Rising Income Inequality and the Fertility of Highly Skilled Women

DongIk Kang* Youngjae Lee†

September 3, 2024

Abstract

The relative fertility rates of high-income females in the U.S. have risen over the last few decades. We argue that this increase is a consequence of the rise in income inequality during this period. When households have heterogeneous preferences for children, greater returns to skill disproportionately increases both the fertility and income of highly skilled women who desire children, by allowing them to have more babies while also earning higher income. Using a quantitative model that incorporates this channel, we find that almost all of the changes in the income-fertility relationship that are observed in the data can be attributed to increased income inequality.

JEL Codes: D31, J13, J24

Keywords: Income inequality, Fertility rate, Fertility preference

*School of Economics, Yonsei University. Email: dikang@yonsei.ac.kr

†Corresponding author. Kyungpook National University. Email: leeyoungjae@knu.ac.kr

1 Introduction

An increase in income inequality raises the opportunity cost of having children for highly skilled women, as children often limit how much a mother can work. As such, a concern associated with rising income inequality in the U.S. was that it would widen the fertility differences across households by reducing the fertility of highly skilled women.[1] The worry lies in the fact that poor families have more children but, on average, invest less in their children. Then, as average education levels fall with greater income-fertility differentials, there could be a negative impact on growth (De la Croix and Doepke (2003), Moav (2005)). However, recent studies by Hazan and Zoabi (2015a) and Bar et al. (2018) have shown that the rise in income inequality in the U.S. has had the opposite effect on fertility differences across households; highly educated and high-income women have increased their fertility in relative terms.

In this paper, we propose a novel explanation for this reduction in the income-fertility differential. We find that when households have heterogeneous preferences for children, wage increases for highly skilled women allows females with both high skill *and* preference for children to not only have more babies but also greater income. With lower wages, loosely speaking, highly skilled but child-loving women must choose between either working or looking after their children. Those of these women who choose to have more babies are assigned to lower income groups while those who choose to work more have low fertility. However, higher wages alleviates this trade-off between work and childcare. It enables these women to earn more while working less, and in turn to reallocate this newfound time towards having more children. Thus, from the perspective of the overall economy, highly skilled women who like and have more children become a part of the high-income group replacing less skilled women who do not like (or have) children who instead invest more time toward their

[1]There is a negative relationship between income and fertility across households that has been well established both theoretically and empirically in the literature (Jones and Tertilt, 2008).

work. As a result the relative fertility rates of high-income households increase. We show that this change in the composition of high-income households, together with the relative fall in the market price for childcare services that income inequality entails,[2] can account for almost the entirety of the reduction in fertility differences across income quantiles, as the consequence of rising income inequality over the past few decades.

We document three empirical facts that motivate our results. First are the well-known changes to the income structure in the United States, where the skill premium (measured as returns to education) increases and the gender wage gap decreases. Secondly, closely related to the results of Hazan and Zoabi (2015a) and Bar et al. (2018), we document the changes in the income-fertility relationship in the U.S. between 1990 and 2018. In particular, we demonstrate that the fertility rates of women in the top twentieth percentile in terms of female labor income have increased in relative terms. Thirdly, we show that the relationship between income and preference for children has changed in a manner similar to the income-fertility relationship. We find that the average number of children desired by households in the top twentieth percentile in terms of female labor income have also increased relative to the average over time. To the best of our knowledge, this third finding is novel to our paper. Then using a toy-model, we consider the effects of the change in income structure, as an increase in the wage for a unit of productivity, and demonstrate that the aforementioned changes in the composition of households lead to an increase in the relative fertility rates and fertility preferences of high-income households, tying the three facts together.

To quantify these effects, we build a quantitative life-cycle model in which there is household heterogeneity in the intrinsic cost of education, labor productivity, and affinity for children. In our model, going to college increases the expected level of labor productivity,

[2]See Hazan and Zoabi (2015a) and Bar et al. (2018) for a detailed discussion on the fall of the market price for childcare and its effect on fertility. They argue that by substituting market time for their own time in raising children and running their households, highly skilled women are able to have more children and work longer hours as the relative price of childcare falls.

which raises future income, while having more children requires greater time commitment by the female members of the household, potentially lowering their income. Balancing these concerns, the households in our model decide how much education to obtain, how much to work and consume, how many children to have, how much to invest in the quality of their children, and how much to save for retirement, as they move through the different phases of their life-cycle. Our model also incorporates the marketization mechanism introduced in Bar et al. (2018) by allowing female members of the household to substitute their personal time with market services in the care for their children. The calibrated model does a good job of matching key targeted and non-targeted moments of the U.S. economy.

Using this model, we quantify how much of the observed changes in the income-fertility patterns can be explained by the rise in income inequality. We change the income structure in the model to match the observed changes in the skill premium and the gender wage gap and study how the economy changes. We find that the model can account for almost all of the observed changes in the income-fertility relationship under reasonable calibrations. The model demonstrates that the fertility level of college graduates increases as well. In addition, the model also predicts that the relationship between income and investment into child quality will become steeper. Households with high female labor income will increase their investment into child quality while lower income households decrease their investment.

In comparison, we find that without accounting for the change in the composition of high-income households, the model cannot accurately account for the observed changes in the income-fertility relationship as a response to the increases in income inequality. When we fix the composition of households in each income group by household income levels preceding the changes in the income structure, the model severely under-predicts both the relative increase in the fertility of high-income women and the relative decrease in the fertility rates women between the 20th and the 60th percentile in terms of income. Furthermore,

the presence of market services and the decrease in its relative price, as suggested by Bar et al. (2018), also plays an important role in our model. When we further restrict the model and fix the relative price of outsourcing child care for each household to their values before the rise in income inequality, we find that the fertility rates of households of all income groups fall significantly with the applied changes in the income structure. The effects of the changes in the composition of high-income households and better access to market services are complementary, and together their combined effects are able to almost completely account for the changing patterns in the income-fertility relationship.

As previously mentioned, two closely related but conceptually distinct phenomena have contributed to the rising income inequality across households over the last 40 years. The rise in the skill premium has increased income inequality between high-skilled and low-skilled workers, and the fall in the gender wage gap has contributed to household income inequality by widening the earnings differences among women in particular. Determining which phenomenon is the primary driver of the decreasing fertility differentials could have potentially important implications with regard to policy. Thus, analyzing the effect of the rise in the skill premium versus the reduction in the gender wage gap separately, we find that both are major contributors but in different ways.

We find that the change in the skill premium accounts for 65% of the total variation of the full model. Its rise incentivizes young women with greater affinity toward children to attend college not only to capture the effect of increased returns to skill but also because college-educated women are more likely to be matched to college-educated men, and more income from their partners allows them to have more children. When sorting households according to female labor income, this leads to increased fertility rates for households with not only the highest female income levels but also the lowest, as college-educated women who do not work but are matched with high-income partners also see an increase in their number

of kids. In contrast, the narrowing gender wage gap strictly favors high-income females with regard to fertility, as its effect can only be captured by working women. We find that as the gender wage gap narrows, high-income women are incentivized to have more children while women with low income have fewer children. Taken together, the rise in the skill premium and the fall of the gender wage gap work in the same direction to increase the fertility rates of high-income women while their effects counteract each other for the lowest income women. For women in the middle income groups, the skill premium effects dominate, meaning their relative fertility rates fall.

This study is closely related to the literature that studies the changing fertility behavior of highly educated and high-income women. Recent studies find that highly educated and high-income women are having more children, which is surprising given the fact that the opportunity cost of having children in terms of female labor income is increasing.[3] Shang and Weinberg (2013) conjectures that increases in the use of personal services may have reduced the burden of having children. Bar et al. (2018) and Hazan and Zoabi (2015a) argue that the fall in the relative price of market services that can be used to outsource childcare has changed the relationship between the mother's education level and fertility rates. Siegel (2017) argues that fertility rates may be sustained by husbands spending more time at home as women's wages increase. These studies revolve around a common idea that the increasing prevalence of substitutes for the mother's time in childcare allows highly skilled women to have more babies. Our paper differs in that the mechanism we propose concerns heterogeneity in the preferences and productivity across workers and not the rise of substitutes for a woman's time.

We also contribute to the literature on the relationship between fertility rates and preferences. Miller et al. (2010) show empirically that having greater fertility motivations actually

[3]Please refer to Doepke et al. (2023) for a detailed discussion of recent empirical findings on changing fertility patterns, as well as an extensive survey of models on fertility choice.

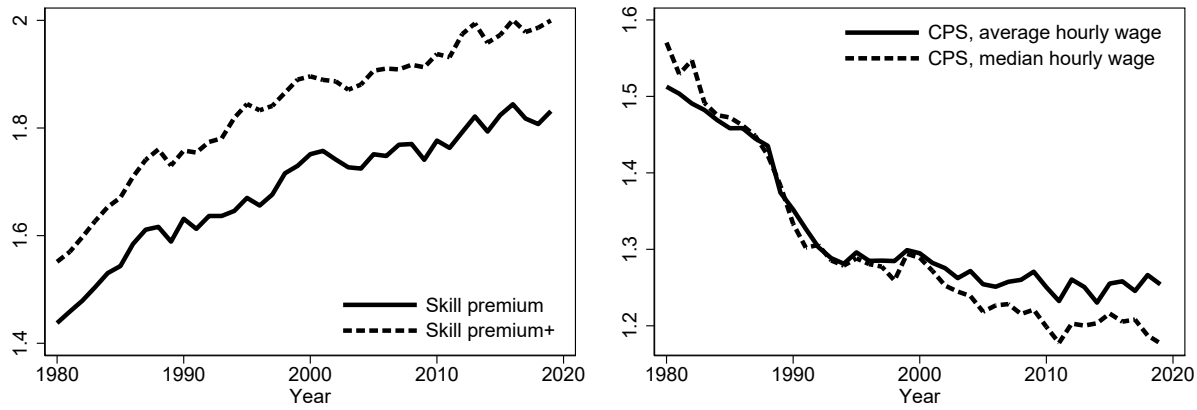
increases the number of children women have. Jones et al. (2008) suggests that heterogeneity in the desire for children may contribute to a negative income-fertility relationship because families with greater affinity for children will invest less in developing their skills. Gobbi (2013) argues that intergenerational transmission of preferences is key to understanding why decreasing fertility rates do not lead to more childlessness. In the field of sociology, Hakim (2000, 2003) argues that preference per se is the most crucial factor in determining a woman's fertility behavior in modern and developed societies.

Lastly, our mechanism is related to the self-selection framework found in the seminal works of Roy (1951) and Borjas (1987), who study the job selection problem of agents with different skill sets. Using this framework, Burtless and Hausman (1978) and Blomquist (1983) introduced preference heterogeneity to investigate the effect of taxation on the labor supply. Their work highlights the effects of differences in the magnitudes of the income and substitution effects between different preference groups. In our framework, households balance earnings ability and the preference for children to select the optimal fertility and income levels.

The rest of this study is organized as follows. In Section 2, we present three empirical facts that motivate our channel and discuss how they are connected. In Section 3 we present our quantitative model, which we use to discuss the effects of rising income inequality on the income-fertility relationship in Section 4. Section 5 concludes.

2 Empirical Facts

In this section, we document three facts about changes in income, preferences and fertility rates and discuss how these three facts are related to our channel using a toy model.



(a) Skill premium (b) Gender wage gap

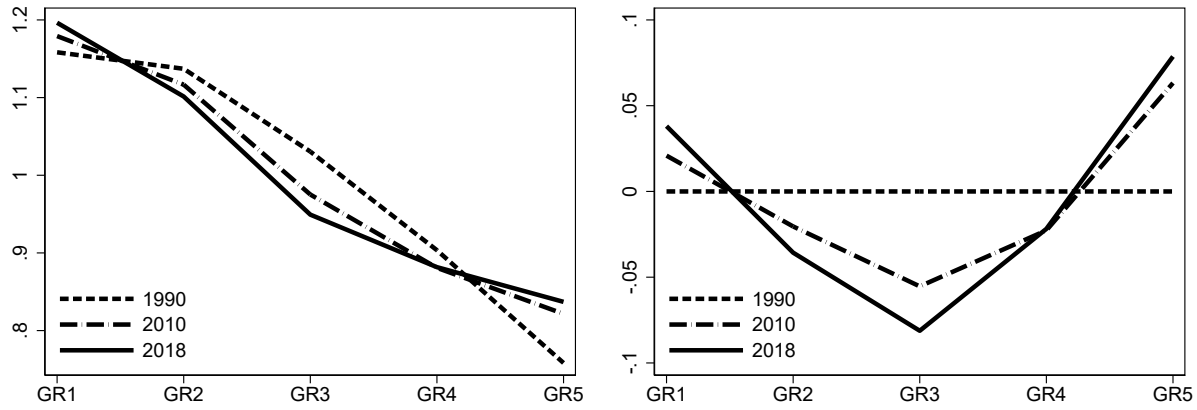
Figure 1: Changes in Income Structure

Notes: In panel (a), the skill premium is defined as the ratio of the hourly wage of people who hold a bachelor's degree (and not more advanced degrees) to the wages of high school graduates. The wage premium between workers with a bachelor's degree or more and workers with high school degrees or less is denoted as skill premium+. In panel (b), the gender wage gap is defined as the ratio of the average or median hourly wage of men to that of women.

2.1 Changes in Income, Preferences, and Fertility Rates

First, we present the well known facts about the changes in the skill premium and the gender wage gap from 1980 to 2020. Panel (a) of Figure 1 illustrates how the skill premium of college graduates over those with only a high school education has changed. The skill premium, defined as the ratio of the hourly wage of people who hold a bachelor's degree (and not more advanced degrees) to the wages of high school graduates, has increased from around 1.4 in 1980 to around 1.8 in 2020. The wage premium between workers with a bachelor's degree or more and workers with high school degrees or less, denoted as skill premium+, has increased from below 1.6 in 1980 to approximately 2 in 2020. In comparison, the gender wage gap has narrowed from about 1.5 to 1.25 in terms of the average hourly wage, and from just below 1.6 to 1.2 in terms of the median hourly wage.

Secondly, as first documented by Hazan and Zoabi (2015a) and Bar et al.(2018), we illustrate how the income-fertility relationship has changed during the past few decades.



(a) Comparing CFRs by female income quintiles (b) Changes in the normalized CFR

Figure 2: Changes in Fertility

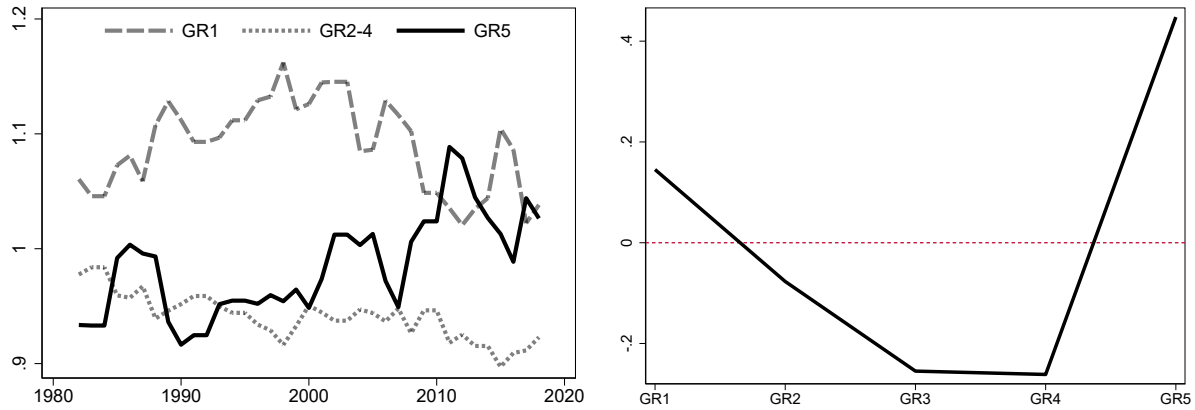
Notes: Panel (a) shows the normalized average completed fertility rate (CFR) by income quintile for the years 1990, 2010, 2018. Average CFR for each income group is normalized by dividing the average CFR within the income group by the average CFR across all women for each year. Panel (b) plots the change in the normalized CFR using the year 1990 as the benchmark.

We compute the average completed fertility rate (henceforth CFR) for each female labor income quintile, using data from the U.S. Census and the American Community Survey (henceforth ACS).[4] Following Autor et al. (2008), we restrict our attention to married white women between the ages of 36 and 40, with husbands that work at least 35 hours a week and 40 weeks a year.[5] Panel (a) of Figure 2 plots the normalized average CFRs for each income quintile in 1990, 2010, and 2018, respectively. We normalize by dividing the average CFRs within income group by the average CFR across all women for each year.[6] The relationship between income and fertility is strongly negative in all three years, as has been well documented. However, comparing across years, it becomes apparent that the fertility

[4]See Appendix A for a discussion of the data and the construction of the variables.

[5]We focus on married white women as our mechanism is focused on the changes in the fertility of married households, and out-of-wedlock fertility is relatively low for white women compared to other races (see Wildsmith et al. (2018)).

[6]In Appendix B.1, we show that the general patterns we find in the income-fertility relationship remains robust when we use raw (not normalized) CFRs, focus on women between the ages of 46 and 50, or include other races.



(a) Changes in fertility preferences (GSS) (b) Changes in fertility preferences (NLSY79)

Figure 3: Changes in the Preference to Have Children

Panel (a) shows the moving average of fertility preferences within a 11-year window (current and past ten years) for the lowest income, the middle-income, and the highest income quintiles using data from the GSS. Panel (b) illustrates the changes in the income-fertility preference relationship using the survey data from the NLSY79. It shows the change in the responses between the cohort born in 1958-59 to the cohort born in 1964-65, made when each cohort was between the ages 17 and 21. The income group to which each respondent belongs is determined by their average labor income between the ages of 37 to 48.

rates in the lowest and highest income quintiles have increased over the years, while that of the middle income quintiles have fallen in relative terms. Panel (b) in Figure 2 highlights these changes over time by plotting the change in the normalized CFR in relation to the values in 1990. Thus the CFR values in the year 1990 serves as the benchmark. Looking at the relative changes, a U-shaped difference curve emerges showing that the relative fertility rates of the highest and lowest income groups increased over time, while relative CFRs in the middle of the income distribution fell. Between 1990 and 2018, the CFR of the lowest income quintile increased by close to 5 percentage points relative to the average rate and that of the highest income quintile increased by 8 percentage points relative to the average fertility rate. Households in the second to fourth income quintiles saw their relative CFR drop.

Finally, we provide evidence showing that the relationship between female labor income

and fertility preferences has changed over the years, in a manner similar to the observed changes in the income-fertility relationship. To the best of our knowledge, this finding is novel to our paper. The General Social Survey (GSS) includes a question that asks "What do you think is the ideal number of children for a family to have?" We use the answers to this question, which we interpret as representing the household's desired number of children, to compute the household preferences. However, because the number of married white women are much smaller in the GSS data, we expand the age range by 5 years and include women between the ages of 31 and 45. We normalize the responses by dividing by the average across all households for that year and then compute the average values for each income group and year pair. In panel (a) of Figure 3 we plot the moving average of this variable for the lowest income quintile, the middle income group consisting of quintiles 2 through 4, and the highest income quintile.[7] It shows that the relative preference for children in women with low labor income remains high throughout the entire analysis period from 1982 to 2018, while that of the highest income group steadily rises and the middle income group steadily falls.

It has been noted by some that the GSS survey question that we use may not be entirely capturing the *personal* preferences of the respondents.[8] It may be possible that the respondent interprets the GSS question as asking for the "socially" ideal number of children, and answer that the social ideally number of children is, for example, two even though they may not want to have any babies at all. To alleviate such concerns, we bolster our analysis using the National Longitudinal Survey of Youth 1979 (NLSY79) data which asks "How many children do you want to have altogether?" and "What do you think is the ideal number of children for a family?" as separate questions. Following the definition from the OECD Family Database, SF 2.2, we interpret the first question as pertaining to personal fertility

[7]The moving average is over a 11 year window including the current and 10 prior years.

[8]See Blake (1966), Philipov and Bernardi (2011), Sobotka and Beaujouan (2014) and others for a discussion regarding the interpretation of survey results regarding fertility preferences.

preferences and the second as to general fertility preferences. Thus, we focus on the first.[9] As before, we normalize the households' answers regarding their desired number of children by dividing by the average across all households for that year.

In panel (b) of Figure 3, we show how the income-fertility preference relationship has changed with time. We compare the responses of the cohort born in 1958-59 to the cohort born in 1964-65, made when the respondents were between the ages of 17 and 21. The income group to which each respondent belongs is determined by their average labor income between the ages of 37 and 48. We subtract the average preference by income group of the 1958-59 cohort from the average preference of the 1964-65 cohort, which is what we show. As can be seen, we find that the average fertility preference of women with high labor income is greater for the 1964-65 cohort compared to the 1958-59 cohort, whereas the fertility preferences of the middle income households fell. The fertility preference of the the lowest income households in terms of female labor income remained relatively steady. These findings are similar to those in panel (a).

2.2 Interpreting the Facts

In the previous section we presented three facts related to the changes in income, preferences, and fertility in the United States during the past 40 years. We now discuss how these facts are related. To make things clear, we present a simple toy-model illustrating the key mechanism linking the rise in the wages of high-skilled women to the observed changes in the income-fertility relationship and the income-preference relationship.

Suppose a household consisting of a single female worker derives utility from consumption c and from their children, proportional to the number of offspring they have n. Each household derives a different level of utility from their children, captured by ζ, and are

[9]However, we find that the changes in the general fertility preferences, reported in Appendix B.2, shows a similar pattern to the changes in the personal fertility preference shown in Figure 3b.

endowed with (heterogeneous) labor productivity ϵ. For the sake of simplicity, we assume that these idiosyncratic characteristics of the households are uniformly distributed in the set $\big\{(\epsilon, \zeta) \mid \epsilon, \zeta \in [0, 1]\big\}$. Without any children, the household works for a unit time, while each child reduces a woman's working hours by time t. Thus given the wage for a unit of efficient labor w, household income can be computed as $w\epsilon(1 - tn)$. The household solves,

$$\max_{c,n} \quad u = c + \zeta \log n \qquad s.t \qquad c = w\epsilon(1 - tn) \geq 0,$$

resulting in the optimal number of children

$$n^* = \min\Big(\frac{\zeta}{w\epsilon t}, \frac{1}{t}\Big). \tag{1}$$

The optimal number of children decreases in wage, productivity, and the time cost of a child, all of which increases the opportunity of cost of having a child.

Panel (a) of Figure 4 shows the iso-income curves of the households in (ϵ, ζ) space. The iso-income curve of households can be derived by substituting in the optimal number of children in equation (1) into the expression for household income $w\epsilon(1 - tn)$ to get,

$$\zeta = w\epsilon - I,$$

where I is the income level of the household. As shown in panel (a), the income of households increases as the iso-income curves shift down and to the right in (ϵ, ζ) space. This is because household income is increasing in labor productivity ϵ while it decreases in the number of children, which is positively related to ζ. Households in the shaded area all have the income level $I = 0$. In panel (b), we show the iso-fertility curves of the households given by equation (1). The fertility of households increases as the slope of the iso-fertility curve increases around the origin, as households must trade off their preference for children ζ with their

opportunist cost from missed labor $w\epsilon$, requiring higher preference levels to maintain the number of kids as labor productivity increases.

Panel (c) of Figure 4 illustrates our main point. We assume $w \geq 1$ and focus on the comparison between households above and below the median income, which greatly simplifies the analysis. Suppose that the wage increases from w to w' ($> w$). This implies that the slope of the iso-income curve at the income level of the median household must increase and its ζ-intercept (equal to negative income) must fall. The two iso-income curves of median income at wage levels w and w' will cross at the mid-point of the support $(0.5, 0.5)$. Then, we can see that the households in the shaded area labeled A will now be above median income at the new wage level w' whereas their income was below median at wage w. Similarly, households in the shaded area labeled B will be below the median income level at wage w' while they were above the median income level at wage w. This change in the composition of households between high and low income groups occurs because an increase in the wage implies an increase in the returns to productivity, meaning households with higher productivity but lower labor supply (due to higher ζ) gain income in relative terms. Due to the way we have constructed this problem, the areas A and B will be of the exact same size.

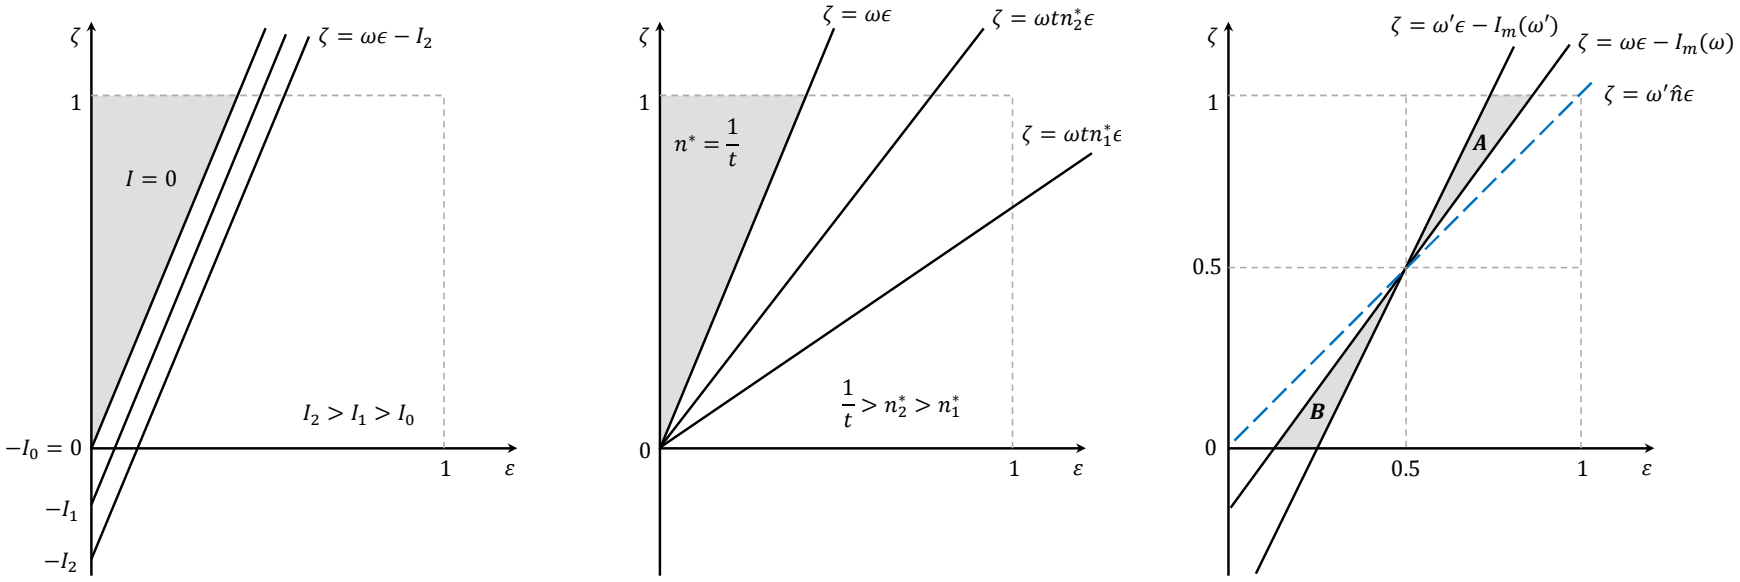
Then, as the households constituting area A have a strictly higher ζ than those in area B, this change in the composition of households will imply that the average value of ζ in the high income group will increase and that of the low income group will decrease. This change in average preference reflects the changes we observed in the empirical analysis shown in Figure 3. The effect of the composition change on realized fertility isn't as straightforward, as a higher wage also increases the opportunity cost of a child. However, we can show unambiguously that, in relative terms, the fertility rate of high income households will increase while the fertility rate of low income households will decrease compared to the mean. The blue line shows the iso-fertility curve at the new wage level w' for fertility level $\hat{n} = \frac{1}{w't}$.

Recall that the fertility level increases as the slope of the iso-fertility curve increases. Thus, we can see that households in area A will always have greater fertility than households in area B given the new wage w', implying the change in composition will increase the relative fertility of households in the high income group. This corresponds to the empirical results shown in Figure 2.

Intuitively, an increase in the wage increases the returns to labor productivity. When wages for highly productive women increase, they can now earn more than less productive women, even if they work fewer hours in order to take care of their children. Thus more households with relatively high productivity *and* relatively high fertility preferences now constitute a larger fraction of high income households, than was previously the case when returns to productivity was lower.[10]

While the simple model presented in this section successfully captures the key mechanism of our framework, it is nevertheless a toy model in which some of the analysis is dependent on adhoc assumptions made for the sake of clarity, such as the assumption of a uniform distribution of households in (ϵ, ζ) space. Therefore, in the following section we present a more realistic quantitative model that incorporates the main mechanism of the toy model as well as other potentially important socio-economic features, and analyze the quantitative importance of the channels presented here.

[10]Note that this effect disappears without preference heterogeneity. For example, if the only dimension of heterogeneity across households was from productivity ϵ, and preference ζ was homogeneous, the relative rank of households in terms of income would be determined completely by productivity. Fertility rates will be decreasing in income (productivity) and identical for households with the same income. In this case, an increase in the wage would not alter the composition of households in different income groups.



(a) Iso-income curves (b) Iso-fertility curves (c) The effect of wage increases

Figure 4: Toy Model

Notes: Panel (a) shows the iso-income curves of households in (ϵ, ζ) space. Panel (b) depicts the iso-fertility curves of the households. Panel (c) illustrates the change in the composition of households below median income and above median income groups, in response to an increase in the wage.

3 Quantitative Model

We now present a dynamic general equilibrium model that incorporates heterogeneity in labor productivity and fertility preferences, building upon the discussions of the preceding sections. We construct an overlapping generations model that builds on the home production model proposed by Greenwood et al. (2005). The model consists of households who supply labor, own capital, and make fertility decisions as well as a representative firm that produces using standard constant returns-to-scale technology. There is no aggregate uncertainty in the model. In addition to the heterogeneity across households, the model also incorporates the quality-quantity trade-off first introduced by Becker (1960) and allows households to use market services to outsource a portion of their childcare as first introduced by Hazan and Zoabi (2015a). We utilize our quantitative model to analyze how the composition changes to income groups contribute to the changing income-fertility relationships seen in the data and examine how it relates to previously established mechanisms in the literature.

3.1 The Household

There are a continuum of households of a unit measure, each of whom live for four periods $\{0, 1, 2, 3\}$. A period in our model corresponds to approximately twenty years in reality. In period 0, a household consisting of a single woman is born and makes their education choice of whether or not to attend college. At the start of period 1, the woman is randomly matched to a man and together they determine how many children they wish to have and how much to work. This is the only period in which the household can decide to have children. In the following period, both members of the household work and they jointly determine how much to consume and save. In period 3, both man and woman retire and consume their savings until they die at the end of the period. There is no divorce in the model. Figure 5 illustrates the timeline of events.

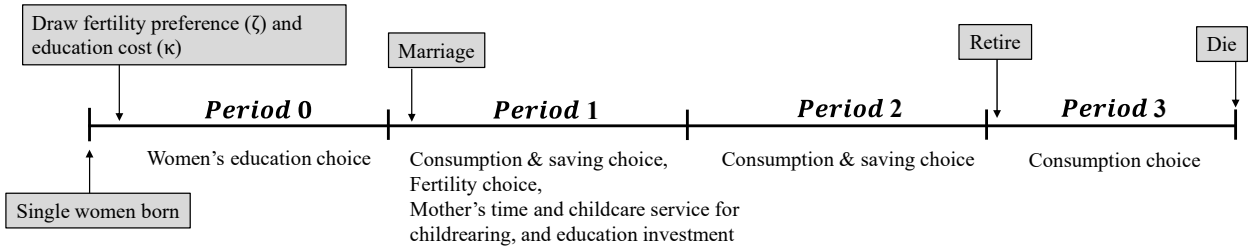


Figure 5: Timeline of Events

Period Zero

At the start of period 0, a household consisting of a single woman is born. At birth, they randomly draw their preference for having children ζ and their education cost κ from the joint distribution:

$$\begin{pmatrix} \zeta \\ \log \kappa \end{pmatrix} \sim \mathbf{N}\left(\begin{bmatrix} \mu_\zeta \\ \mu_{\log \kappa} \end{bmatrix}, \begin{bmatrix} \sigma_\zeta^2 & \sigma_{\zeta,\log \kappa} \\ \sigma_{\log \kappa,\zeta} & \sigma_{\log \kappa}^2 \end{bmatrix} \right). \tag{2}$$

We think of κ as representing intellectual capabilities or personal (non-monetary) valuation of educational opportunities, and ζ represents differences in the innate preference for children.

Given the realized values of (ζ, κ), the household decides whether or not to go to college. She can choose to pay the fixed cost κ to obtain a college degree or receive a high school education at no cost. The expected utility of the household from receiving a high school education only is,

$$V_0^h(\zeta) = \pi_{hh} \cdot \beta \mathbb{E}_0 V_1^{hh}(\zeta) + \pi_{hc} \cdot \beta \mathbb{E}_0 V_1^{hc}(\zeta), \tag{3}$$

while going to college would yield,

$$V_0^c(\zeta, \kappa) = \pi_{ch} \cdot \beta \mathbb{E}_0 V_1^{ch}(\zeta) + \pi_{cc} \cdot \beta \mathbb{E}_0 V_1^{cc}(\zeta) - \kappa, \tag{4}$$

where π_{ij} denotes the probability a woman with education level $i \in \{h, c\}$ marries a man with education level $j \in \{h, c\}$ where h indicates a high school graduate and c indicates a college graduate $V_1^{ij}(\zeta)$ is the expected utility of a household with a woman with education level i and man with j in period one. The conditional expectation $\mathbb{E}_0$ is over the realization of the labor productivity of the household in the following periods. Thus in period zero the household solves:

$$V_0(\zeta, \kappa) = \max\left(V_0^h(\zeta), V_0^c(\zeta, \kappa)\right). \tag{5}$$

Note that we only consider the educational choice of the woman and the educational status of the man is determined by marriage. Nevertheless, we attempt to capture the education decisions of men as well as the consequences of assortative matching by calibrating the exogenous matching probabilities π_{ij} to match the observed matching probabilities in the data.[11] The education levels do not change once they are set.

Period One

At the start of period one, a household consisting of a single woman is randomly matched with a man who's education level is governed by the aforementioned probabilities $\pi_{ij}, i, j \in \{h, c\}$. After being matched, the man and woman each draw their period one labor productivity $(\epsilon_1^f, \epsilon_1^m)$ from the distributions:

$$\begin{aligned} \epsilon_1^f, \epsilon_1^m &\sim N\left(0, \left(\frac{\sigma_{\nu,h}}{1-\rho_h}\right)^2\right) \qquad \text{if high school graduate,} \\ &\sim N\left(0, \left(\frac{\sigma_{\nu,c}}{1-\rho_c}\right)^2\right) \qquad \text{if college graduate.} \end{aligned}$$

[11] We also note that the matching probabilities π_{ij} attempt to indirectly incorporate the consequences of various factors that affect the educational choices of men and women including how changing economic structures affect female educational choice, which has been extensively explored in the literature by Galor and Weil (1996), Sauré and Zoabi (2014), and Hazan and Zoabi (2015b) among others.

The parameters $(\rho_h, \rho_c, \sigma_{\nu,h}, \sigma_{\nu,c})$ are the auto-regressive coefficients and the across-time volatility of labor productivities, to be shown below.

The household in period one solves

$$V_1^{ij}(\epsilon_1^f, \epsilon_1^m; \zeta) = \max_{c_1, a_2, n, q, t_f, t_b} \ln(c_1) + \zeta \ln(nq) + \beta \mathbb{E} V_2^{ij}(a_2, \epsilon_2^f, \epsilon_2^m; \zeta) \tag{6}$$

subject to the budget constraint

$$c_1 + \varphi(t_b) + \vartheta(q)n + a_2 = w_f^i \epsilon_1^f (1 - t_f) + w_m^j \epsilon_1^m, \tag{7}$$

the production function for children

$$n = Z \left[\phi (t_f)^\gamma + (1 - \phi) (t_b)^\gamma \right]^{1/\gamma}, \tag{8}$$

and the law of motions governing the productivity processes

$$ln(\epsilon_2^f) = \rho_i ln(\epsilon_1^f) + \nu_i, \qquad \nu_i \sim N(0, \sigma_{\nu,i}^2) \tag{9}$$

$$ln(\epsilon_2^m) = \rho_j ln(\epsilon_1^m) + \nu_j, \qquad \nu_j \sim N(0, \sigma_{\nu,j}^2) \tag{10}$$

where c denotes consumption, a is the household's assets, n is the number of children the household chooses to have, t_f and t_b is the amount of time devoted to the children by the mother and from market services (b stands for babysitter) respectively, and w_f^i and w_m^j denotes the market wage of female and male workers with education levels $i, j \in \{h, c\}$ respectively. Parameters (Z, ϕ, γ) govern the process through which investment into child care determines the number of children the household has. The function $\varphi(t_b)$ denotes the household's cost of buying a unit of t_b from the market and $\vartheta(q)$ denotes the per child cost of increasing child quality. Both costs are determined in the competitive market for child

care services introduced below.[12]

Equation (6) shows that the household's utility increases in consumption as well as in the number of children n. The degree to which an additional child adds to the household's utility is determined by the quality of the child q, as well as by preference ζ. The production function for children in equation (8) indicates that the number of children the household ultimately chooses to have is a direct outcome of their willingness to invest in the care of children as a function of time t_f and t_b. However, market services are an imperfect substitute for the mother's time, and the degree of substitutability is determined by the parameters ϕ and γ. The household's labor productivities evolve each period according to the processes (9) and (10).

While having children raises the utility of the household, it comes with various costs. Time invested in the child t_f decreases female labor income, as the mother must sacrifice her time to take care of her children. Outsourcing childcare to market services is costly. Furthermore the household must invest in the quality of their children. The household balances these concerns to make the optimal decisions regarding consumption, savings, and investment into children.

We solve the household's optimization problem and get the first-order conditions:[13]

$$(1+\zeta)c_1^* = w_f^i \epsilon_1^f + w_m^j \epsilon_1^m - a_2, \tag{11}$$

$$\zeta c_1^* = w_f^i \epsilon_1^f t_f + \varphi(t_b) + \vartheta(q) n^*, \tag{12}$$

where c^* and n^* denotes the optimal values of consumption and number of children, respec-

[12] We follow the modeling choices of studies such as Sommer (2016) and Park (2018), in that investment into child quality by the parent is not linked to the innate educational ability of the next generation captured by κ. As in previous studies, we omit this inter-generational link for the sake of simplicity and in order to focus on our main mechanism, the effect of composition changes on income-fertility relations.

[13] The remaining first-order conditions of the household in period one and the derivations of equations (11) and (12) can be found in Appendix C.1.

tively. Note that given a_2, the expression on the right-hand-side of (11) is the potential consumption level of the household without children, and the term on the right-hand-side of (12) is the total household expenditure on their children. Equations (11) and (12) show that (temporarily setting aside dynamic concerns), households will choose their number of children such that they devote fraction $\frac{1}{1+\zeta}$ of their non-savings to actual consumption and fraction $\frac{\zeta}{1+\zeta}$ to expenditures on their children.

Periods Two and Three

At the start of period two, labor productivities of the household ϵ_2^f and ϵ_2^m are realized according to the law of motions (9) and (10). During this period, the households can no longer have babies but continue to participate in the labor market. They solve,

$$V_2^{ij}(a_2, \epsilon_2^f, \epsilon_2^m) = \max_{c_2, a_3} \quad \ln(c_2) + \beta \ln\big((1+r)a_3\big) \tag{13}$$

such that

$$c_2 + a_3 = w_f^i \epsilon_2^f + w_m^j \epsilon_2^m + (1+r)a_2. \tag{14}$$

The term $\ln\big((1+r)a_3\big)$ represents the household's utility in period three, as the household retires at the start of period three and dies with certainty at the end of period three, implying that they will consume their entire savings $(1+r)a_3$. Thus in period two the household consumes

$$c_2 = \frac{w_f^i \epsilon_2^f + w_m^j \epsilon_2^m + (1+r)a_2}{1+\beta},$$

and saves

$$a_3 = \beta c_2.$$

3.2 The Representative Firm

Following previous studies such as (Katz and Murphy, 1992; Heckman et al., 1998; Heathcote et al., 2010; Hong et al., 2019), we assume that the representative firm aggregates the labor supply of various workers via the following constant-elasticity of substitution function:

$$L \equiv \left[\lambda_h\left(\lambda_f L_f^h + (1-\lambda_f)L_m^h\right)^{\frac{\psi-1}{\psi}} + (1-\lambda_h)\left(\lambda_f L_f^c + (1-\lambda_f)L_m^c\right)^{\frac{\psi-1}{\psi}}\right]^{\frac{\psi}{\psi-1}}, \tag{15}$$

where L_f^i and L_m^j indicate the total labor supply of either female or male workers with education level $i, j \in \{h, c\}$, and λ_f and λ_h indicate the weights on female workers and workers with a high school diploma respectively. L_f^i and L_m^j are computed as simply the sum of the labor of all workers of that type. The parameter ψ governs the elasticity of substitution between skilled (college) and unskilled (high school) workers.

The representative firm produces final goods using a standard Cobb-Douglas production function

$$Y = AK^{\alpha}L^{1-\alpha}. \tag{16}$$

We assume that the markets for capital and labor are competitive, resulting in the standard demand for capital:

$$r + \delta = \alpha\left(\frac{L}{K}\right)^{1-\alpha}, \tag{17}$$

and the demand for labor:

$$w_f^h = \lambda_f \lambda_h (1-\alpha) A \left(\frac{K}{L}\right)^{\alpha} \left(\frac{L}{\lambda_f L_f^h + (1-\lambda_f) L_m^h}\right)^{\frac{1}{\psi}}, \tag{18}$$

$$w_m^h = (1-\lambda_f) \lambda_h (1-\alpha) A \left(\frac{K}{L}\right)^{\alpha} \left(\frac{L}{\lambda_f L_f^h + (1-\lambda_f) L_m^h}\right)^{\frac{1}{\psi}}, \tag{19}$$

$$w_f^c = \lambda_f (1-\lambda_h) (1-\alpha) A \left(\frac{K}{L}\right)^{\alpha} \left(\frac{L}{\lambda_f L_f^c + (1-\lambda_f) L_m^c}\right)^{\frac{1}{\psi}}, \tag{20}$$

$$w_m^c = (1-\lambda_f)(1-\lambda_h)(1-\alpha) A \left(\frac{K}{L}\right)^{\alpha} \left(\frac{L}{\lambda_f L_f^c + (1-\lambda_f) L_m^c}\right)^{\frac{1}{\psi}}. \tag{21}$$

3.3 Child Care and Quality Sectors

There are many child care firms in a competitive market that take inputs in the form of final goods and produces child care time using technology

$$t_b = \varphi^{-1}(i_t),$$

that turns i_t units of the final good into time t_b. Likewise child quality firms produce child quality q using technology,

$$q = \vartheta^{-1}(i_q),$$

that turns i_q units of the final good into quality q. Since the firms are competitive, the households must pay $\varphi(t_b)$ units of the final good to purchase t_b units of child care time and $\vartheta(q)$ units of the final good to purchase q units of child quality. This is as shown in the household budget constraint (7). For our quantitative analysis we use the functional forms:

$$\varphi^{-1}(i_t) = \frac{i_t}{p_b}, \qquad \vartheta^{-1}(i_q) = \left(\frac{i_q}{p_q}\right)^{\theta}.$$

3.4 Equilibrium

In equilibrium, households behave optimally choosing their education level, consumption each period, time spent on raising their children, spending on market services for childcare, and investment in the quality of their children. The representative firm chooses capital and labor demand for each worker type. Prices and wages are determined such that the market for goods, child care, child quality, capital, and each labor type clears. In Appendix C.2, we describe the equilibrium of this model in more detail, including a formal definition of the equilibrium and a presentation of the full set of equilibrium conditions.

Table 1: Externally Calibrated Parameters

Parameters	**Value**	**Description**	**Source**
α	0.36	Capital share	Aiyagari (1994)
δ	0.06	Capital depreciation rate	Prescott (1986)
Z	1.00	Child Production Technology	Normalized to one
ψ	1.43	The elasticiy of substitution between skilled and unskilled	Katz and Murphy (1992)
$(\pi_{hh},\ \pi_{hc})$	$(0.79, 0.21)$	Matching probability of high school graduated women	The U.S. Census
$(\pi_{cc},\ \pi_{ch})$	$(0.71, 0.29)$	Matching probability of college graduated women	The U.S. Census
$(\rho^h,\ \rho^c)$	$(0.986, 0.984)$	Persistency of AR(1) labor productivity process for high school and college graduates	Hong et al. (2019)
$(\sigma^2_{\nu,h},\ \sigma^2_{\nu,c})$	$(0.01, 0.02)$	Variance of labor productivity process for high school and college graduates	

Note: Parameter values denoted in annual units.

Table 2: Internally Calibrated Parameters

Parameters	**Benchmark Economy**	**Comparison Economy**	**Description**
(A) Household Utility and Education			
β	0.99	-	Time discount factor
$(\mu_\zeta, \sigma^2_\zeta)$	$(1.81, 0.86)$	-	Mean and variance of fertility preference ζ
$(\mu_{\log\kappa}, \sigma^2_{\log\kappa})$	$(3.50, 0.15)$	-	Mean and variance of (log) education cost κ
$\sigma_{\zeta,\log\kappa}$	-0.02	-	Covariance of fertility preference and education cost
(B) Child Care, Labor Supply and Demand			
p_b	0.24	-	Price of child care services
p_q	0.49	-	Price of child education services
θ	0.04	-	Child quality production technology
ϕ	0.42	-	Weight on the mother's time in the child production function
$\frac{1}{1-\gamma}$	1.08	-	Elasticity of substitution between the mother's time and market services
(λ_f, λ_h)	$(0.40, 0.53)$	(0.52, 0.44)	Weight on women and high school graduates worker

Table 3: Targeted Moments

Target	Data	Benchmark Economy	Comparison Economy
(A) Household Utility and Education			
Capital to output ratio (K/Y)	3.00	3.01	*2.99*
Average CFR	1.89	1.92	*1.80*
Slope of CFR-income relationship	-0.20	-0.25	*-0.17*
The CFR of high school to college graduates ratio	0.99	0.97	*1.03*
Female college graduates share	0.27	0.27	*0.33*
Slope of female income-college graduate share relationship	0.05	0.17	*0.22*
(B) Child Care, Labor Supply and Demand			
Average child care cost to household income ratio	0.20	0.20	*0.20*
Average educational investment to household income ratio	0.02	0.02	*0.02*
Slope of education investment-income relationship	0.36	0.23	*0.25*
Average female labor supply (normalized by male labor supply)	0.52	0.53	*0.58*
Slope of female labor supply-income relationship	0.20	0.10	*0.09*
Skill premium (1973-1990)	1.50	1.49	-
Skill premium (1991-2010)	1.72	-	1.72
Gender wage gap (1973-1990)	1.51	1.50	-
Gender wage gap (1991-2010)	1.25	-	1.25

Notes: Three moments of the Comparison Economy have been targeted in the calibration: the female college graduates share in 2010, the skill premium, and the gender wage gap. Moments that have not been targeted directly are italicized.

3.5 Calibration and Results

We calibrate our benchmark model to match key features of the U.S. economy using 1990 as our base year, and our "comparison economy," which we use to analyze the impact of rising income inequality, to reflect the changes in the income structure leading up to the year 2010.[14] First, we present the parameters calibrated using values derived from external sources. We set the capital share α to the standard value of 0.36 as in Aiyagari (1994) and the capital depreciation rate δ to 0.06 annually as in Prescott (1986). The parameter governing the level of the production function of children Z is normalized to one and the parameter governing the elasticity of substitution between skilled and unskilled labor is adopted from Katz and Murphy (1992) and set at 1.43. The matching probabilities for marriage and the population shares of each agent type is computed using the U.S. Census data for the year 1990. The probability that a woman with a high school degree marries either a high school graduate or a college graduate, respectively, is $(0.79, 0.21)$. The probability that a women with a college degree marries a college school graduate and a high school graduate is $(0.71, 0.29)$. The proportion of females with high school and college degrees relative to the population is computed as $(0.36, 0.14)$ and the same proportion for males is $(0.33, 0.17)$. To calibrate the productivity process, we use the estimates from Hong et al. (2019) that estimate the persistence of productivity to be $(0.986, 0.984)$ for high school and college graduates, respectively, and the volatility to be $(0.01, 0.02)$, respectively. Table 1 summarizes these results.

We use simulated method of moments to calibrate the remaining parameters.[15] We calibrate 13 parameters to match 13 moments of the benchmark economy. We note that while our discussion below introduces each parameter as if they are being calibrated to match a

[14] Appendix A.2 discusses how the target moments, shown in Table 3, and data relationships depicted in Figures 6 and 7 were constructed.

[15] We describe our calibration method in more detail in Appendix D.

particular target moment, in practice they are jointly determined to simultaneously match all of the targeted moments.

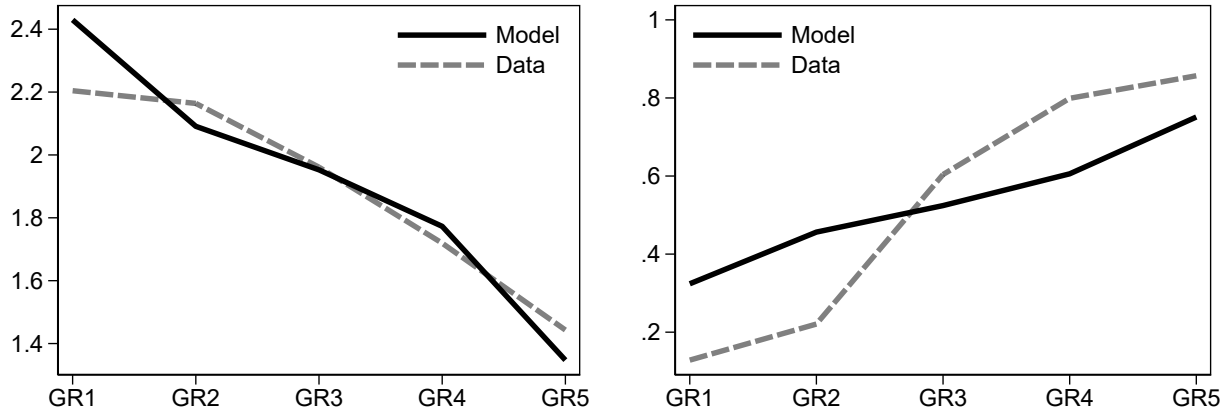
There are 6 parameters related to household utility and education choice. The time discount factor β is calibrated to 0.99 in order to match the capital output ratio $K/Y = 3.0$. The mean and variance of the distribution governing fertility preference $(\mu_\zeta, \sigma_\zeta^2)$ targets the average CFR level and the slope of the relationship between CFR and female labor income computed using the U.S. Census data, and is set to $(1.81, 0.86)$. The parameters governing the distribution of education costs $(\mu_{\log\kappa}, \sigma_{\log\kappa}^2)$ are set to $(3.50, 0.15)$ and the covariance between fertility preferences and educational costs $(\sigma_{\zeta,\log\kappa})$ is set to -0.02. These parameters target the college graduation rate in 1990, the CFR ratio between high school and college graduates in 1990, and the correlation between female labor income and college graduate share computed from the U.S. Census data.

The remaining parameters govern child care, labor supply, and labor demand. Note that the two parameters governing child care and child quality technology p_b and p_q determine their price. Thus we calibrate these parameters to match the child-rearing cost per child and the educational investment per child as a fraction of household labor income, which is 20% and 2%, respectively, in the data. The parameter θ in the quality production function is set to 0.04 in order to match the slope of the relationship between education expenditures and household income. The parameter ϕ in the production function for children is set to 0.42 to target the average of female labor supply as a ratio of male labor supply, which is equal to 0.52 in the data. The elasticity of substitution between the mother's time and market services, $\frac{1}{1-\rho}$, is set to 1.08 and targets the slope of the relationship between female labor supply and income of 0.20. We calibrate the weights (λ_f, λ_h) of the labor aggregation function (15) in the benchmark economy to target the average values of the gender wage gap and the skill premium computed from hourly wages between 1973 and 1990, using data from

the Economic Policy Institute's State of Working America Data Library. Table 2 summarizes these results.

We note that the elasticity of substitution between the mother's time and market services for child care, estimated to be 1.08 in our model, is lower than the values used by Bar et al. (2018) of 2.4 and by Park (2018) of 3.3 in quantitative models that are similar to ours in spirit. It is difficult to find estimates for the elasticity between a mother's time and market services in general child care in the empirical literature. However, the empirical literature investigating the relationship between the mother's time and market services as inputs in child skill accumulation finds that the elasticity of substitution is very low, in the range of 0.2 to 0.5 (Molnar, 2023; Caucutt et al., 2021). The literature that investigates the elasticity between time and market services for household chores (such as cleaning) find that the elasticity is around 2 (Aguiar and Hurst, 2016; Guryan et al., 2008). Child care in our model, represents the full range of activities associated with raising a child, from basic household chores to educational activities for young children (such as improving the child's vocabulary by talking and reading to them). Therefore, we believe that our estimated value for the elasticity, that falls in between the two extremes, is reasonable.

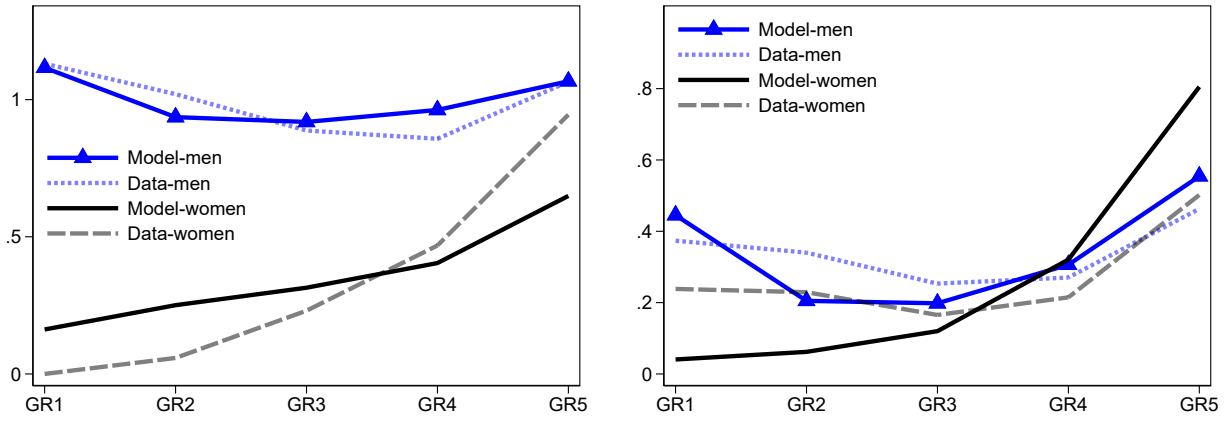
Table 3 shows the target moments, comparing values in the data to those of the model. The benchmark model generally matches the target moments closely. In Panel (a) of Figure 6, we plot the CFRs by female labor income. While we target the average relationship (i.e., the slope) between income and CFR, we find that the model also does a good job of matching the values of each income quintile separately, as the relationship is close to linear in the data. In Panel (b), we plot female labor supply by labor income. While our model does a reasonable job of matching the average relationship, the model slightly overestimates labor supply for the low-income individuals and underestimates labor supply from the high-income individuals.



(a) CFR by female labor income (b) Female labor supply by female labor income

Figure 6: CFR and Female Labor Supply

Panel (a) shows the CFRs by female labor income. Panel (b) shows female labor supply, normalized by the average labor supply of males.



(a) Normalized labor income (b) Fraction of college graduates

Figure 7: Income and Education

The results in panel (a) are normalized using the average labor income of males. Panel (b) shows the proportion of male and female college graduates in each income group.

In panel (a) of Figure 7 we plot the average female and male labor incomes of households, by groups sorted by female income. The dashed lines show the relationship in the data and the solid lines show the relationships in the model. The results show that, while not targeted directly, the model still does a good job of capturing the income patterns of of female and male workers in the household. In particular, the male labor income of households sorted by female labor income exhibits a U-shape in the data and this is pattern captured by our model.

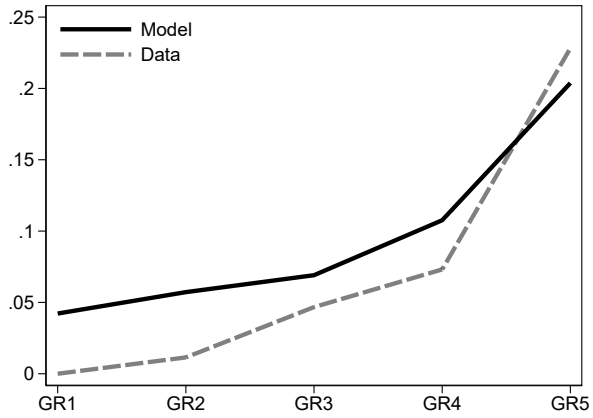
Panel (b) plots the fraction of college-educated workers by gender and income group. Our model does a reasonable job of matching the observed relationship between education and labor income in the data. The targeted education-income relationship of female workers increases a bit more sharply with income in the model compared to the data, but is otherwise comparable. We do not directly target the relationship between female labor income and education levels for male workers, but the the model does a good job of capturing the U-shaped relationship found in the data.

4 The Effects of Rising Income Inequality

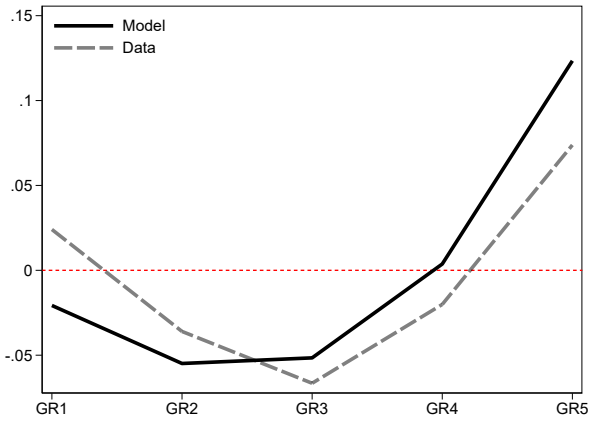
Using the model presented in Section 3, we investigate how much the rise in income inquality experienced in the United States over the past few decades has contributed to the changes in the fertility preferences and fertility rates across income groups that we presented in Section 2.

4.1 The Comparison Economy

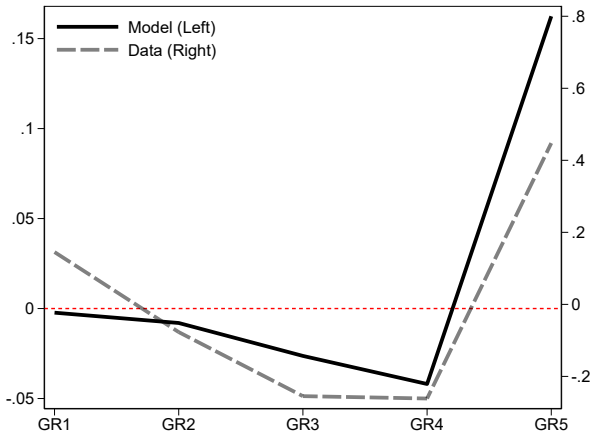
To simulate the change in the income structure, we re-calibrate the parameters (λ_f, λ_h) for the comparison economy so that the wage structure shows an increase in the skill premium and a decrease in the gender wage gap commensurate with the changes found in the data.



(a) Changes in female labor income



(b) Changes in the average CFR



(c) Changes in average fertility preference (ζ)

Figure 8: Effects of Rising Income Inequality

For the comparison economy, we target the average value of the skill premium and the gender wage gap between 1991 and 2010 which is equal to 1.72 and 1.25, respectively, in the data. The column labeled "Comparison Economy" in Table 2 shows the new parameter values.

To study the impact of the changes in the income structure on fertility, we analyze how the income-fertility relationship has changed from the benchmark economy to the comparison economy. In panel (a) of Figure 8 we compare the *changes* in female labor income observed in the data versus the model. By adjusting the parameters (λ_f, λ_h), our model is able to replicate the disproportionate increase in labor income across income groups, which is large for high income groups and small for low-income groups. The increase in wage inequality mainly raises the wages of highly productive women and thus increases the income differences between groups. Furthermore, higher wages allow households to outsource a greater proportion of child care to market services, allowing highly skilled females to work more hours. Panel (b) illustrates the changes to the CFR for each income quintile. The model is successful in generating the pattern of change, where the fertility rates of high income groups increase, the middle income groups decrease, and the low income groups see relatively little change. Panel (c) shows the changes to the income-fertility preference relationship in the model by plotting the average value of ζ for each income group. The general patterns showing increased preference for fertility in the high-income groups and decreased average fertility preferences in the middle-income groups resemble the patterns found in the data as depicted by the dashed line.[16]

As previously discussed, these changes are the result of the fact that with higher wages, women with high skill levels can now earn higher income even while allocating some time to raising children, and also because they can work more hours by outsourcing some of their

[16]The model results are in units as shown on the left vertical axis of panel (c), and the data results are in units as shown on the right vertical axis.

child care duties to market services as the relative price of market services fall. We discuss the relative contributions of each mechanism in the following section.

In addition, we find that women who have a relatively higher preference for children become more likely to attend college when returns to education increases because it not only allows them to increase consumption, but also the number of their children. This can be seen when we compare the ratio of the average value of ζ of high school over college graduates, in the benchmark and the comparison economies. The value of the ratio $(\bar{\zeta}_c/\bar{\zeta}_h)$, where $\bar{\zeta}_c$ is the average ζ among college graduates and $\bar{\zeta}_h$ among high school graduates, is equal to 0.81 in the benchmark economy whereas it increases to 0.82 in the comparison economy. This leads to a similar increase in the relative CFR of college graduates, where the ratio of the average CFR of college graduates to high school graduates is 0.97 in the benchmark economy and 1.03 in the comparison economy. This is especially true due to assortative matching, because high-income partners allow women to have more children and not worry as much about working.

Our model highlights the fact that the effect of increased income inequality on fertility occurs primarily through changes in the composition of high and middle income groups, where the relative income of highly skilled households with high fertility preferences increases and that of less skilled households with low fertility preferences falls in relative terms. How this change affects the educational investment of households is relevant for its influence the accumulation of human capital and consequently economic growth. Figure 9 compares average education investment of the benchmark model to the comparison model, sorted by female labor income quintiles. We find that education investment per child increases for the higher income households (groups 4 and 5), and the rise becomes steeper as income group increases. Meanwhile, it remains largely unchanged for households in lower income groups (groups 1 to 3).

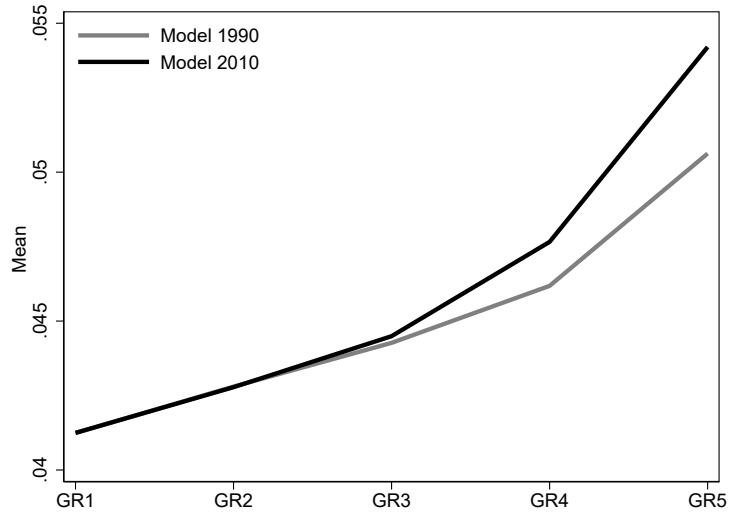


Figure 9: Education Investment per Child

4.2 Composition Changes and Market Services

We now decompose the effect of the two major mechanisms in our model, changes in the composition of households constituting each income group and the fall in the relative price of market services, to understand the relative importance of each channel in explaining the observed changes to the income-fertility relationship as a response to increased income inequality. To do so, we use our model to first simulate the changes to the income-fertility relationship when neither channel is present, then simulate the results when the relative price of market services fall, and then finally compare with our full model where we also incorporate the composition changes.

First, to see the effect of rising income inequality in an economy without both the composition changes and improved access to market services, we set the price of market service time p_b so that the ratio of this price to the woman's hourly wage remains unchanged for each household. In other words, we set the price p_b separately for both high school and college educated women so that the ratio p_b/w_f^i takes on the value from the benchmark economy. Then, we fix the composition of households in each income group using their income levels in

the benchmark economy to compute fertility levels for each group after changes to the wage structure. This allows us to disregard the changes in the income-fertility relationship that result from the changes in the composition of households with different preferences, skills, and education costs. The blue dotted line in Figure 10 shows this result. Without either channel, we find that the fertility rates of households in all income groups fall by similar proportions when income inequality rises. This leads to a significant under-prediction of fertility rates for the lowest income households in group 1 as well as for the high income households in groups 4 and 5.

The solid green line in Figure 10 shows the results when we allow the relative price of market services to fall, as in the comparison economy. When the relative price of market services fall, the fertility rates of all households increase by similar magnitudes across income groups. While this may lead to better predictions for the change in the average fertility rates across all households, the model still under-predicts the increase in the fertility rates of households with high female labor income and also under-predicts the degree to which the relative fertility rates of households in the middle income groups fall.

Finally, when we include both channels, as shown by the solid black line, our model is able to explain the differential changes to the fertility patterns of households by income. The model can account for both the increase in the fertility of high income households and the decrease in the fertility of households in income quintiles 2 and 3. Households in income groups 1 and 4 do not see a large change in their fertility rates, both in the model and the data.

4.3 Changes to the Skill Premium and the Gender Wage Gap

So far, we have considered rising income inequality as a whole, taking into account the effects of increasing skill premia and the decreasing gender wage gap together. However,

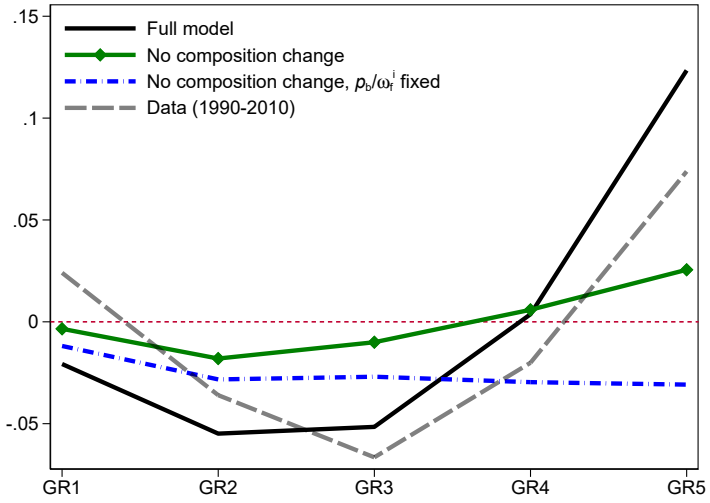


Figure 10: A Decomposition of Effects by Mechanism

when considering potential implications for policy, whether rising skill premia or reductions in the gender wage gap is more central to the changes may become a meaningful distinction. Thus in this section, we investigates the contributions of each phenomenon separately. To do so, we change the value of either λ_f to 0.44 or λ_h to 0.52, which are the values for the comparison economy, and keep the other parameter values as was in the benchmark economy. We then study how the income-fertility relationship changes in the model economy.

Figure 11 shows the decomposition of the effects of rising income inequality into the contributions from the increase in the skill premia and the contribution of the decrease in the gender wage gap. The solid line re-illustrates the aggregate effect from the full model and each bar shows the decompositions. As can be seen, we find that the increase in the skill premium explains a large portion of the variation in the change in the income-fertility relationship. The rise in the skill premium by itself explains 61% of the total variation generated by the full model.[17] It works to increase the fertility of both the highest and lowest income households while lowering the income of households in quintiles 2 and 3. The

[17]We describe how we measure how much of the variation is explained by each component in Appendix E.

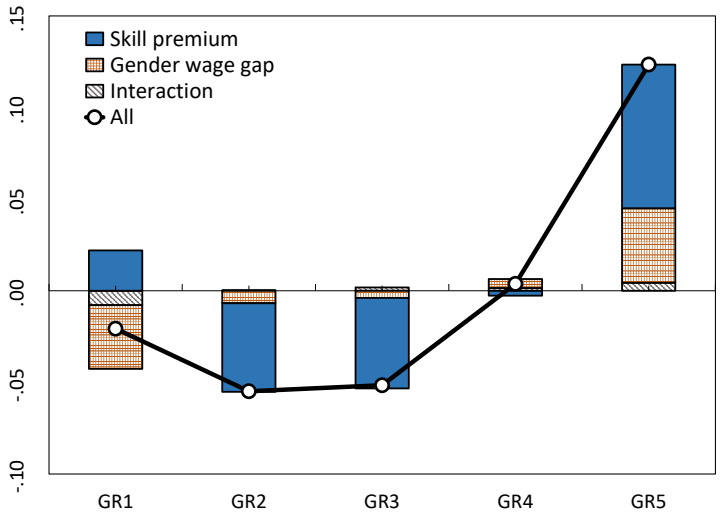


Figure 11: Decomposing the Effect of Changes to Income Structure

reason the skill premium raises the fertility of the households with the lowest levels of female income is because this group includes many college-educated women that do not work but are married to college-educated men whose income increases with the skill premium. Thus the fertility of these households increase.

The gender wage gap acts mostly on the lowest and the highest income quintiles. This is due to the fact that a decrease in the gender wage gap causes high income women to have more children as the wage of high-skilled, child loving women increase (and they are able to outsource child care to a greater degree). The reduction in the gender wage gap lowers the relative fertility rates of low income households as these households are unable to capture the benefits of increasing female wages, as the women in this group do not work very many hours.

In total, the effect of an increasing skill premium and a decreasing gender wage gap cancel each other out in the lowest income quintile while they reinforce each other in the highest income quintile, resulting in the shape of the income-fertility relationship as shown by the solid line.

5 Conclusion

In this paper we document three empirical facts found in the data: that income inequality has risen, that the relative fertility of high-income households has risen, and that the fertility preference of high-income households has risen in the United States over the past few decades. We find that these three phenomena can be jointly explained when we consider how rising income inequality differently affects households that are heterogeneous in terms of their fertility preferences and labor productivity. Our quantitative model is able to account for nearly all of the changes observed in the income-fertility relationship.

This study focuses on the explanation for the change in the observed relationship between female labor income and fertility. A natural next step would be to explore the consequences of this change in the income-fertility relationship, on the accumulation of human capital and growth. Exploring how rising income inequality affects human capital accumulation and growth through the changes that it imposes on household fertility patterns may lead to a new and better understanding of how to address the various issues that arise with greater income inequality.

References

Aguiar, M. and E. Hurst (2016). The Macroeconomics of Time Allocation. In *Handbook of Macroeconomics*, Volume 2, pp. 203–253. Elsevier.

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics 109*(3), 659–684.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in U.S. Wage Inequality: Revising the Revisionists. *Review of Economics and Statistics 90*(2), 300–323.

Bar, M., M. Hazan, O. Leukhina, D. Weiss, and H. Zoabi (2018). Why Did Rich Families Increase Their Fertility? Inequality and Marketization of Child Care. *Journal of Economic Growth 23*(4), 427–463.

Becker, G. (1960). An Economic Analysis of Fertility. In *Demographic and Economic Change in Developed Countries*, pp. 209–240. National Bureau of Economic Research, Inc.

Blake, J. (1966). Ideal Family Size Among White Americans: A Quarter of a Century's Evidence. *Demography 3*(1), 154–173.

Blomquist, N. S. (1983). The Effect of Income Taxation on the Labor Supply of Married Men in Sweden. *Journal of Public Economics 22*(2), 169–197.

Borjas, G. J. (1987). Self-Selection and the Earnings of Immigrants. *American Economic Review 77*(4), 531–553.

Briley, D. A., F. C. Tropf, and M. C. Mills (2017). What Explains The Heritability of Completed Fertility? Evidence from Two Large Twin Studies. *Behavior Genetics 47*(1), 36–51.

Burtless, G. and J. A. Hausman (1978). The Effect of Taxation on Labor Supply: Evaluating

the Gary Negative Income Tax Experiment. *Journal of Political Economy 86*(6), 1103–1130.

Caucutt, E., L. Lochner, J. Mullins, and Y. Park (2021). Child Skill Production: Accounting for Parental and Market-Based Time and Goods Investments. *Working Paper*.

De la Croix, D. and M. Doepke (2003). Inequality and Growth: Why Differential Fertility Matters. *American Economic Review 93*(4), 1091–1113.

Doepke, M., A. Hannusch, F. Kindermann, and M. Tertilt (2023). The Economics of Fertility: A Mew Era. In *Handbook of the Economics of the Family*, Volume 1, pp. 151–254. Elsevier.

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

Gobbi, P. E. (2013). A Model of Voluntary Childlessness. *Journal of Population Economics 26*(3), 963–982.

Greenwood, J., A. Seshadri, and G. Vandenbroucke (2005). The Baby Boom and Baby Bust. *American Economic Review 95*(1), 183–207.

Guryan, J., E. Hurst, and M. Kearney (2008). Parental Education and Parental Time with Children. *Journal of Economic Perspectives 22*(3), 23–46.

Hakim, C. (2000). *Work-lifestyle Choices in the 21st Century: Preference Theory*. Oxford University Press.

Hakim, C. (2003). A New Approach to Explaining Fertility Patterns: Preference Theory. *Population and Development Review 29*(3), 349–374.

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

Hazan, M. and H. Zoabi (2015b). Sons or Daughters? Sex Preferences and the Reversal of the Gender Educational Gap. *Journal of Demographic Economics 81*(2), 179–201.

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

Heckman, B. J. J., L. Lochner, and C. Taber (1998). Tax Policy and Human-Capital Formation. *The American Economic Review 88*(2), 293–297.

Hong, J. H., B. H. Seok, and H. M. You (2019). Wage Volatility and Changing Patterns of Labor Supply. *International Economic Review 60*(2), 1091–1113.

Jones, L. E., A. Schoonbroodt, and M. Tertilt (2008). Fertility Theories: Can They Explain the Negative Fertility-income Relationship? In J. B. Shoven (Ed.), *Demography and the Economy*, pp. 43–100. University of Chicago Press.

Jones, L. E. and M. Tertilt (2008). An Economic History of Fertility in the United States: 1826–1960. In P. Rupert (Ed.), *Frontiers of Family Economics*, Volume 1, pp. 165–230. Emerald Publishing.

Katz, L. F. and K. M. Murphy (1992). Changes in Relative Wages, 1963-87: Supply and Demand Factors. *Quarterly Journal of Economics 107*, 35–78.

Kornrich, S. and F. Furstenberg (2013). Investing in children: Changes in parental spending on children, 1972–2007. *Demography 50*(1), 1–23.

Lino, M. (1996). Expenditures on Children by Families, 1995. *Family Economics and Nutrition Review 9*(3), 1–20.

Lino, M. (1998). Expenditures on Children by Families, 1997. *Family Economics and Nutrition Review 11*(3), 1–25.

Miller, W., J. Rodgers, and D. Pasta (2010). Fertility Motivations of Youth Predict Later Fertility Outcomes: A Prospective Analysis of National Longitudinal Survey of Youth Data. *Biodemography and Social Biology 56*, 1–23.

Moav, O. (2005). Cheap Children and the Persistence of Poverty. *The Economic Journal 115*(500), 88–110.

Molnar, T. L. (2023). Costs of Daycare, Complementarities, and Heterogeneous Productivity of Parenting Time in Child Skill Formation. *Working Paper*.

Park, S. (2018). A Structural Explanation of Recent Changes in Life-Cycle Labor Supply and Fertility Behavior of Married Women in the United States. *European Economic Review 102*, 129–168.

Philipov, D. and L. Bernardi (2011). Concepts and Operationalisation of Reproductive Decisions Implementation in Austria, Germany and Switzerland. *Comparative Population Studies 36*(2-3).

Prescott, E. C. (1986). Theory Ahead of Business-cycle Measurement. *Carnegie-Rochester Conference Series on Public Policy 25*, 11–44.

Roy, A. D. (1951). Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers 3*(2), 135–146.

Sauré, P. and H. Zoabi (2014). International Trade, the Gender Wage Gap and Female Labor Force Participation. *Journal of Development Economics 111*, 17–33.

Shang, Q. and B. A. Weinberg (2013). Opting for Families: Recent Trends in the Fertility of Highly Educated Women. *Journal of Population Economics 26*(1), 5–32.

Siegel, C. (2017). Female Relative Wages, Household Specialization and Fertility. *Review of Economic Dynamics 24*, 152–174.

Sobotka, T. and É. Beaujouan (2014). Two Is Best? The Persistence of a Two-Child Family Ideal in Europe. *Population and Development Review 40*(3), 391–419.

Sommer, K. (2016). Fertility Choice in a Life Cycle Model with Idiosyncratic Uninsurable Earnings Risk. *Journal of Monetary Economics 83*, 27–38.

U.S. Census Bureau (2020). *Understanding and Using American Community Survey Data: What All Data Users Need to Know*. Government Publishing Office.

Wildsmith, E., J. Manlove, and E. Cook (2018). Dramatic Increase in the Proportion of Births Outside of Marriage in the United States from 1990 to 2016. *Child Trends*.

Appendix for Online Publication

A Data and Variable Construction

In this section, we discuss the data and the construction of variables used throughout our analysis.

A.1 Empirical Facts

We begin by discussing how we construct the variables used in the empirical facts presented in Section 2. First, for Figure 1, we compute the skill premium and gender wage gap using "Wages by education" and "Median/average hourly wage" series from the State of Working America Data Library, provided by the Economic Policy Institute. These series are constructed using the Current Population Survey Outgoing Rotation Group microdata.

We use the U.S. Census data and American Community Surveys from IPUMS USA to compute our income and fertility measures used in Figure 2. In particular, we use the U.S. Census 5% samples for the years 1980, 1990, 2000, and the ACS 5-year estimates data for 2018.[18] The CFR of each household is simply their number of children (NCHILD) and we use the Wage and Salary Income variable (INCWAGE) for the labor income of each household member. We then sort households by the labor income of the women of the household to construct groups by female labor income quantile. We then compute the average CFR of the households in each income group. The normalized CFR is computed by dividing the CFR of each income group by the average CFR across all households in a given year.

[18]The U.S. Decennial Census, conducted since 1790, used a short form survey for basic demographic information of the population, and a long form survey with detailed socioeconomic and housing questions for a sample of households until 2000. After the 2000 Census, the ACS replaced the long form survey. The ACS collects detailed socioeconomic and housing data every year. In addition, ACS 5-year estimates provide data on the average characteristics of the U.S. population over a 5-year period, based on information collected during a 60 month time frame. See U.S. Census Bureau (2020) for a detailed discussion.

We construct the change in fertility preferences shown in Panel (b) of Figure 3 using the NLSY79 data from the National Longitudinal Surveys. The NLSY79 has three race categories: Hispanic, Black, and non-Hispanic and non-black. We include only households where the women is non-Hispanic and non-black and was continuously married to the same person between the ages of 33 and 40. We find that our results are not sensitive to this age range (not reported). As discussed in Section 2, the NLSY79 asks, "How many children do you want to have altogether?" (FER-1A) and "What do you think is the ideal number of children for a family?" (FER-1B). We use the answers to the second question to construct our preference measure. The above questions are only asked twice during the survey period, once in 1979 and then again in 1982. Therefore, in order to keep the age of the respondent as similar as possible, for the cohort born between 1958-59, we use the responses given in 1979 when their ages were between 20-21, and for the 1964-65 cohort we use the responses given in 1982 when they were of the age 17-18. We compute labor income using the variable "total income from wages and salaries in the previous calendar year" (Q13-5).

A.2 Target Moments

We now discuss how we compute our target moments in the data. We compute the capital-to-output ratio (K/Y) using "Fixed Assets and Consumer Durable Goods" and "Gross Domestic Product" from the Bureau of Economic Analysis. The capital-to-output ratio fluctuates between 3.0 and 3.5 during the 1990s. We choose the standard value of 3.0 as the target moment.

We use the Census and ACS data discussed in the previous section to compute the rest of the target moments. The average CFR and the slope of the CFR-income relationship is constructed as was described in the previous section, in the construction of Figure 2. Female labor income is normalized by dividing income by average male income for that given year.

We compute the CFR of high school to college graduates ratio and the female college graduate share in 1990 using the educational attainment variables EDUC and EDUCD. We classify someone as a college graduate if they have a Bachelor's degree. Labor supply is measured by annual hours worked. We multiply the variables WKSWORK2 and UHRSWORK to construct annual hours worked, where WKSWORK2 represents the number of weeks worked and UHRSWORK indicates the number of hours worked per week. All labor supply variables are normalized by dividing by the average hours worked for males for that given year. These variables are also used in the construction of the data relationships in Figures 6 and 7.

We compute the average child care cost to household income ratio and average educational investment to household income ratio from the average estimated annual expenditures on a child by husband-wife families as well as average household income given in Lino (1996, 1998). The slope of the education investment-income relationship is taken from of Kornrich and Furstenberg (2013). The skill premium and gender wage gap are computed as discussed in Appendix A.1.

B Robustness

B.1 Changes in Fertility: Robustness

Figure B.1 shows the income-fertility relationship using the (not normalized) CFRs as the fertility measure. Panel (a) shows the average CFRs by female income quintile and panel (b) shows the change in the CFRs in relation to the value in 1990. As can be seen, the depicted income–fertility relationship patterns are very similar to those computed using normalized CFRs in Figure 2.

While we focus our analysis on women aged between 36 and 40, as the fertility rates of women over 40 is very low (Briley et al., 2017), our results remain robust even if we instead

analyze women between the ages of 46 and 50. This result is illustrated in Figure B.2. Panels (a) and (b) show the results using raw CFRs and panels (c) and (d) show results using normalized CFRs. Both panels (a) and (c) show that the income–fertility relationship is also negatively correlated and panels (b) and (d) show that the changes to the income–fertility patterns are also similar to those found in women between 36 and 40, where CFR of the middle income group falls in relative terms and those of high and low income groups rise.

Figure B.3 shows the income–fertility relationship when we include women of all races instead of focusing only on white women. Again, we find that these figures are very similar to the patterns shown in Figures 2 and B.1. In contrast, the patterns revealed in Figure B.4, which shows the income–fertility relationship of non-white women, diverges from the previously observed patterns. In particular, the fertility rates of low income women decreases considerably, unlike in the previous analyses. Nevertheless, the relative fertility rate of high-income females increase, as shown in panel (d), and this pattern is consistent with the main channel proposed in this paper. We note that Doepke et al. (2023) also finds similar patterns in the fertility rates by education levels for both black and non-white women.

B.2 Changes in Fertility Preference

Figure B.5 illustrates the changes in the income–fertility preference relationship using the answers to the general fertility preference question "What do you think is the ideal number of children for a family?" in the NLSY79 survey. The figure shows the change in fertility preference responses between the cohort born in 1958-59 and the cohort born in 1964-65, with responses collected when each cohort was between the ages 17-21. The income group to which each respondent belongs is determined by their labor income between the ages of 37-48. The change in the general fertility preference shown here is similar to that of personal

fertility preference shown in Figure 3.

C Further Discussions of the Model

C.1 Household Optimization in Period One.

As introduced in Section 3.3, we use the functional forms

$$\varphi^{-1}(i_t) = \frac{i_t}{p_b}, \qquad \vartheta^{-1}(i_q) = \left(\frac{i_q}{p_q}\right)^{\theta},$$

for the child care and child quality firms' technologies. We can then rewrite the household's period one problem as follows:

$$V_1^{ij}(\epsilon_1^f, \epsilon_1^m; \zeta) = \max_{c,a_1,n,q} \ln(c_1) + \zeta \ln(nq) + \beta \mathbb{E} V_2^{ij}(a_2, \epsilon_2^f, \epsilon_2^m; \zeta) \tag{22}$$

such that

$$c_1 + (p_b t_b + w_f^i \epsilon_1^f t_f) + p_q q^{\frac{1}{\theta}} n + a_2 = w_f^i \epsilon_1^f + w_m^j \epsilon_1^m. \tag{23}$$

We define $p^*(n) \equiv p_b t_b^*(n) + w_f^i \epsilon_1^f t_f^*(n)$ as the minimized child-rearing cost function given n, where $(t_b^*(n), t_f^*(n))$ denotes the optimal choice of (t_b, t_f) given n. The households minimize cost as follows:

$$p^*(n) = \min_{t_f, t_b} \quad p_b t_b + w_f^i \epsilon_1^f t_f \tag{24}$$

such that

$$n = Z \left[\phi t_f^{\gamma} + (1-\phi)\, t_b^{\gamma}\right]^{1/\gamma} \tag{25}$$

If a household has zero or negative ζ, $t_f = t_b = n = 0$. Otherwise, the optimal choices for t_f

and t_b are:

$$t_f^*(n) = \min\left(\frac{\phi^{\frac{1}{1-\gamma}}}{Z\left[\phi^{\frac{1}{1-\gamma}} + (1-\phi)^{\frac{1}{1-\gamma}}\left(\frac{w_f^i \epsilon_1^f}{p_b}\right)^{\frac{\gamma}{1-\gamma}}\right]^{\frac{1}{\gamma}}} \cdot n, 1\right) \tag{26}$$

$$t_b^*(n) = \min\left(\frac{(1-\phi)^{\frac{1}{1-\gamma}}}{Z\left[\left(1-\phi\right)^{\frac{1}{1-\gamma}} + \phi^{\frac{1}{1-\gamma}}\left(\frac{p_b}{w_f^i \epsilon_f^i}\right)^{\frac{\gamma}{1-\gamma}}\right]^{\frac{1}{\gamma}}} \cdot n, 1\right). \tag{27}$$

The relative wage, $(\frac{w_f^i \epsilon_f^i}{p_b})$, determines the division of time between raising the child personally and investing in market services. If the relative wage increases, more time is allocated toward market services and less of the mother's time is used to raise the child.

Substituting (26) and (27) back into the expression for $p^*(n)$ we get

$$p^*(n) = \hat{p} \cdot n \tag{28}$$

where

$$\hat{p} = \frac{1}{Z}\left[\phi^{\frac{1}{1-\gamma}}\left(w_f^i \epsilon_f^i\right)^{\frac{\gamma}{\gamma-1}} + \left(1-\phi\right)^{\frac{1}{1-\gamma}}\left(p_b\right)^{\frac{\gamma}{\gamma-1}}\right]^{\frac{\gamma-1}{\gamma}}.$$

Note that $\hat{p}$ represents the cost of a child to the household.

Substituting (28) into (23) we can solve the utility maximization problem to find child quality,

$$q^* = \left(\frac{\theta}{1-\theta} \cdot \frac{\hat{p}}{p_q}\right)^{\theta} \tag{29}$$

Thus, we see that educational investment has a positive relationship with the marginal cost of a child, $\hat{p}$, and a negative relationship with the price of education, p_q. Hence, as the marginal cost of raising a child increases, households will invest more into the education of each child, indicating that the quantity-quality trade-off exists in our model. When the price

of education becomes more expensive households will decrease investment.

Then, we can solve for the optimal number of children given a_2 as,

$$n^* = \frac{\zeta}{1+\zeta} \cdot \frac{w_f^i \epsilon_f^i + w_m^j \epsilon_m^j - a_2}{\hat{p} + p_q (q^*)^{\frac{1}{\theta}}}. \tag{30}$$

The number of children a household has increases is their preference ζ, the household's earnings potential if they had no children $w_f^i \epsilon_f^i + w_m^j \epsilon_m^j$. On the other hand the amount of required savings a_2, the marginal cost of raising a child $\hat{p}$, and the cost of educational investment per child, $p_q (q^*)^{\frac{1}{\theta}}$ all lower fertility.

Finally, note that from (30) we have

$$\hat{p} n^* + p_q (q^*)^{\frac{1}{\theta}} n^* = \frac{\zeta}{1+\zeta} \cdot (w_f^i \epsilon_f^i + w_m^j \epsilon_m^j - a_2),$$

which we can then substitute into the budget constraint (23) which gives us the expressions (11) and (12).

C.2 Equilibrium

Let μ_k be the externally given number of households in the kth period of their lives. For convenience we refer to these households as of being age k. Let $F_k^{ij}(\epsilon_k^f, \epsilon_k^m, a_k, \zeta)$ be the joint cumulative distribution of households of age k with education levels (i, j) and labor productivities $(\epsilon_k^f, \epsilon_k^m)$, assets a_k, and preference ζ. Then the capital market clearing condition can be given as:

$$K = \sum_{k=1}^{2} \sum_{ij} \mu_k \int a_{k+1}^{ij}(\epsilon_k^f, \epsilon_k^m, a_k, \zeta) dF_k^{ij}(\epsilon_k^f, \epsilon_k^m, a_k, \zeta) \tag{31}$$

where $a_{k+1}^{ij}(\epsilon_k^f, \epsilon_k^m, a_k, \zeta)$ is the optimal savings of the household in period k. Note that households of age $k = 0, 1$ hold no assets.

We define Λ_c as the fraction of households that choose to attend college given (ζ, κ) in period 0. Then the fraction of high school educated women are $\Lambda_h \equiv 1 - \Lambda_c$ and the fraction of high school and college graduates among men are $\Lambda_h \pi_{hh} + \Lambda_c \pi_{ch}$ and $\Lambda_h \pi_{hc} + \Lambda_c \pi_{cc}$ respectively. Then, the labor market clearing conditions are:

$$L_f^i = \mu_1 \Lambda_i \sum_{j \in \{h,c\}} \pi_{ij} \int [(1 - t_f^{ij}(\epsilon_1^f, \epsilon_1^m, \zeta))\epsilon_1^f] dF_1^{ij}(\epsilon_1^f, \epsilon_1^m, 0, \zeta) \tag{32}$$
$$+ \mu_2 \Lambda_i \sum_{j \in \{h,c\}} \pi_{ij} \int \epsilon_2^f dF_2^{ij}(\epsilon_2^f, \epsilon_2^m, a_2, \zeta), \qquad \text{for } i \in \{h, c\}$$

$$L_m^j = \sum_{k=1}^{2} \sum_{i \in \{h,c\}} \mu_k \Lambda_i \pi_{ij} \int \epsilon_1^m dF_k^{ij}(\epsilon_k^f, \epsilon_k^m, a_k, \zeta), \qquad \text{for } j \in \{h, c\} \tag{33}$$

where $t_f^{ij}(\epsilon_1^f, \epsilon_1^m, \zeta)$ denotes the optimal choice of t_f given $(\epsilon_1^f, \epsilon_1^m, \zeta)$ and education levels (i, j). The goods market clears automatically by Walras law. Given the market clearing conditions we define the equilibrium of the model as follows.

Equilibrium Definition

For each $k \in \{0, 1, 2, 3\}$ and education level $i, j \in \{h, c\}$, the recursive stationary equilibrium is: the set of value functions V_0 and V_k^{ij}; education choice in period 0, policy functions c_k^{ij}, a_{k+1}^{ij}, n^{ij}, q^{ij}, t_f^{ij}, t_b^{ij}; prices r, ω_f^i, ω_m^j; firm's capital and labor demand K, L_f^i, and L_m^j; and distributions F_k^{ij} such that:

1. *Given prices r, ω_f^i, ω_m^j, and distributions F_k^{ij}, the value functions and the policy functions solve the household's utility maximization problems* (5), (6) *and* (13),
2. *Given prices r, ω_f^i, and ω_m^j the representative firm's demand for capital and labor satisfies* (17) - (21), *and*
3. *Markets clear.*
   (a) *The market for capital clears according to* (31).

(b) The market for labor clears according to (32) *and* (33)

(c) The goods market clearing condition is

$$Y = C + I + I_t + I_q,$$

where

$$\begin{aligned}
C &= \sum_{k=1}^{3} \sum_{ij} \mu_k \int c_k^{ij}(\epsilon_k^f, \epsilon_k^m, a_k, \zeta) dF_k^{ij}(\epsilon_k^f, \epsilon_k^m, a_k, \zeta) \\
I &= \delta K, \\
I_t &= \mu_1 \sum_{i,j} \int p_b t_b^{ij}(\epsilon_k^f, \epsilon_k^m, 0, \zeta) F_1^{ij}(\epsilon_f, \epsilon_m, 0, \zeta) \\
I_q &= \mu_1 \sum_{i,j} \int p_q q^{ij}(\epsilon_k^f, \epsilon_k^m, 0, \zeta)^{\frac{1}{\theta}} F_1^{ij}(\epsilon_f, \epsilon_m, 0, \zeta).
\end{aligned}$$

4. Distributions are consistent with individual behavior:

$$\begin{aligned}
F_1^{ij}(\epsilon_1^f, \epsilon_1^m, 0, \zeta) &= \Phi\Big(0, \Big(\frac{\sigma_{v,i}}{1-\rho_i}\Big)^2\Big) \times \Phi\Big(0, \Big(\frac{\sigma_{v,j}}{1-\rho_j}\Big)^2\Big) \qquad \forall \zeta, \\
F_2^{ij}(\epsilon_2^f, \epsilon_2^m, a_2, \zeta) &= \int Prob(\epsilon_2^f, \epsilon_2^m, a_2 | \epsilon_1^f, \epsilon_1^m) dF_1^{ij}(\epsilon_f, \epsilon_m, 0, \zeta)
\end{aligned}$$

where $\Phi(0, x)$ *is the cumulative normal distribution with mean zero and variance* x*, and* $Prob(\epsilon_2^f, \epsilon_2^m, a_2 | \epsilon_1^f, \epsilon_1^m)$ *indicates the conditional probability of realizing the values* $(\epsilon_2^f, \epsilon_2^m, a_2)$ *given* $(\epsilon_1^f, \epsilon_1^m)$ *according to the law of motions* (9) *and* (10)*.*

D Calibration Strategy

In this section, we describe how we calibrate the parameters listed in Table 2 using simulated method of moments. As discussed in the main text, the 13 parameters are calibrated jointly to simultaneously match all of the targeted moments. Each target moment is listed in Table

3. To calibrate the parameters, we first impute the externally calibrated parameters and guess the remaining 13 parameters that govern the benchmark economy. Then we compute the 13 target moments of the benchmark economy using these parameters and compare them to the moments from the model with the data.

In this manner we estimate the set of 13 parameters jointly so that the difference between the model and data moments are minimized. To do so, we define the percentage difference in moments as:

$$\Delta m_i(\Theta) = \frac{m_i^{data} - m_i^{model}(\Theta)}{m_i^{data}}$$

where m_i^{data} is the ith moment computed from the data and m_i^{model} is the generated moment from the model simulation. Θ is the set of parameters to be estimated. Next, we construct the objective function of the simulated method of moments as follows:

$$Q(\Theta) = \sum_{i=1}^{13} \omega_i \left(\Delta m_i(\Theta)\right)^2 \tag{34}$$

where ω_i is the weight for each moment. We set the weights ω_i to one, for all moments. Finally, we estimate the set of parameters, $\hat{\Theta}$, which minimizes the objective function in Equation (34) as:

$$\hat{\Theta} = \arg\min_{\Theta} Q(\Theta).$$

E Measure of Explanatory Power

We define a measure of explanatory power used to evaluate how well the model explains the observed changes as follows. First, define Δ_i as the change in the CFR between 2010 and 1990 for a given income group i that we wish to explain as follows,

$$\Delta_i = n_{i,2010} - n_{i,1990},$$

where $n_{i,t}$ is the average CFR of income group i in the year t. Then, we define total change Δ, as the sum of the absolute values of Δ_i across all income groups as follows:

$$\Delta = \sum_i |\Delta_i| .$$

We take the absolute value of Δ_i to prevent the positive and negative changes for each income group from canceling out. Likewise, we define $\hat{\Delta}_i$ and $\hat{\Delta}$ as described above, for the changes observed in the model simulations.

Next, we define U_i, a measure of changes observed in the data that is unexplained by the model, as

$$U_i = \Delta_i - \hat{\Delta}_i.$$

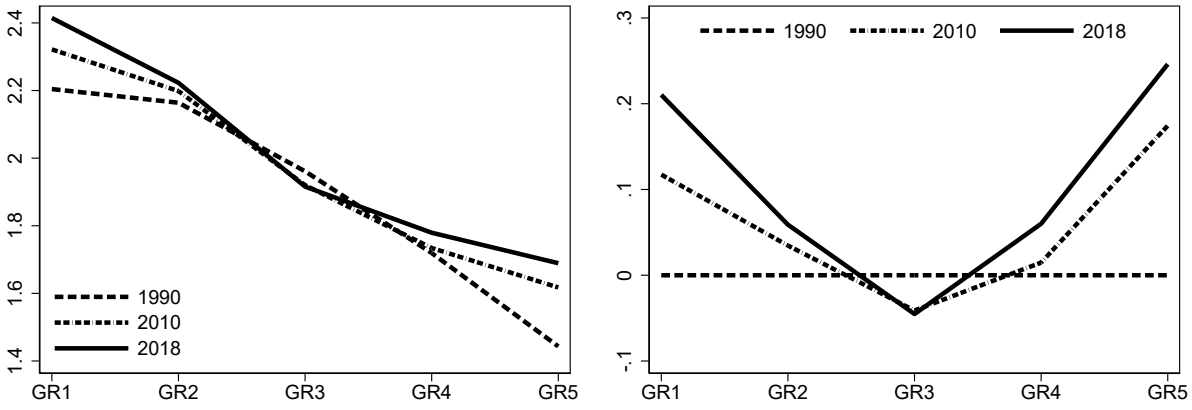
Again, we define the total unexplained variation as summation of U_i,

$$U = \sum_i |U_i| .$$

Finally, the proportion of the changes that is successfully explained by the model is measured as,

$$E = \frac{\Delta - U}{\Delta}.$$

We compute this measure when only the skill premium increases and when only the gender wage gap decreases separately for the decomposition of the effect of each channel. As described in Section 4.3, the change in the skill premium alone is able to explain 65% of the total variation generated by the full model.

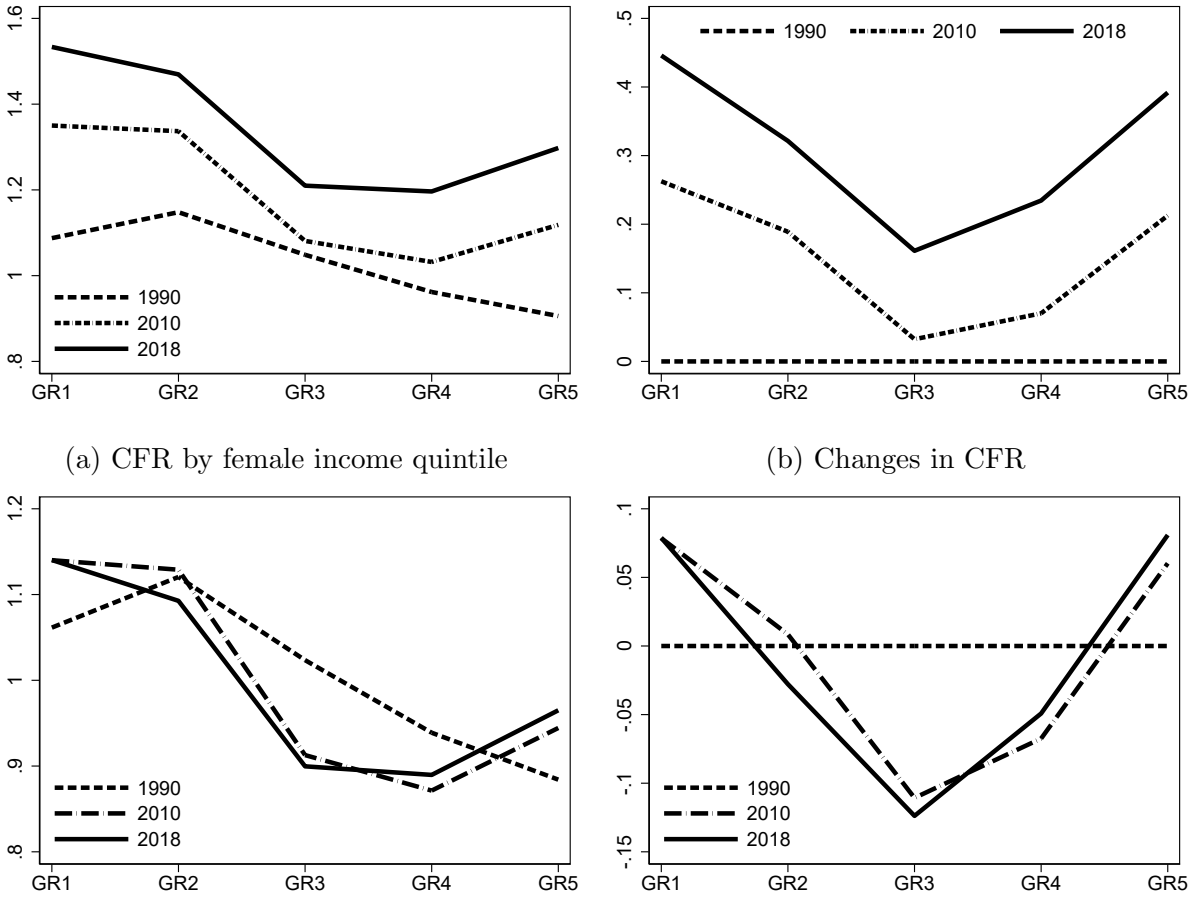


(a) CFR by female income quintile

(b) Changes in CFR

Figure B.1: Changes in Fertility (Not normalized)

Notes: Panel (a) shows the average CFR by income quintile for the years 1990, 2010, 2018. Panel (b) plots the change in the CFR using the year 1990 as the benchmark.



(a) CFR by female income quintile

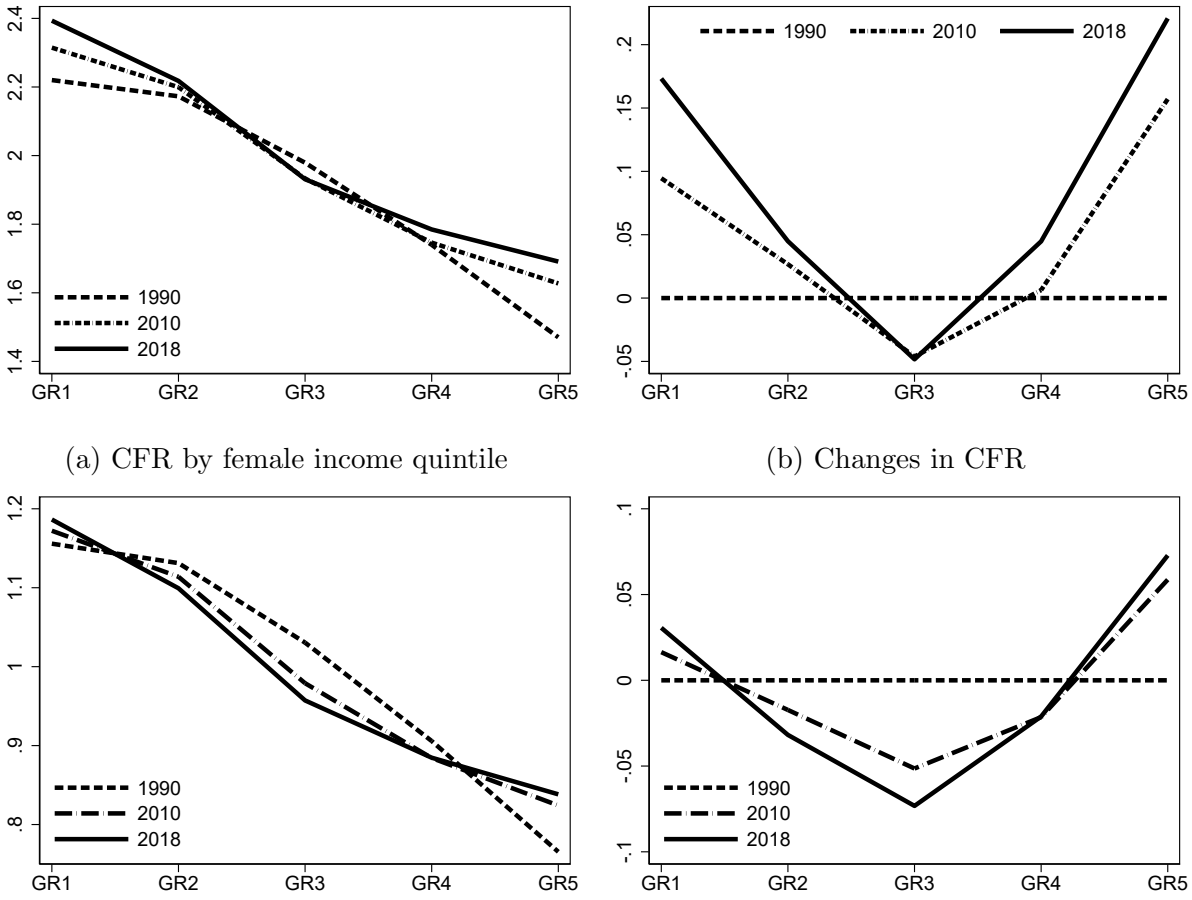
(b) Changes in CFR

(c) Normalized CFR by female income quintile

(d) Changes in normalized CFR

Figure B.2: Changes in Fertility: Age range 46-50

Notes: This figure shows the average CFR for each income quintile for married white women aged 46 to 50 whose husbands work at least 40 weeks a year and 35 hours a week on average for the data construction.



(a) CFR by female income quintile

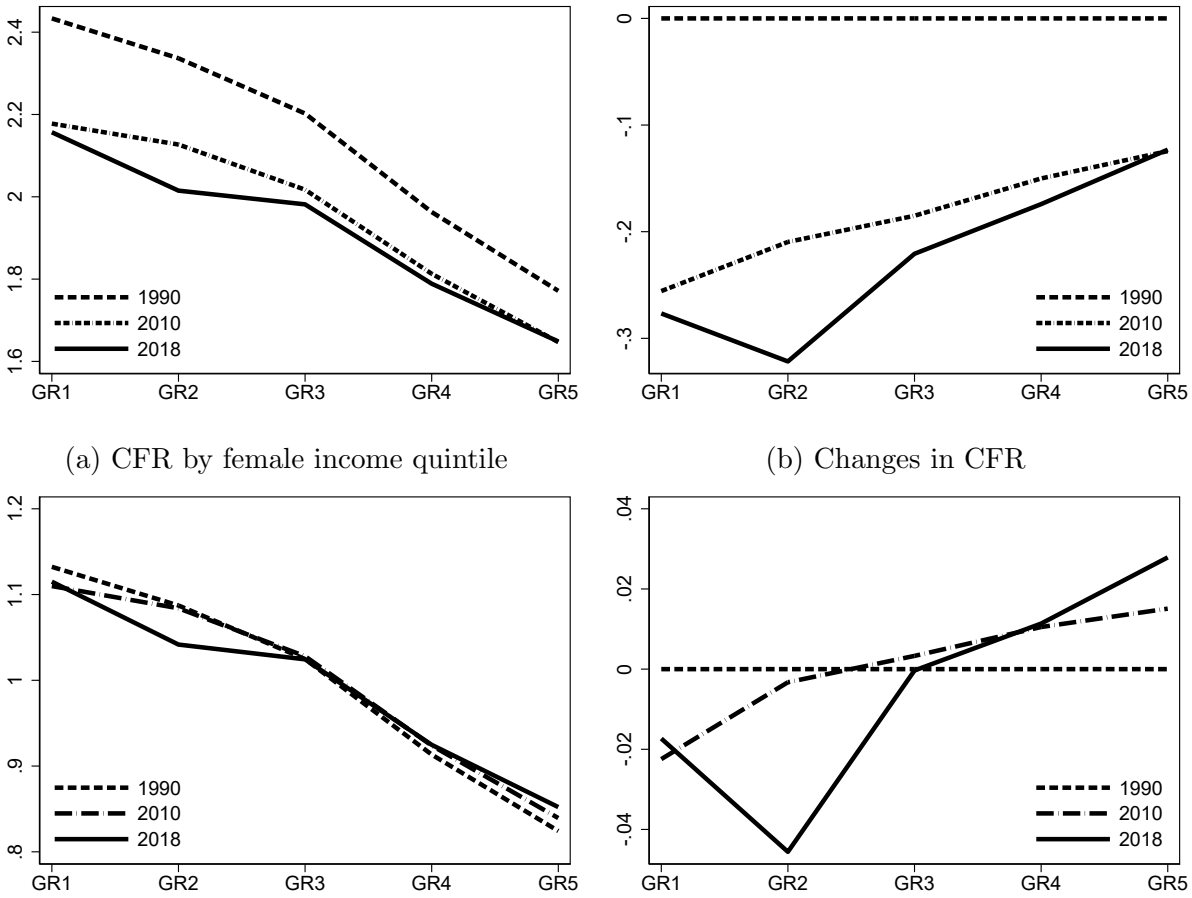
(b) Changes in CFR

(c) Normalized CFR by female income quintile

(d) Changes in normalized CFR

Figure B.3: Changes in Fertility: All races

Notes: This figure shows the average CFR for each income quintile for married women of all races aged 36 to 40 whose husbands work at least 40 weeks a year and 35 hours a week on average for the data construction.



(a) CFR by female income quintile

(b) Changes in CFR

(c) Normalized CFR by female income quintile

(d) Changes in normalized CFR

Figure B.4: Changes in Fertility: Non-white women

Notes: This figure shows the average CFR for each income quintile for non-white married women aged 36 to 40 whose husbands work at least 40 weeks a year and 35 hours a week on average for the data construction.

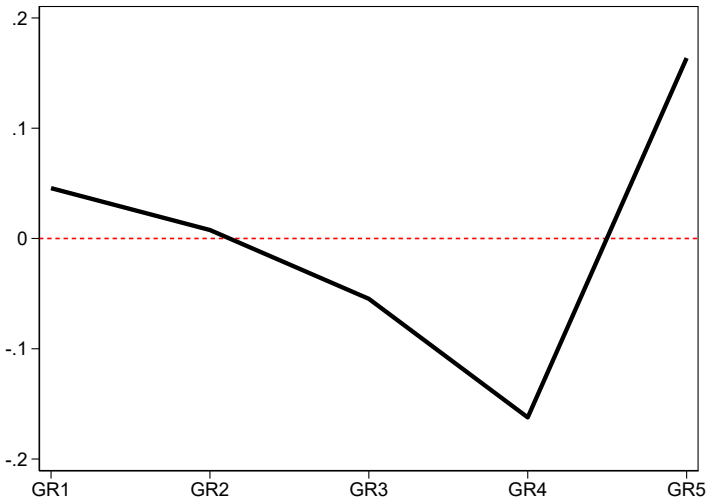


Figure B.5: Changes in General Fertility Preference

Notes: This figure illustrates the changes in the income-fertility preference relationship using the survey question asking about general fertility preference in the NLSY79. It shows the change in the responses between the cohort born in 1958-59 to the cohort born in 1964-65, made when each cohort was between the ages 17 and 21. The income group to which each respondent belongs is determined by their average labor income between the ages of 37 to 48.