, Y. (2008). Selection, Investment, and Women's Relative Wages over Time. *The Quarterly Journal of Economics*, *123*(3), 1061–1110.

Ngai, L. R., & Petrongolo, B. (2017). Gender Gaps and the Rise of the Service Economy. *American Economic Journal: Macroeconomics*, *9*(4), 1–44.

Olivetti, C. (2006). Changes in women's hours of market work: The role of returns to experience. *Review of Economic Dynamics*, *9*(4), 557–587.

Pistaferri, L. (2003). Anticipated and Unanticipated Wage Changes, Wage Risk, and Intertemporal Labor Supply. *Journal of Labor Economics*, *21*(3), 729–754.

Ragan, K. S. (2013). Taxes and Time Use: Fiscal Policy in a Household Production Model. *American Economic Journal: Macroeconomics*, *5*(1), 168–92.

Ramey, V. A., & Francis, N. (2009). A Century of Work and Leisure. *American Economic Journal: Macroeconomics*, *1*(2), 189–224.

Rendall, M. P. (2017). *Brain versus brawn: the realization of women's comparative advantage* (tech. rep. No. No. 491). Revised version.

Shore, S. H. (2015). The co-movement of couples' incomes. *Review of Economics of the Household*, *13*(3), 569–588.

United States Bureau of Labor Statistics. (1993). *Labor composition and U.S. productivity growth, 1948-90*. Bulletin of the United States Bureau of Labor Statistics. U.S. Dept. of Labor, Bureau of Labor Statistics.

United States Bureau of Labor Statistics. (2019). U.S. Bureau of Labor Statistics: Consumer Price Index for All Urban Consumers: All Items [CPIAUCSL]. Retrieved May 4, 2019, from https://fred.stlouisfed.org/series/CPIAUCSL.

Waldfogel, J. (1998). The Family Gap for Young Women in the United States and Britain: Can Maternity Leave Make a Difference? *Journal of Labor Economics*, *16*(3), 505–45.

A Data Appendix

For my analysis, I use two customized microdata extracts obtained from the Integrated Public Use Microdata Series (IPUMS): One of the U.S. Current Population Survey (CPS) by Flood et al. (2018) and one of the American Heritage Time Use Study (AHTUS) by Fisher et al. (2018).

A.1 CPS

The CPS is a monthly household labor force survey conducted in the U.S. and covers the period from 1964 to 2018. I use the public IPUMS version of the CPS (IPUMS-CPS), because the data is harmonized for comparisons over time. More specifically, I work with the Annual Social and Economic (ASEC) supplement, also referred to as March Supplement, due to the monthly surveys only being available from 1976. For generating descriptive statistics on the individual and household level, I apply the provided ASEC sample weights *ASECWT* and *ASECWTH* respectively.

First of all, I impose an age restriction of 25 to 55 years, similar to the one used by Mulligan and Rubinstein (2008). The lower bound is set to 25 years because of education being exogenous in my model. Hence, 25 ensures that most individuals will have completed any undergraduate university education. The upper bound is set to 55 such that retirement does not effect the labor supplies and home production hours.

Furthermore, following Mulligan and Rubinstein (2008) and Herrmann and Machado (2012), I restrict the sample to non-Hispanic whites in order to mitigate the effects of changing demographics. For example, Fernández (2013) notes that black women have higher labor force participation early on in their lives compared to white women. While restricting on whites can be done for all years using the *RACE* variable, identifying Hispanics is much more difficult because they are not assigned their own category in *race*, but rather coded as whites in the data. The variable *HISPAN* solves this, but is only available from 1971. Therefore, it was only possible to drop Hispanic individuals from the data for all years from 1971. Moreover, I drop military personnel (who do not report hours worked in any case), again following Mulligan and Rubinstein (2008) and Herrmann and Machado (2012).

As my analysis focuses on couples, I only keep observations which reported being either the household head, spouse, or unmarried partner. The latter category was only introduced in 1995. Before, it included ”partner/roommate”, which not necessarily constitutes a cohabiting romantic partner. Thus, identifying cohabiting couples (and subsequently distinguishing them from married ones) is only possible from 1995.

To keep the results comparable across households, I only consider households where the male partner is working for income (reporting non-zero earnings and hours worked). Subsequently, I restrict the sample on heterosexual couples by dropping couples consisting of same-sex partners.

Labor earnings are only reported as the wage and salary income of the previous calender year (*INCWAGE*). In order to compute hourly wages, I therefore need to adjust the yearly earnings with the weeks and usual weekly hours worked of the previous year. Before doing so, two problems of missing data need to be solved: Firstly, weeks worked in the year before the survey are only available accurately from 1976 on. Prior to 1976, weeks worked last year were reported in six intervals and must therefore be imputed. To do so, I follow United States Bureau of Labor Statistics (1993) and take the average by gender of each interval in 1976 (where discrete weeks were being reported for the first time), using the sample weight *ASECWT*. I then assign these gender-hour averages to each of the six intervals in 1964 to 1975.

The second case of missing data concerns the usual hours worked per week in the preceding calender year, which are also missing from 1964 to 1975. Again, I use the same method as United States Bureau of Labor Statistics (1993) and construct the missing work hours as the product of a dummy variable for having worked last year and the actual work hours reported for the week before the survey. For the case that an individual has worked last year (as indicated by the dummy variable), but has not worked the week before the survey, I impute the hours worked last year to be 40.

I construct hourly wages from the reported annual earnings by dividing *INCWAGE* by the weeks and usual weekly hours worked last year. For individuals who report earnings, but do not report weeks or hours worked, I impute weeks worked last year to be 52 and hours worked last year to be either hours worked last week or 40. As is common in the literature, for example in Herrmann and Machado (2012) and Attanasio, Battistin, and Ichimura (2007), I take a 1% sample of earnings and wages by using the sample weight to generate the percentiles for each year and dropping the top and bottom 1%. The reason for doing this is both to deal with outliers and earnings being top-coded. In addition, I deflate earnings and wages using CPI data obtained from United States Bureau of Labor Statistics (2019) and expressing them in constant 2010 U.S. Dollars. Household earnings are constructed as the sum of the trimmed and deflated earnings of both partners. I calculate annual hours worked as the product of weeks worked in the previous year and the hours worked in the week before the survey, because this is the only continually available hours measure in the CPS, as Ngai and Petrongolo (2017) note.

Regarding the skill level cut-off, I follow the definition of Cortes et al. (2018) and Greenwood et al. (2016) and consider individuals with at least four years of college as high-skilled and the remainder as low-skilled. Using four years of college rather than graduating from college is necessary because the education variable in the CPS (*EDUC*) is made up of the variables *HIGRADE* for all years prior to 1992 and *EDUC99* for the following years. Unlike *EDUC99*, *HIGRADE* only contains the highest grade the respondent has completed. As 4 years is the usual duration of an undergraduate degree in the U.S., I use at least 4 years of college education for defining high-skilled individuals.

To obtain the shares of the four different household types (u-u, s-u, u-s and s-s), I generate the proportion of each household type within each year, applying the household weight *ASECWTH*. I use the same procedure to obtain the share of married and cohabiting couples for each year. Table A.1 lists the number of individual observations left in the CPS data after imposing the aforementioned sample restrictions. The number of households in each year is exactly half of the number of observations due to the restriction on married and cohabiting heterosexual couples.

Table A.1: Number of Observations in the CPS Data per Year After Sample Restrictions

Year	Obs.	Year	Obs.
1964	15,646	1992	24,002
1965	15,422	1993	23,830
1966	30,872	1994	23,520
1967	19,730	1995	24,724
1968	30,086	1996	21,656
1969	30,452	1997	21,592
1970	29,012	1998	20,874
1971	23,706	1999	20,906
1972	22,122	2000	20,738
1973	20,752	2001	37,046
1974	21,310	2002	36,280
1975	22,234	2003	35,480
1976	22,126	2004	34,208
1977	26,000	2005	33,148
1978	24,740	2006	31,770
1979	24,358	2007	30,960
1980	28,396	2008	29,982
1981	28,274	2009	29,132
1982	24,486	2010	28,194
1983	24,402	2011	26,968
1984	24,842	2012	26,242
1985	25,302	2013	25,916
1986	24,930	2014	24,910
1987	24,800	2015	24,344
1988	25,438	2016	22,162
1989	24,284	2017	21,966
1990	25,554	2018	20,882
1991	24,700	**Total**	1,405,408

Source: Flood et al. 2018

A.2 AHTUS

The only variable needed for my quantitative analysis that is not included in the CPS is the amount of time spent in home production. I obtain this variable from the American Heritage Time Use Study (AHTUS) provided by IPUMS (Fisher et al. 2018). The IPUMS-AHTUS contains U.S. time use data for the years 1965, 1975, 1985, 1995, 1998 and 2003 to 2012, based on time-diary samples and harmonized by the Centre for Time Use Research(CTUR) at the University of Oxford. For each year, respondents were asked to track their time use over 24 hours in a diary and report their activities in minutes spent during those 24 hours. People could report multiple days per year, but unlike for the CPS, only one person per household reports their own time use.

Using the AHTUS-X data extract builder, I created a custom time use variable containing the sum of the reported unpaid domestic work (0400), child care (0500) and adult care (0640). This constitutes what I refer to as home production. I explicitly do not include time spent on personal care and sleep, as these do not generate mutual benefits to the household to the same extent as doing chores and taking care of children and adults.

Consistent with the CPS, I impose the same restrictions on age (25-55) and ethnicity (non-Hispanic whites). The only limitation here is the ethnicity variable *ETHNIC2* not being available in 1985, but given the scarcity of time use data I keep the 1985 data regardless of this issue. Furthermore, I restrict the sample on married and unmarried people, which is done through the variable *CIVSTAT*, where "married" includes both legally married and cohabiting people. I drop observations with missing education data as well as all men who are not working. Respondents also report the employment status of their partner in *EMPSP*, which allows me to also drop all women whose partner is not working.

For generating the average hours spent on home production for the different subgroups for each year, I apply the recommended sample weight *RCWGHT*. Same as in the CPS, I consider all individuals with at least four years of college education as high-skilled. Unfortunately, the AHTUS does not provide data on the partner's education. Therefore, due to only having one respondent per household, it is not possible to determine average household production hours by household type. Hence, home production hours will only differ by gender and own skill level, rather than the skill level of the partner.

Finally, I convert the weighted means of the home production time for each gender-skill combination from minutes per day to hours per week by dividing by 60 and multiplying by 7. The sample weights I applied automatically adjust for over- and undersampling of certain weekdays.

Table A.2 lists the number of record 24-hour time diary observations in the AHTUS data after the sample restrictions. These numbers are not equivalent to individuals, but to days recorded, as repondents could report multiple days in a year.

Table A.2: Number of Observations in the AHTUS Data per Year After Sample Restrictions

Year	Obs.
1965	1,059
1975	1,474
1985	1,073
1995	358
1998	870
2003	6,743
2004	4,549
2005	4,208
2006	4,012
2007	3,838
2008	3,776
2009	3,772
2010	3,585
2011	3,284
2012	3,297
Total	45,898

Source: Fisher et al. 2018

B Mathematical Appendix

B.1 Household Problem

The household optimization problem can be formulated as:

$$\max_{c_{i,j},l_{i,j,m},l_{i,j,f}} U(c_{i,j}, l_{i,j,m}, l_{i,j,f}) = \frac{c_{i,j}^{1-\eta}}{1-\eta} + \psi\frac{(1 - l_{i,j,m} - \overline{h_{i,j,m}})^{1-\sigma}}{1-\sigma} + \psi\frac{(1 - l_{i,j,f} - \overline{h_{i,j,f}})^{1-\sigma}}{1-\sigma}$$

$$s.t.\ c_{i,j} = w_i l_{i,j,m} + w_j l_{i,j,f}$$

$$d_{i,j,m} + d_{i,j,f} = 2 - l_{i,j,m} - l_{i,j,f} - \overline{h_{i,j,m}} - \overline{h_{i,j,f}}$$

As consumption is determined by the household's labor income, the optimization problem can be rewritten by substituting $c_{i,j}$ in the utility function for the budget constraint:

$$\begin{aligned}\max_{l_{i,j,m},l_{i,j,f}} U(l_{i,j,m}, l_{i,j,f}) =& \frac{(w_i l_{i,j,m} + w_j l_{i,j,f})^{1-\eta}}{1-\eta} + \psi\frac{(1 - l_{i,j,m} - \overline{h_{i,j,m}})^{1-\sigma}}{1-\sigma} \\ &+ \psi\frac{(1 - l_{i,j,f} - \overline{h_{i,j,f}})^{1-\sigma}}{1-\sigma}\end{aligned}$$

Since this leaves the household with only two choice variables ($l_{i,j,m}$ and $l_{i,j,f}$), solving the optimization problem yields two first-order conditions with respect to male and female labor supply:

$$\begin{aligned}(1) : \frac{\partial U}{\partial l_{i,j,m}} &= w_i(w_i l_{i,j,m} + w_j l_{i,j,f})^{-\eta} - \psi(1 - l_{i,j,m} - \overline{h_{i,j,m}})^{-\sigma} \overset{!}{=} 0 \\ (2) : \frac{\partial U}{\partial l_{i,j,f}} &= w_j(w_i l_{i,j,m} + w_j l_{i,j,f})^{-\eta} - \psi(1 - l_{i,j,f} - \overline{h_{i,j,f}})^{-\sigma} \overset{!}{=} 0\end{aligned}$$

By dividing $\frac{(1)}{(2)}$, I obtain the gender wage ratio as a function of both labor supplies:

$$\frac{(1)}{(2)} \Rightarrow \frac{w_i}{w_j} = \left(\frac{1 - l_{i,j,m} - \overline{h_{i,j,m}}}{1 - l_{i,j,f} - \overline{h_{i,j,f}}}\right)^{-\sigma}$$

From this equation, I can then derive individual female labor supply as a function of both partners' wages, household hours, and male labor supply:

$$\begin{aligned}
\left(\frac{w_i}{w_j}\right)^{-\frac{1}{\sigma}} &= \frac{1 - l_{i,j,m} - \overline{h_{i,j,m}}}{1 - l_{i,j,f} - \overline{h_{i,j,f}}} \\
\Leftrightarrow \left(\frac{w_i}{w_j}\right)^{\frac{1}{\sigma}} &= \frac{1 - l_{i,j,f} - \overline{h_{i,j,f}}}{1 - l_{i,j,m} - \overline{h_{i,j,m}}} \\
\Leftrightarrow (1 - l_{i,j,m} - \overline{h_{i,j,m}}) \left(\frac{w_i}{w_j}\right)^{\frac{1}{\sigma}} &= 1 - l_{i,j,f} - \overline{h_{i,j,f}} \\
\Leftrightarrow l_{i,j,f} &= 1 - \overline{h_{i,j,f}} - (1 - l_{i,j,m} - \overline{h_{i,j,m}}) \left(\frac{w_i}{w_j}\right)^{\frac{1}{\sigma}}
\end{aligned}$$

B.2 Firm Problem

The firm maximizes profits subject to the production function and the technology constraint:

$$\begin{aligned}
\max_{A_u, A_s, L_u, L_s, k} \pi(A_u, A_s, L_u, L_s, k) &= y - w_u L_u - w_s L_s - rk \\
s.t.\ y &= k^{\alpha} \left[(A_u L_u)^{\theta} + (A_s L_s)^{\theta}\right]^{\frac{1-\alpha}{\theta}} \\
(A_s)^{\omega} + \gamma (A_u)^{\omega} &\leq B
\end{aligned}$$

This can be solved by substituting the production function into the profit function for y, and then setting up the Lagrangian with the second (technology set) constraint:

$$\mathcal{L} = k^{\alpha} \left[(A_u L_u)^{\theta} + (A_s L_s)^{\theta}\right]^{\frac{1-\alpha}{\theta}} - w_u L_u - w_s L_s - rk + \lambda \left[B - (A_s)^{\omega} - \gamma (A_u)^{\omega}\right]$$

The resulting first-order conditions (excluding the one for capital, which is irrelevant for my analysis) are:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial L_u} &= k^{\alpha} \frac{1-\alpha}{\theta} \left[(A_u L_u)^{\theta} + (A_s L_s)^{\theta}\right]^{\frac{1-\alpha-\theta}{\theta}} \theta A_u^{\theta} L_u^{\theta-1} - w_u \overset{!}{=} 0 \\
\frac{\partial \mathcal{L}}{\partial L_s} &= k^{\alpha} \frac{1-\alpha}{\theta} \left[(A_u L_u)^{\theta} + (A_s L_s)^{\theta}\right]^{\frac{1-\alpha-\theta}{\theta}} \theta A_s^{\theta} L_s^{\theta-1} - w_s \overset{!}{=} 0 \\
\frac{\partial \mathcal{L}}{\partial A_u} &= k^{\alpha} \frac{1-\alpha}{\theta} \left[(A_u L_u)^{\theta} + (A_s L_s)^{\theta}\right]^{\frac{1-\alpha-\theta}{\theta}} \theta A_u^{\theta-1} L_u^{\theta} - \lambda \gamma \omega (A_u)^{\omega-1} \overset{!}{=} 0 \\
\frac{\partial \mathcal{L}}{\partial A_s} &= k^{\alpha} \frac{1-\alpha}{\theta} \left[(A_u L_u)^{\theta} + (A_s L_s)^{\theta}\right]^{\frac{1-\alpha-\theta}{\theta}} \theta A_s^{\theta-1} L_s^{\theta} - \lambda \omega (A_s)^{\omega-1} \overset{!}{=} 0
\end{aligned}$$

Dividing the second by the first and simplifying yields the firm's first optimality condition, pinning down the skill premium:

$$\frac{w_s}{w_u} = \left(\frac{A_s}{A_u}\right)^{\theta} \left(\frac{L_s}{L_u}\right)^{\theta-1}$$

Likewise, the second optimality condition can be obtained by dividing the fourth by the third first-order condition and simplifying:

$$\left(\frac{A_s}{A_u}\right)^{\omega-\theta} = \gamma \left(\frac{L_s}{L_u}\right)^{\theta}$$

C Codes

I have used Stata to clean the CPS and AHTUS data, both for the chapter on empirical facts and to obtain the average labor supplies, household hours, wages and population shares by gender and skill level to be put into my model. Solving the model, testing the female labor supply candidates, and the counterfactual analysis of married and cohabiting households was done in MATLAB. Both the Stata and MATLAB codes are available upon request.