

Female Relative Wages, Household Specialization and Fertility

Christian Siegel*

January 31, 2017

Abstract

Falling fertility rates have often been linked to rising female wages. However, over the last 40 years the US total fertility rate has been rather stable while female wages have continued to grow. Over the same period, women's hours spent on housework have declined, but men's have increased. I propose a model in which households are not perfectly specialized, but both men and women contribute to home production. As the gender wage gap narrows, the time allocations of men and women converge, and while fertility falls at first, the decline stops when female wages are close to male's. Rising relative wages increase women's labor supply and due to higher opportunity cost lower fertility at first, but they also lead to a reallocation of home production and child care from women to men, and a marketization. I find that both are important in understanding why fertility did not decline further. In a further quantitative exercise I show that the model performs well in matching fertility over the entire 20th century, including the overall decline, the baby boom, and the recent stabilization.

JEL codes: D13, E24, J13, J22

Keywords: Fertility, Female Labor Supply, Household Production, Intrahousehold Allocations

*University of Kent, School of Economics, Canterbury, Kent, CT2 7NP, UK; C.Siegel@kent.ac.uk. This is a substantially revised version of a chapter of my PhD thesis at the London School of Economics entitled 'Female Employment and Fertility: The Effects of Rising Female Wages'. I am grateful to Francesco Caselli for his advise on this project. I would also like to thank two anonymous referees, the editor Matthias Doepke, Zsófia Bárány, Wouter den Haan, Monique Ebell, Paula Gobbi, Moshe Hazan, Ethan Ilzetzki, Winfried Koeniger, John Knowles, Rigas Oikonomou, Albert Marcet, Michael McMahon, Rachel Ngai, Silvana Tenreyro, as well as seminar and conference participants at Bonn, Cologne, Exeter, LSE, Uppsala, the SED 2013 meeting in Seoul and the Demographic Economics Conference at the University of Iowa for many comments and suggestions.

1 Introduction

Between the 1960s and today, we have seen enormous changes to the economic and demographic structure in all Western countries. There has been a decline in total fertility rates¹ and an increase in women’s market hours (see figures 1 and 2 for US data).

Many authors explain both with a rise in female wages (e.g. Galor and Weil (1996), Doepke, Hazan, and Maoz (2015)). An apparent puzzle, however, is that while female wages and market hours have continued to grow, since the early 1970s fertility has stopped falling.² Understanding the underlying fertility decisions is important since they affect population growth, labor force composition and social security systems, and thereby economic outcomes. In this paper I argue that the common driving force behind the trends in fertility and in female employment is the narrowing of the gender wage gap (shown in figure 3), rather than the level of female wages per se, since it changes the division of labor within the family. The explanation I propose is based on imperfect specialization of households, such that both men and women contribute to home and market production, implying that a reallocation has a nonlinear effect.

My explanation is based on the observation that men’s home hours have increased, allowing women’s home hours to fall. In the 1960s, when the wage gap was large, the catching up of female wages increased women’s labor supply. The associated increase in the opportunity cost of women’s time, who shouldered most of child care, lowered fertility, as argued by Becker (1960) or Galor and Weil (1996). But as relative wages become more equal over time, specialization in the household decreases. Consequently, male home hours increase, a father’s time at home becomes more important for raising children, and the allocation of time between home and market work becomes more evenly balanced for men and women. I show that when the complementarity between mother’s and father’s time working at home is sufficiently large, the marginal utility cost of having an additional child can become constant, or even fall, despite women working more hours in the market.

Circumstantial evidence in favor of this mechanism is provided by data on non-market hours. Using data from Aguiar and Hurst (2007), based on the American Time Use Survey since 1965, I show the trends in hours spent working at home in figure 4. I focus on the sum of time spent non-market work, defined by Aguiar and Hurst as home production plus obtaining goods, and basic child care.³ The data displays a shift in household production

¹The total fertility rate (TFR) is the average number of children that would be born if all women lived to the end of their childbearing years and bore children according to the current age-specific birth rates.

²Most of the recent rise in the official total fertility rate is driven by the effects of immigration. For US-born women the increase is much less and fertility virtually flat since the late 1970s. The details on the decomposition of TFR by mothers’ birthplace are given in the online appendix, section A.1.

³The measure is constructed as weekly hours spent on home production in a narrow sense, on obtaining goods, and on basic child care. Arguably, for many people child care is an activity that is more enjoyable than other forms of housework. Excluding child care from the measure of market work does not change the qualitative trends. This can be seen in the regression results of table A-1 in the online appendix.




Figure 1: Male and female market hours

This graph shows average (per capita) yearly market hours worked for men and women aged 20 to 60. The data is taken from the PSID and refers to self-reported hours worked in the previous year. In the panel on the left, all individuals of age 20 to 60 are included. The panels to the right restrict the sample to married and to single individuals respectively. Per capita hours worked are computed as the simple average of the average hours worked in 5-year age bins, in order to take out a potential composition effect due to changing cohort sizes.

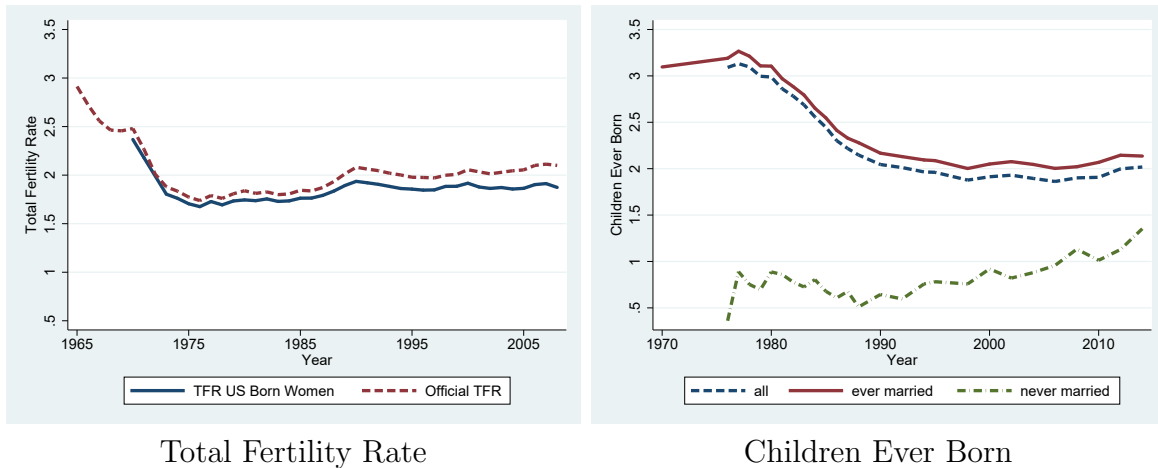


Figure 2: Total fertility rate and children ever born

In the left panel, the red dashed line shows the official total fertility rate (TFR) for the United States, taken from the World Bank (based on national statistics and the United Nations Demographic Yearbook), and the blue solid line the author's computation of TFR for US-born women, based on the Vital Statistics of the United States combined with population estimates from the US Census; see online appendix for details.

The right panel shows children ever born to women of age 40 to 44; the blue dashed line for all women of age 40 to 44, the red solid line and the green dot-dashed line the breakdown for ever married and never married women respectively. The data is taken from the United States Census Bureau's Fertility Data Historical Time Series and based on the CPS.




Figure 3: Female Relative Wages

Ratio of median female to male hourly wages, computed in a PSID subsample of men and women aged 20 to 60 who work at least one hour in paid employment and report positive earnings.

from women to men, with an overall reduction of hours worked at home. As a consequence, the ratio of married men’s to married women’s hours devoted to working at home rose from 0.25 to 0.54 over 1965–2003. There has also been a shift in time investments into children. Sayer, Bianchi, and Robinson (2004) find in time diary data that the ratio of married mothers’ to married fathers’ time in child care declined in all primary child care activities since mid-1960s. This is consistent with the facts documented by Robinson and Godbey (2008), who find in time use data over 1965–1995 evidence for convergence in activities across gender.⁴ The tendencies for men and women’s time spent on home production to converge have been at work throughout the last century, as documented by Ramey (2009).

The importance of considering intrahousehold allocations can be seen in the right panels of figure 4. While home hours of married men have increased over time, single men’s hours of home production are constant, after an initial change between 1965 and 1975. This is consistent with the explanation I propose. Single men, who are without a female partner, are not affected by rising female wages, whereas married men spend more hours working at home –a trend in US data that has not received much attention by researchers yet, with the notable exception of Knowles (2013).⁵ Single women, on the other hand, spend less time working at home, but the decline is not as pronounced as for

⁴They also report that average parental hours spent on child-care per child has been roughly constant. Ramey and Ramey (2010), on the other hand, include time spent teaching children and document a rise since the mid 1990s, especially for college educated parents.

⁵Burda, Hamermesh, and Weil (2007) study time use data across various developed countries, and find that total work, the sum of home and market work, is virtually the same for men and women. For the US, Ramey and Francis (2009) report that total work of men and women has been constant throughout the 20th century.




Figure 4: Male and female home hours worked

This graph shows average (per capita) weekly hours spent on nonmarket work and basic child care for men and women aged 20 to 60. The data is taken from Aguiar and Hurst (2007). Per capita hours worked are computed as the simple average of the average hours worked in 5-year age bins, in order to take out a potential composition effect due to changing cohort sizes.

married women, whose husbands devote more of their time to home production.⁶

Distinguishing household types is also important in the fertility data. Figure 2 shows that the total fertility rate (left panel) has been rather stable over the last forty years (and even increased slightly) and that the number children ever born (CEB) for women of age 40–44 (right panel) has been flat for the last 25 years.⁷ However, the trends in CEB for ever-married and never-married women differ substantially. While overall fertility has been flat, there has been an increase in the fertility of never married women. However, the rise in single women’s fertility is not the sole driver of the flattening out in aggregate data. To the contrary, the total figure of children ever born follows more closely the number for ever-married women, which has been stable for the last 25 years. Understanding the evolution of married fertility is the focus of this paper.

I present a model matching the observed patterns of married fertility and hours worked through an exogenous decrease in the gender wage gap. As the main interest of this paper

⁶In the Aguiar and Hurst (2007) data marital status is defined in a legal sense, and it is not possible to disentangle cohabiting from other singles. In figure A-4 of the online appendix I show that home hours have changed more for singles with children (who are more likely to be cohabiting) than for singles without children.

⁷CEB has flattened out only from 1990 on, since it shows life-time fertility of women aged 40-44, whereas TFR is a measure of fertility across women of all ages at a point in time. Since TFR is computed by adding up the age-specific fertility rates of all women (in their child-bearing years), it is potentially subject to a tempo effect, which is that an increase in the age of childbearing mechanically lowers the computed total fertility rate. CEB is not affected by changes in the timing of fertility, as it measures actually completed fertility. In general, TFR and CEB follow similar trends though.

is in understanding the recent time series of fertility, starting with the first available time use data in 1965, I abstract from many factors that might have been important for the fertility decline in earlier periods, such as a reduction in infant mortality or quality-quantity considerations in the wake of technological changes during the 19th century (Becker, Murphy, and Tamura (1990), Galor and Weil (2000)). However, in an extension I show that the model –under parameters calibrated to 1965–2005– performs well in replicating the whole 20th century data. Moreover, since the focus is on the time series, I do not model heterogeneity within a cohort. Nonetheless it is worth noting that Hazan and Zoabi (2015) find in recent U.S. cross-sectional data that fertility is not monotonically declining in women’s education, but U-shaped.

For each cohort I consider a representative household that maximizes the sum of a male and a female member’s utilities. Both members can work in the market and at home. Due to complementarities in home production the household is not perfectly specialized. When women’s wages are lower than men’s, they contribute more time to home production. A rise in female relative wages directly increases female labor supply and lowers female home production, whereas more male time is devoted to home activities and less to market work. Initially, when the gender wage gap is large, there is an overall drop in home labor and a couple devotes less time to having and raising children.

However, when the gender wage gap is fairly small and shrinks further, fertility stabilizes and might even increase. The reason for this differential reaction to improvement in female relative wages is in the degree to which households are specialized. Initially, a husband’s labor supply was much higher than his wife’s, and his time spent working at home much lower than hers. As a consequence, the opportunity cost of having children was largely determined by her wage (which was low compared to his), and when female wages improved the cost of having children increased substantially. In this situation the substitution effect of higher female wages dominates the positive income effect, leading to a fall in fertility. But as relative wages become more equal, the increase in the opportunity cost of having children becomes smaller, since the cost of reallocating home production from the wife to the husband falls. As a consequence, the substitution effect weakens with the narrowing of the gender wage gap. On top of this, with the improvement in the wife’s earnings, a couple can acquire more parental time-saving inputs, and the larger her labor supply already is, the more the household gains when her wage increases.

Both the rise in male home labor and the higher use of parental time-saving inputs into home production is what I find key in explaining why the fertility decline ended. First I use a simplified model to show analytically that the combined substitution effect becomes weaker than the income effect before the gender wage gap has fully closed if complementarities between both parents’ time spent in the household are sufficiently strong and if marketization is possible. Then I calibrate the full model to study the quantitative implications, which suggest further that the degree of marketization is important for the

timing of the fertility stabilization.

To the best of my knowledge, the only previous paper that has noted the flattening out in total fertility rates and informally suggested an explanation in terms of increased male home production is Feyrer, Sacerdote, and Stern (2008).⁸ My paper formally models and quantifies the endogenous response of male and female hours and their implications for fertility. In a working paper Regalia and Ríos-Rull (1999) discuss the stability in fertility and use a rich model to assess the effects of changes in wage premia, including the gender wage gap, for fertility decisions. My paper differs along two margins. On the one hand, focusing on married fertility, I assume that the number of children is a joint decision the couple takes. On the other hand, since both parents can contribute to home production, the implicit cost of having children is endogenous and a function of both parents' wages.

Other papers that have studied the implications of the decline in the gender wage gap for both male and female hours include Jones, Manuelli, and McGrattan (2003) and Knowles (2013), who highlights the rise of male home production and on whose work I build. In a calibrated life-cycle model Attanasio, Low, and Sanchez-Marcos (2008) find that a reduction in the gender wage gap alone cannot explain the change in labor supply of young mothers and that also a fall in the cost of children is needed. However, none of these papers have explored the implications for fertility. In my model, the rise of men's time spent on home production –in response to improved female wages– endogenously lowers mothers' cost of having children and the fertility decline ends.⁹ Galor and Weil (1996) present a unified framework to explain the rise in female employment and the fall in fertility that we observed until the early 1970s. Since their mechanism links fertility decisions to the market value of women's disposable time¹⁰, it cannot explain why fertility stopped falling when female wages continued improving. My work therefore highlights that the existing literature that assumes perfect specialization within families, such that women shoulder all of home production or child care, has overlooked important implications of intrahousehold allocations.

2 Cross-sectional Time Use Data

To support the view that the reallocation of home production from women to men is linked to relative wages, I use data from the American Time Use Survey (ATUS) for the

⁸Feyrer, Sacerdote, and Stern (2008) explain the increase in male home hours with intrahousehold bargaining. I show that bargaining is not necessary; a change in relative wages per se implies not only a reallocation of work in the market, but also at home.

⁹Attanasio, Low, and Sanchez-Marcos (2008) report that data on child care costs are not consistently available before 1980. For that reason, they calibrate the drop in the price of market child care using their model. Since they assume that only mothers provide domestic child care, they are likely to overstate the required reduction in child care costs as they do not take into account responses in fathers' time.

¹⁰Other papers linking fertility decisions to the market value of women's disposable time include Greenwood, Seshadri, and Vandenbroucke (2005) and Doepke, Hazan, and Maoz (2015), besides others.

Table 1: Cross-sectional regressions of time spent on home production

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Nbr. of Kids in Hh	59.11*** (7.01)	60.04*** (8.00)	59.39*** (7.34)	61.69*** (8.15)	62.02*** (7.22)	61.91*** (8.35)	60.31*** (7.98)	62.46*** (8.06)
(Nbr. of Kids in Hh) ²	-8.81*** (2.02)	-8.55*** (2.35)	-9.09*** (2.11)	-8.97*** (2.39)	-8.93*** (2.03)	-8.57*** (2.40)	-8.67*** (2.34)	-9.12*** (2.36)
Own market hours	-2.98*** (0.28)	-2.69*** (0.36)			-2.91*** (0.28)		-2.71*** (0.36)	
Spouse's market hours	1.88*** (0.30)	2.06*** (0.39)			1.77*** (0.30)		2.02*** (0.39)	
Own wage		-0.39 (0.28)		-0.49* (0.29)		-0.46 (0.29)		
Spouse's wage		0.66*** (0.22)		0.65*** (0.23)		0.54** (0.23)		
Own earnings			-0.03*** (0.01)					-0.03*** (0.01)
Spouse's earnings			0.04*** (0.01)					0.04*** (0.01)
Own age					-1.82** (0.72)	-2.62*** (0.81)		
Spouse's age					2.55*** (0.72)	2.94*** (0.82)		
Constant	204.76*** (17.51)	174.74*** (23.96)	145.56*** (8.41)	148.55*** (8.96)	174.17*** (23.07)	137.35*** (18.96)	183.74*** (22.53)	143.65*** (9.16)
Race Dummies	No	No	No	No	Yes	Yes	No	No
Sample	All	Empl	All	Empl	All	Empl	Empl	Empl
Observations	2691	2108	2417	2108	2691	2108	2108	2108
R^2	0.09	0.09	0.07	0.06	0.10	0.07	0.09	0.08

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

year 2011, provided by Hofferth, Flood, and Sobek (2013). The ATUS provides data on the time use of only the respondents themselves but not of other family members. It is therefore not possible to directly study how home production chores are split within a household. However as the ATUS respondents are sampled from individuals who completed the Current Population Survey (CPS), it is possible to link the respondents' time use data to the previous CPS data which include labor market characteristics of both the respondent and their partner. The lag between the ATUS and CPS interviews spans between 2 and 5 month, with an average of around 3 month.

I restrict the sample to married or cohabiting couples in which both partners are of age 21 to 65, and drop households with family businesses. To compute hourly wages for both partners, I divide weekly earnings by usual hours worked. As in figure 4, I define time spent working at home as the sum of home production, purchasing goods, and caring for household members.¹¹ Then I regress individuals' home work on their own and their partner's labor market characteristics, which come from the earlier CPS sample, and a set of other controls including number of children in the household. Denoting individual

¹¹In the ATUS, the respondents answer how many minutes they have spent on the various activities over a 24-hour period.

i 's time spent on household production by h_i , I estimate

$$h_i = \beta_0 + \beta_1 b_i + \beta_2 b_i^2 + \gamma_o X_i^o + \gamma_s X_i^s + error_i$$

where b_i are the number of children living in the household, and X_i^o individual i 's own and X_i^s their spouses' labor market characteristics, such as market hours worked, hourly wages, and weekly earnings (all taken from the previous CPS). Table 1 shows the results for different specifications.

The results of column 1 imply that individuals who work longer hours in the market tend to spend less time on home production. However, *ceteris paribus*, they spend more time on home production if their partner works more. The relationship is robust to include wages or age as further covariates, as seen in columns 2 and 5 respectively.

Moreover, columns 3, 4 and 6 show that individuals tend to devote less time to household production if they have higher hourly wages or weekly earnings, but more if their spouse earns more. Using wages in the analysis restricts the sample to households in which both the respondent and their spouse are working and wages are observed. Due to a potential self-selection into labor market participation, the result for the interdependence between home production time and spouses' wages are potentially not representative for the overall population. However, the results for the relationship with market hours worked (column 1) and with earnings (column 3) include households in which a spouse is not in employment. Excluding observations with zero market hours or earnings to restrict the sample to couples in which both are employed, gives the results in columns 7 and 8. For the correlation between home hours and market hours, or earnings, whether to include individuals out of employment, or not, does not matter much for the results.

To summarize, table 1 highlights that married individuals' time allocated to home production depends not only on their own labor market characteristics, but also on their spouses'.¹² These results are in line with the conjecture that higher female wages, or higher female market hours, are associated with a reduction in women's time in home production and an increase in their husband's home hours.

To the extent that better educated women earn higher wages, this is consistent with Hazan and Zoabi (2015)'s finding that married men's time spent providing childcare at home increases with mothers' education. However Hazan and Zoabi do not explore the implications of fathers' time for the fertility decision. In the next section I propose a model in which both parents' time matters, and provide a link between the rise in male home production and the flattening out of fertility over time.

¹²Splitting the regression sample by the gender of the respondent does not change the general pattern; see online appendix A.4.

3 The Model

3.1 Assumptions

To explain the time series of fertility as well as of male and female time allocations for married individuals, I propose a parsimonious model. For tractability I assume that each cohort can be represented by a representative married couple household that decides on how to allocate male and female time and on a continuous number of children to have. In the main text I focus on a version in which the couple lives for one period only, such that their choices are static. In the online appendix I present a life-cycle version of this model and show quantitatively that the general conclusions regarding fertility and time allocations remain the same. The advantage of the simple model presented in the main text is that it allows for an analytical characterization of the mechanism that links fertility to women’s relative wages.

3.1.1 Agents and Households

All agents, men and women, derive utility from the consumption of a market good (c), from having children (b), and from leisure, which is the time spent neither working in the labor market (n) or at home (h). Each individual is endowed with one unit of time. Agents differ in terms of their gender ($g \in \{m, f\}$) and their age (j). I assume that all economic active men and women live in couple households, formed by one man (‘husband’) and one woman (‘wife’).

I apply a model of collective household behavior, as introduced by Chiappori (1988). The male and the female partner have their own preferences, and derive felicity u_m and u_f respectively. Since they have children together, they solve a joint maximization problem. In particular, the couple household solves a Pareto program with relative weight θ attached to the husband’s and $1 - \theta$ to the wife’s utility. As will be seen later, under the model assumption made the spouses’ relative time spent on home production is independent of the bargaining weight, but solely based on comparative advantage. As a consequence, θ has no effect on the implicit cost of having children. I therefore restrict attention to the unitary model with $\theta = 0.5$ throughout.

3.1.2 Fertility

Child-care requires more home production (x), which could be done by the father, the mother, or both.¹³ A further input to home production are goods acquired in the market (e). Although I refer to this home-labor saving input as home appliances, in a broader sense this could include paid domestic help, such as hiring nannies. I assume that the

¹³The former assumption is similar to Erosa, Fuster, and Restuccia (2010), the latter is as in Knowles (2007) and similar to Greenwood and Seshadri (2005).

amount of the home good needed is

$$(1) \quad \bar{x}(b) = \kappa_0 + \kappa_1 b^{\kappa_2}$$

with $\kappa_0, \kappa_1, \kappa_2 > 0$, which implies $\bar{x}'(b) > 0$.

3.1.3 Household Production

Following Olivetti (2006) and in particular Knowles (2013), I assume that the home good is produced using a technology that is consistent with substitution among household member's time and home appliances according to

$$(2) \quad x = e^\gamma H^{1-\gamma},$$

with $0 \leq \gamma < 1$, where the home labor input is

$$(3) \quad H = (z_m h_m^{1-\rho} + z_f h_f^{1-\rho})^{\frac{1}{1-\rho}},$$

and z_m and z_f are the male and female home labor productivities. The parameter $\rho > 0$ is the inverse of the elasticity of substitution between male and female inputs. For finite ρ , this definition of H as a CES aggregator allows for substitutability between male and female inputs to home production – as long as both $z_m > 0$ and $z_f > 0$.

With $\gamma > 0$, this technology allows, to some degree, for a marketization of inputs to home production.¹⁴ Moreover, this home production function is consistent with the findings by Ramey (2009). From regressing the home capital-labor ratio on appliances prices, Ramey finds that the production function is more or less consistent with a Cobb-Douglas technology over home capital and labor.

3.1.4 Preferences

Agents derive utility from consuming the market good, from having children, and from leisure time. In particular, assume that preferences are additively separable and given by¹⁵

$$(4) \quad u(c_g, n_g, h_g, b) = \log(c_g) + \phi_l \frac{(1 - n_g - h_g)^{1-\eta} - 1}{1 - \eta} + \phi_b \frac{b^{1-\sigma_b} - 1}{1 - \sigma_b}$$

where c_g , n_g , and h_g are specific to a spouse $g \in \{m, f\}$, but the number of children b is common to the couple.

¹⁴In a structural transformation framework, Ngai and Pissarides (2008) study substitutions between home and market production, but do not distinguish between male and female labor.

¹⁵Adding utility from children in this additive form is generalizing Galor and Weil (1996) and Greenwood and Seshadri (2005), who assume $\ln(b)$.

Note that in this model setup individuals derive utility from leisure, i.e. the time spent not working in the market or at home, and both spouses contribute to home production due to complementarities.¹⁶

I assume that there is a subsistence level in consumption of the home produced good, which is increasing in the number of children living in the household (b). As a simplifying assumption, following Knowles (2007), agents do not derive any further utility from home production, and this constraint will be binding, $x = \bar{x}(b)$.

3.2 Couple Household's Optimization

The representative couple household maximizes the sum of male and female utilities, by choosing male and female hours of home production and of market labor supply and how many children to have, subject to a budget constraint, the home production requirement (1) and the home production technology (2). The household can earn a labor income $w_m n_m + w_f n_f$, where w_m is the male wage and w_f the female wage, which can be used to purchase consumption goods, c_m and c_f , and time-saving inputs to home production, e .

The representative couple's maximization problem is therefore given by

$$(5) \quad \max_{\substack{c_m, c_f, b, e \\ n_m, n_f, h_m, h_f}} \theta u_m(c_m, n_m, h_m, b) + (1 - \theta) u_f(c_f, n_f, h_f, b)$$

subject to the constraint set:

$$(6) \quad c_m + c_f + e = w_m n_m + w_f n_f$$

$$(7) \quad \bar{x}(b) = e^\gamma (z_m h_m^{1-\rho} + z_f h_f^{1-\rho})^{\frac{1-\gamma}{1-\rho}}$$

$$(8) \quad n_m + h_m \leq 1 \text{ and } n_f + h_f \leq 1$$

$$(9) \quad n_g \geq 0 \text{ and } h_g \geq 0 \text{ for } g \in \{m, f\}$$

where (6) is the household's budget constraint and (7) the relationship between the required amount of the home good and the home production inputs. The inequalities in (8) constrain the male and female total time working in the market and at home to be no larger than the time endowment of 1. As the marginal utility of leisure goes to infinity when leisure goes to zero, these constraint will be slack for both spouses. The inequalities in (9) are non-negativity conditions on the spouses' time spent working in the market and at home. As the home production function (3) features complementarity between male and female home hours, optimality requires h_m and h_f to be positive and the associated constraints never bind. The non-negativity constraints on market hours could in principle

¹⁶In an earlier working paper version, Siegel (2012), I had proposed an alternative to generate imperfect specialization within the household based on preferences featuring imperfect substitutability in the disutilities from working at home or in the market. Both models generate qualitatively the same predictions for the effects of changing relative wages on the couple's time allocation and on the fertility choice.

bind, but since I focus on a representative couple, I will choose parameter values for which both n_m and n_f are positive (to capture the trends in figure 1). As a consequence all time constraint will be slack in equilibrium and I will focus throughout on interior solutions. Also note that since I am focusing on a representative couple, the household can choose any non-negative continuous quantity of children b .

Combining (7) with (6), the effective budget constraint to the household's maximization problem can be written as

$$(10) \quad c_m + c_f + \left(\frac{\bar{x}(b)}{(z_m h_m^{1-\rho} + z_f h_f^{1-\rho})^{\frac{1-\gamma}{1-\rho}}} \right)^{\frac{1}{\gamma}} = w_m n_m + w_f n_f$$

Combining these optimality conditions implies

$$(11) \quad c_f = \frac{1-\theta}{\theta} c_m$$

and for the optimal time allocations

$$(12) \quad h_m = \left(\frac{z_m}{w_m} \right)^{\frac{1}{\rho}} \left(\frac{1-\gamma}{\gamma} \right)^{\gamma} \left[z_m^{\frac{1}{\rho}} w_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} w_f^{\frac{\rho-1}{\rho}} \right]^{\frac{\gamma\rho-1}{1-\rho}} \bar{x}(b)$$

$$(13) \quad h_f = \left(\frac{z_f}{w_f} \right)^{\frac{1}{\rho}} \left(\frac{1-\gamma}{\gamma} \right)^{\gamma} \left[z_m^{\frac{1}{\rho}} w_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} w_f^{\frac{\rho-1}{\rho}} \right]^{\frac{\gamma\rho-1}{1-\rho}} \bar{x}(b)$$

$$(14) \quad n_m = 1 - h_m - \left(\frac{\phi_l c_m}{w_m} \right)^{\frac{1}{\eta}}$$

$$(15) \quad n_f = 1 - h_f - \left(\frac{1-\theta}{\theta} \frac{\phi_l c_m}{w_f} \right)^{\frac{1}{\eta}}$$

The optimal fertility choice satisfies

$$(16) \quad \phi_b b^{-\sigma_b} = \underbrace{\frac{\theta}{c_m} \left(\frac{1}{\gamma} \right)^{\gamma} \left(\frac{1}{1-\gamma} \right)^{1-\gamma} \left[z_m^{\frac{1}{\rho}} w_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} w_f^{\frac{\rho-1}{\rho}} \right]^{\frac{(1-\gamma)\rho}{\rho-1}}}_{C_x(w_m, w_f)} \bar{x}'(b)$$

Equations (10) to (16) fully characterize the couple's utility maximizing choices. When choosing a fertility plan, the couple is outweighing benefits and costs from having children. The left-hand side in (16) is the marginal benefit of having an additional child, which at the optimum has to equal the marginal cost (in utility terms), the right-hand side. The term $C_x(w_m, w_f) = \left(\frac{1}{\gamma} \right)^{\gamma} \left(\frac{1}{1-\gamma} \right)^{1-\gamma} \left[z_m^{\frac{1}{\rho}} w_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} w_f^{\frac{\rho-1}{\rho}} \right]^{\frac{(1-\gamma)\rho}{\rho-1}}$ is the minimal cost to produce one unit of the home good x . The marginal cost of having more children lies in the need for more home production. To increase home production, the couple devotes more time to home labor and uses more purchased inputs. Both adjustments reduce consumption of the parents. Notice that $C_x(w_m, w_f)$ is increasing in the male and female wage rates as

these represent the opportunity cost of male and female home labor; devoting more time to home production lowers labor income and thereby reduces consumption. Since how costly it is to forgo consumption depends negatively on the original consumption level, the right-hand side of (16) is decreasing in c_m . Ceteris paribus, higher wages imply an increase in the cost of having children, whereas higher consumption implies a lower cost.

Also note that C_x is independent of the bargaining weight θ . The reason for this is that the spouses' relative time spent on home production is solely based on comparative advantage. From (12) and (13) it can be seen that the optimal ratio of male to female home hours is $\frac{h_m}{h_f} = \left(\frac{z_m w_f}{z_f w_m} \right)^{\frac{1}{\rho}}$.

3.3 Effect of Higher Female Wages

When female relative wages improve, the couple finds it optimal to allocate a larger share of home production to the husband. Holding the number of children and the required amount of the home good constant, this leads to a reallocation of men's time from market to home labor, and for women from home to market labor.

This impacts the marginal cost of having children in (16). Higher female wages directly increase the opportunity cost of home production, but they also improve labor income and thereby consumption. While the former channel tends to reduce fertility as the couple substitutes from having children towards consumption of market goods, the latter channel tends to increase it as the household experiences higher income and children are normal goods in the utility function. The first effect is the *substitution* effect and the second one the *income* effect.

To derive a condition under which the substitution effect dominates and fertility falls, use the parametrization of home production requirement (1) to write the optimal number of children as

$$(17) \quad b = \left(\frac{\phi_b c_m}{\kappa_1 \kappa_2 \theta} \underbrace{\gamma^\gamma (1-\gamma)^{1-\gamma} \left[z_m^\rho w_m^{\rho-1} + z_f^\rho w_f^{\rho-1} \right]^{\frac{(1-\gamma)\rho}{1-\rho}}}_{=1/C_x(w_m, w_f)} \right)^{\frac{1}{\sigma_b + \kappa_2 - 1}}$$

Note that (17) implies that, holding consumption constant fertility falls in response to higher wages (due to the substitution effect) only if $\sigma_b + \kappa_2 - 1 > 0$, which is a restriction on parameters I will assume throughout. Consumption itself is pinned down by the budget constraint (10) given the optimal choices according to (11) to (15). In general this gives a system of two non-linear equations which has no closed form solution.

On the one hand, higher female wages increase the cost of home production, which induces a substitution effect that implies a reduction in fertility. But on the other hand, they also have a positive income effect, which implies an increase in consumption and in fertility. When C_x increases by more than c_m , fertility falls. In the next subsection I

focus on a special case, for which one can derive a closed-form solution for the couple's optimal number of children. I show under what conditions fertility falls in response to higher female wages, and when it does not. In the subsection thereafter I show that the drawn conclusions generalize.

3.3.1 Special Case: $\sigma_b = 1$ and $\phi_l = 0$

Consider the special case of $\sigma_b = 1$ and $\phi_l = 0$. This means that agents have log utility from children and do not value leisure time, such that $n_g + h_g = 1$ for $g \in \{m, f\}$. These assumptions allow to derive a closed form solution. The optimal fertility choice is

$$b = \left(\frac{\phi_b}{\kappa_1 \kappa_2 + \phi_b} \frac{1}{\left((w_m + w_f) \gamma^\gamma (1 - \gamma)^{1-\gamma} \left[z_m^\rho w_m^{\rho-1} + z_f^\rho w_f^{\rho-1} \right]^{\frac{(1-\gamma)\rho}{1-\rho}} - \kappa_0 \right)} \right)^{\frac{1}{\kappa_2}}$$

The sign of $\frac{\partial b}{\partial w_f}$, which captures how the optimal fertility choice responds to increasing female wages, is given by

$$(18) \quad \begin{aligned} & z_f^\rho w_f^{\rho-1} + z_m^\rho w_m^{\rho-1} - (1 - \gamma) z_f^\rho w_f^{\rho-1} - (1 - \gamma) z_f^\rho w_m w_f^{\rho-1} \\ & = \gamma z_f^\rho w_f^{\rho-1} + w_m \left(\left(\frac{z_m}{w_m} \right)^{\frac{1}{\rho}} - (1 - \gamma) \left(\frac{z_f}{w_f} \right)^{\frac{1}{\rho}} \right) \end{aligned}$$

No marketization: First, consider a scenario in which there is no marketization of home production by setting $\gamma = 0$. Then $\frac{\partial b}{\partial w_f} < 0$ if and only if $\frac{w_f}{w_m} < \frac{z_f}{z_m}$. Hence, higher female wages reduce fertility when women's relative wages are low compared to their relative productivities at home. The reason is that when women's comparative advantage lies in home production, the female share in the cost of children is larger than the female share in total income. This implies that higher female wages increase the cost of having children by more than household income. As a consequence, the substitution effect dominates the income effect, and fertility falls.

No male home production: Next, consider a scenario with marketization but without men contributing to home production, which is nested as the optimal choice when $z_m = 0$. Here, $\frac{\partial b}{\partial w_f} < 0$ if and only if $\frac{w_f}{w_m} < \frac{1-\gamma}{\gamma}$. Thus, when the degree of marketization, as measured by γ , is sufficiently low, higher female wages reduce fertility. When the usage of time-saving inputs is limited, the cost of home production increases by more than income. Intuitively this happens since there is a lack of possibilities to exploit the benefits from higher female wages. Since there is no substitution with men's time and substitution with purchased inputs is limited, women's home hours cannot fall much; as a consequence her labor supply and the household's income do not increase much. As a result, the substitution effect is stronger than the income effect. If one assumes that in the long-run female and male wages are equalized, a sufficient condition for fertility to always fall when female wages go up is $\gamma < \frac{1}{2}$. Hence, if there is no male home production

and the degree of marketization is sufficiently low, fertility always falls when women's wages improve.

With marketization and male home production: When $\gamma > 0$ and $z_m > 0$, and the home production technology allows both for marketization and for substitutability between male and female time, $\frac{\partial b}{\partial w_f} < 0$ if and only if

$$(19) \quad \gamma \frac{w_f}{w_m} + \left(\frac{z_m}{z_f} \right)^{\frac{1}{\rho}} \left(\frac{w_f}{w_m} \right)^{\frac{1}{\rho}} < 1 - \gamma$$

Since the left-hand side is increasing in $\frac{w_f}{w_m}$, whereas the right-hand side is constant, this inequality is satisfied if and only if $\frac{w_f}{w_m}$ is sufficiently low. This means that improvements in female wages reduce the optimal number of children only if female relative wages are sufficiently low, echoing the conclusions drawn above. Notice that this inequality is more likely to hold when men's relative productivity in home production is high. To understand how the degree of substitutability between male and female time affects this condition, consider first a scenario where men and women have equal productivities at home ($z_m = z_f$). In this case, the left-hand side of (19) increases in ρ for $\frac{w_f}{w_m} < 1$. This implies that for larger ρ , i.e. when there are stronger complementarities between male and female time in home production, the fertility decline ends at a lower female relative wage threshold. When male and female home productivities differ, there is an additional effect through the term $\left(\frac{z_m}{z_f} \right)^{\frac{1}{\rho}}$. When $z_m < z_f$, which is a reasonable assumption for some child care tasks, this term depends positively on ρ , resulting in an additional channel through which the left-hand side of (19) increases in ρ .¹⁷ This implies that the larger the complementarities between male and female time at home, the earlier the fertility decline ends in response to improvements in women's relative wages. With larger complementarities, the household gains more from increasing the husband's home hours when the wife lowers hers; as a consequence the home production cost, C_x , increases less when w_f rises.

The condition for the fertility decline to end is also more likely to be satisfied when the degree of marketization is higher. A larger γ means that the additional margin, the possibility to substitute between parental time and purchased inputs, is more important, which weakens the strength of the overall substitution effect on fertility relative to the income effect. From inequality (19) one can derive a sufficient condition that ensures a bottoming out of fertility. Fertility stops falling, and potentially starts increasing, before female and male wages are equalized, only if

$$(20) \quad \left(\frac{z_m}{z_f} \right)^{\frac{1}{\rho}} > 1 - 2\gamma$$

¹⁷When $z_m > z_f$ and men are more productive than women in home production, the opposite result emerges and a larger ρ reduces $\left(\frac{z_m}{z_f} \right)^{\frac{1}{\rho}}$, making it less likely that inequality (19) is satisfied.

To summarize, fertility falls in response to higher female wages only if the gender wage gap is large. Moreover, both the substitution of female time spent in home production with male time as well as the substitution with inputs purchased in the market are important. The fertility decline ends before the gender wage gap is fully closed, when men's inputs to home production sufficiently well complement female time and when there is a sufficient degree of marketization. Obviously, once fertility stops falling, it in general increases. Whether the model implies that overall fertility looks rather stable or more like a pronounced U-shape, is a quantitative question whose answer depends on other parameters, and on which I return in section 6.1.

3.3.2 Log-Linearization of the General Case ($\sigma_b \neq 1$)

To study the effect of higher female wages on the fertility decision in the more general case with $\sigma_b \neq 1$, I rely on a first-order approximation. For tractability I continue to assume $\phi_l = 0$, implying $n_g + h_g = 1$ for $g \in \{m, f\}$.

Given the optimal choices for female consumption and for the time allocations as a function of male consumption, according to (11) to (15), the budget constraint (10) implies

$$(21) \quad \frac{c_m}{\theta} + \underbrace{\left(\frac{1}{\gamma}\right)^\gamma \left(\frac{1}{1-\gamma}\right)^{1-\gamma} \left[z_m^\rho w_m^{\frac{\rho-1}{\rho}} + z_f^\rho w_f^{\frac{\rho-1}{\rho}} \right]^{\frac{(1-\gamma)\rho}{\rho-1}}}_{C_x(w_m, w_f)} \bar{x}(b) = w_m + w_f$$

which pins down c_m for a given number of children b . The household chooses b optimally according to (17). These two equations fully characterize the optimal choices.

As there is no closed form solution, I log-linearize the two equations. In particular, I consider the log deviations of a variable between this period and the previous one. To keep the notation simple, I denote previous period's values by a tilde and log deviations between two periods by a hat; for a generic variable Z , the previous period's values is $\tilde{Z} = Z_{t-1}$, and the log deviations between the two periods is $\hat{Z} = \log(Z) - \log(\tilde{Z}) \approx \frac{Z_t - Z_{t-1}}{Z_{t-1}}$.

The log-linearization of the optimality condition for the fertility choice (17) gives

$$\hat{b} = \frac{1}{\sigma_b - \kappa_2 - 1} \left(\hat{c}_m - \hat{C}_x \right)$$

and of the reduced-form budget constraint (21)

$$\hat{C}_x + \frac{\bar{x}'(\tilde{b}) \tilde{b}}{\tilde{b}} \hat{b} = \frac{\tilde{w}_m \tilde{w}_m + \tilde{w}_f \tilde{w}_f - \frac{\tilde{c}_m}{\theta} \tilde{c}_m}{\tilde{w}_m + \tilde{w}_f - \frac{\tilde{c}_m}{\theta}}$$

where

$$\hat{C}_x = \frac{1-\gamma}{z_m^\rho \tilde{w}_m^{\frac{\rho-1}{\rho}} + z_f^\rho \tilde{w}_f^{\frac{\rho-1}{\rho}}} \left(z_m^\rho \tilde{w}_m^{\frac{\rho-1}{\rho}} \hat{w}_m + z_f^\rho \tilde{w}_f^{\frac{\rho-1}{\rho}} \hat{w}_f \right)$$

Solving these equations for $\widehat{w}_m = 0$, gives for \widehat{b} (the log change in fertility) as a function of \widehat{w}_f (the log change in women's wages),

$$\widehat{b} = \frac{1}{\frac{\widehat{c}_m}{\theta}(\sigma_b - \kappa_2 - 1) + \widehat{x}'(\widehat{b})} \frac{1}{\widehat{w}_m + \widehat{w}_f - \frac{\widehat{c}_m}{\theta}} \left(\widehat{w}_f - \frac{(1 - \gamma) z_f^{\frac{1}{\rho}} \widehat{w}_f^{\frac{\rho-1}{\rho}}}{z_m^{\frac{1}{\rho}} \widehat{w}_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} \widehat{w}_f^{\frac{\rho-1}{\rho}}} (\widehat{w}_m + \widehat{w}_f) \right) \widehat{w}_f$$

The sign of $\frac{\widehat{b}}{\widehat{w}_f}$ is therefore given by the term in parentheses, which can re-written as

$$\frac{1}{\widehat{w}_f} \frac{1}{z_m^{\frac{1}{\rho}} \widehat{w}_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} \widehat{w}_f^{\frac{\rho-1}{\rho}}} \left(z_m^{\frac{1}{\rho}} \widehat{w}_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} \widehat{w}_f^{\frac{\rho-1}{\rho}} - (\widehat{w}_m + \widehat{w}_f) (1 - \gamma) z_f^{\frac{1}{\rho}} \widehat{w}_f^{\frac{\rho-1}{\rho}} \right)$$

As the terms outside of these parentheses are necessarily positive, the crucial term is the one inside, which is the very same as in condition (18). Therefore, the conclusions drawn in subsection 3.3.1 above generalize to cases with $\sigma_b \neq 1$, as long as $\sigma_b > 1 - \kappa_2$ (which is the parameter restriction ensuring a normal substitution effect). Higher female relative wages lower fertility when the gender wage gap is large, but the fertility decline might end before the gap is closed, depending on the properties of the home production technology. In the quantitative model calibrated in the next section, I find that these results also go through when allowing for $\phi_l \neq 0$.

4 Calibration

To quantitatively assess the effect of the narrowing of the gender wage gap on fertility, I calibrate the model. I choose parameters such that the model replicates in 1965 features of the data for married couples. I target married men's and married women's hours worked, both at home and in the market, and an indicator for married fertility. All parameters are time-invariant, and the only exogenous change over time is in the gender wage gap. As the Aguiar and Hurst (2007) data is available once per decade, I study the choices of representative couples born ten years apart.

Below I first discuss what data is targeted and then the choice of parameters. I calibrate the model in the following order: the home production function, the required amount of the home good, and the preference parameters.

4.1 Targets

To construct the targets, I need to map the hours worked data of figures 1 and 4 into the model. Since the model is about the decisions of prime-aged couples, I restrict the sample to married individuals of age 20 to 60. In the model, agents have a time endowment of one, which they can split between market work, home work, and leisure time. Assuming that

people need 8 hours of rest a day, 2/3 of time is discretionary. Hence to map the weekly data from Aguiar and Hurst (2007) into my model, I divide all hours by $(2/3 \times 24 \times 7)$. For market hours worked, I use data from the Panel Study of Income Dynamics (PSID) from 1968 to 2007, which I average over 10 year periods.¹⁸ Since in the PSID market hours worked are given per year, I divided these by $(2/3 \times 24 \times 365)$. Not to confound trends in market and home hours with potential composition effects that could be due to demographic changes, I first compute average hours for individuals at different ages in 5 year bins, and then take the simple average over the bins, which gives every age group the same weight.

As the theory I propose applies to couples, the model’s implications are for married fertility. Since the total fertility rate is constructed by age-specific birthrates, it does not allow for a break-down into unmarried and married fertility. I therefore use children-ever born to women at age 40–44, for which there is data by marital history (shown in the right panel of figure 2), published by the United States Census Bureau in the Fertility Data Historical Time Series¹⁹. This is an indicator of completed fertility. However, most couples have children at an earlier point in their lives than at age 40–44 (and thereby take a decision under earlier wages). To construct an indicator for married fertility that corresponds to the model, I shift the series of children-ever born. I take the average number of children an ever-married women at age 40–44 ever had, and shift it back by 16 years, since the mean age of a mother when giving birth is 25.83 in the data.²⁰ The drawback of this indicator for the married fertility choice is that the constructed series ends early; the model’s predictions cannot be evaluated against constructed data for 2005 (yet). However, one advantage of using this measure of completed fertility over other fertility indicators, such as the total fertility rate, is that it is not confounded by changes in the timing of births.

For the gender wage gap I take the ratio of median female to median male wages in the PSID, shown in figure 3. To feed it into the model, where a period corresponds to 10 years, I take simple averages.

4.2 Home Production Technology

I set the elasticity of substitution between male and female time spent in home production to replicate the rise in the ratio of married men’s home hours to married women’s observed in the data. In the model, equations (12) and (13) imply for relative male home hours $\frac{h_{m,t}}{h_{f,t}} = \left(\frac{z_m w_{f,t}}{z_f w_{m,t}}\right)^{\frac{1}{\rho}}$; their ratio between periods T and 0 is therefore $\frac{h_{m,T}/h_{f,T}}{h_{m,0}/h_{f,0}} =$

¹⁸I discuss the advantages of using data from the PSID, rather than from the CPS, and compare the data from these two surveys in section A of the online appendix.

¹⁹<http://www.census.gov/hhes/fertility/data/cps/historical.html>

²⁰The average age of a mother (total over all birth orders) was 24.9 in 1968 and 27.4 in 2003, according to Vital Statistics of the United States, 2003, Volume I, Natality. The average over 1968–2003 is 25.83.

$\left(\frac{w_{f,T}/w_{m,T}}{w_{f,t}/w_{m,t}}\right)^{\frac{1}{\rho}}$. To match the change in married men's relative home hours over 1965 to 2003 from 0.25 to 0.54, given the change in female relative wages from 0.57 to 0.78, this requires for the inverse elasticity of substitution $\rho = 0.4037$. In turn this implies for the relative home productivities $\frac{z_m}{z_f} = \left(\frac{h_{m,1965}}{h_{f,1965}}\right)^{\rho} \frac{w_{m,1965}}{w_{f,1965}}$. Restricting $z_m + z_f = 1$, men's productivity in home production follows as $z_m = 0.496$ and women's as $z_f = 0.504$.

In the literature the range of estimates for the labor share in home production is very wide. Studies that include housing as capital or equipment used for home production typically find a relatively low value, close to the one of market production, e.g. Greenwood, Rogerson, and Wright (1995), while Benhabib, Rogerson, and Wright (1991), who exclude housing, estimate a very high value of 0.92. In my model, the need for home production arises at the margin only from having children (in the household) and does not correspond closely to either study. Since parents can acquire home production inputs in the market, such as hiring nannies or paid domestic help, the share of time-saving inputs acquired in the market, γ , should be higher than the Benhabib, Rogerson, and Wright (1991) value. As a benchmark I consider an intermediate value of $\gamma = 0.29$, but I conduct a series of robustness checks in section 6.1.

4.3 Home Production Requirement

The required amount of home production (1), $\bar{x}(b) = \kappa_0 + \kappa_1 b^{\kappa_2}$ is calibrated in two steps, conditional on the home technology parameters pinned down earlier. First I choose κ_0 to ensure that the level of home hours are matched by the model. Second, I choose the additional home production required when having children, $\kappa_1 b^{\kappa_2}$, to replicate some cross-sectional variation in home hours.

While the parameterization of the home production technology ensures that relative male home hours are replicated by the model, I choose κ_0 to match in 1965 the level of male time spent on home production (as a fraction of the time endowment of one). This yields $\kappa_0 = 0.1693$.

To calibrate by how much the need for home production increases due to children, I make use of the observed variation of married men's and women's home hours against the number of children in the household. In the model, equations (12) and (13) imply for aggregated parental time input (3) $(z_m h_m^{1-\rho} + z_f h_f^{1-\rho})^{\frac{1}{1-\rho}} = \left(\frac{1-\gamma}{\gamma}\right)^{\gamma} \left[z_m^{\frac{1}{\rho}} w_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} w_f^{\frac{\rho-1}{\rho}} \right]^{\frac{\gamma\rho}{1-\rho}} \bar{x}(b)$, which can be rearranged to give

$$(22) \quad \kappa_1 b^{\kappa_2} = \frac{(z_m h_m^{1-\rho} + z_f h_f^{1-\rho})^{\frac{1}{1-\rho}}}{\left(\frac{1-\gamma}{\gamma}\right)^{\gamma} \left[z_m^{\frac{1}{\rho}} w_m^{\frac{\rho-1}{\rho}} + z_f^{\frac{1}{\rho}} w_f^{\frac{\rho-1}{\rho}} \right]^{\frac{\gamma\rho}{1-\rho}}} - \kappa_0$$

This relationship can be used to infer the home production requirement from the Aguiar and Hurst (2007) data, which includes not only time use information, but for some

years also respondents' wages and the number of children in the household. Conditional on parameters already calibrated, I can therefore construct the right-hand side of (22). In particular, I take Aguiar and Hurst's data for 1985 and take averages of male and female home hours and wages by the number of children in the household. To use the hours data in the model, I transform them into fractions of the total time endowment of one. To map the wage data, I divide them by the mean of male wages in this sample, since in the model I fix men's wages to unity (and only let the gender wage gap vary). I construct (22) for households with one child ($b = 1$) and with two children ($b = 2$) to have a set of two equations in the two unknowns κ_1 and κ_2 . Solving these gives $\kappa_1 = 0.0151$ and $\kappa_2 = 0.419$, implying that the need for home production increases in the number of children, but at a decreasing rate.

4.4 Preference Parameters

The four parameters of the utility function (4) are left to be chosen. Two parameters, ϕ_l and η , relate to the preference for leisure, and two parameters, ϕ_b and σ_b , are key for the fertility choice. In 1965, the base year, there are three targets left, male labor supply, female labor supply, and fertility.

As home hours are already determined by the parameters set above, targeting market labor is essentially targeting leisure time. In the model, male and female labor supply are the solution to (14) and (15), where c_m solves (10) and $\theta = 0.5$. Therefore I try to replicate male and female time in paid employment, as function of the time endowment of one, by choosing the leisure preference parameters. This results in $\phi_l = 0.0059$ and $\eta = 10$.²¹

Finally, the two parameters governing fertility have to be chosen. While ϕ_b captures the relative weight the household attaches to having children (compared to consumption), σ_b essentially captures the curvature, which is important for the magnitude by which fertility responds to changes in female wages. Since 1965 fertility is only one target, I need to add a second one. I restrict the calibration by choosing σ_b such that the lowest fertility the model generates in response to the narrowing of the gender wage gap is equal to the lowest married fertility observed, 2.044 (which in the data corresponds to 1985). Conditional on a value for σ_b , I calibrate ϕ_b to the fertility data for 1965. Then I solve the model's transition in response to a catching up of women's relative wages. If the implied minimum fertility is different to the observed one, I update the guess for σ_b until this target is met. This leads to $\phi_b = 0.0500$ and $\sigma_b = 0.6434$. It is worth noting that I do not impose that the model has minimum fertility in the same year as in the data. In fact, under the

²¹While the model gets close, it does not match the 1965 labor supply data perfectly. A perfect match cannot be attained since ϕ_l and η impact n_m and n_f symmetrically. However, this deviation is likely to exert only a minor effect on the fertility choice, as the cost of having children does not depend on market hours worked.

parameters chosen as the baseline calibration, the model matches the lowest fertility when $\frac{w_f}{w_m} = 0.7614$, which is in between the values of the gender wage gap in 1995 and 2005. The analytical derivations of section 3.3.1 show that the condition (20), for whether fertility stops declining or not, is independent of σ_b (at least when there is no leisure margin). In section 6.1 I perform a series of robustness checks, including alternative values for σ_b (and therefore of ϕ_b), which highlight that the value of σ_b matters for how strongly fertility decreases or increases, but not for when the fertility decline ends.

Table 2 summarizes all parameters and table 3 the targeted statistics alongside their model analogue.

Table 2: Model parameters

	Description	Value
ρ	inv. substitutability at home	0.4037
z_m	male home labor productivity	0.496
z_f	female home labor productivity	0.504
γ	share of market inputs in home	0.29
$\bar{x}(b)$	amount of home good needed	$0.1693 + 0.0151 b^{0.419}$
ϕ_l	weight on utility from leisure	0.0059
η	curvature of leisure in utility	10
ϕ_b	weight on utility from children	0.0500
σ_b	curvature of fertility in utility	0.6434
θ	husband's Pareto share	0.5

Table 3: Calibration targets (all for married couples)

1965-Moments to be Matched	Data	Model
Fertility	2.973	2.973
Home hours of men	0.097	0.097
Home hours of women	0.389	0.391
Market hours of men	0.380	0.405
Market hours of women	0.108	0.082
Additional Target to be Matched		
Minimum fertility	2.044	2.044

5 Quantitative Results

5.1 The Equilibrium as a Function of the Gender Wage Gap

To illustrate the model's predictions, I first show in figure 5 the representative couple's optimal choices as a function of female relative wages (on the horizontal axis).

When female relative wages are low, the couple household is very specialized. Based on comparative advantage, the wife spends much more time than the husband working

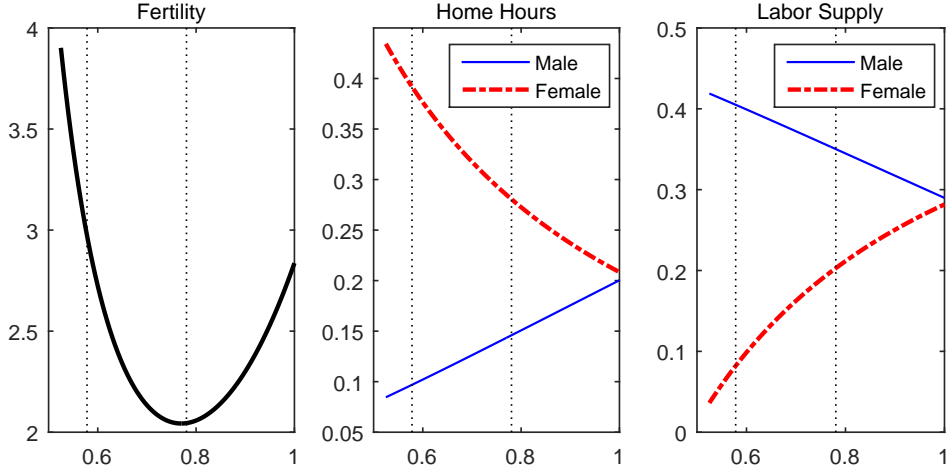


Figure 5: Benchmark model's outcomes against female relative wages

This graph shows the benchmark model's outcomes against female relative wages (on the horizontal axis). The vertical lines correspond to two observed values of female relative wages, the dashed line shows the low value of 1965 and the dotted line the higher value of 2005. Home hours and labor supply are shown as fractions of the time endowment of one.

in the household, and much less time in paid employment. The opportunity cost of having children then depends mainly on the female wage. Improvements in female wages therefore raise the couple's opportunity cost of children substantially. However, they also increase her labor supply and the household's income. Yet, since she worked initially rather little in the market, the impact on household income is relatively small. Overall, the calibrated model therefore implies that when female relative wages are low and improve, the channel through higher cost is more important, and the substitution effect dominates the income effect, such that fertility falls. While the couple optimally reallocates some of the husband's time to home production, this change in the division of labor helps only a little to curb the increase in the cost of children, since his wage rate is substantially higher than hers.

On the contrary, when female relative wages are high and improve further, the cost of having children increases by less. On the one hand, reassigning tasks to the husband is relatively cheaper when the gender gap is smaller. On the other hand, the wife's income is a larger fraction of family income, and an improvement in the female wage has a larger effect on household resources, which allows for larger substitution of parental time with purchased inputs. This means that at higher female relative wages, the substitution effect on the fertility choice becomes weaker relative to the income effect.

In the calibrated model, the income effect overturns the substitution effect before the gender wage gap is fully closed. As a consequence, fertility is U-shaped in figure 5. Fertility stabilizes in the sense that the initial decline slows down, and ultimately turns into an incline. This echoes the analytical results derived in 3.3.1. There I showed, in the absence of a leisure margin, that fertility stops falling before the gender wage gap is closed if inequality (20) is satisfied. Given the calibrated parameters, $\left(\frac{z_m}{z_f}\right)^{\frac{1}{\rho}} > 1 - 2\gamma$,

condition (20) holds and fertility is U-shaped in female relative wages.

5.2 The Transition in the Benchmark Model

To compare the model's prediction against the data, I feed the series of the observed gender wage gap into the model. I also explore what the model would imply if the gap continued to close. Figure 6 shows the implied path of the economy since 1965. During the transition, the gender wage gap narrows, as shown in the upper-left panel. The dotted vertical line corresponds to the last value of the gender wage gap from the data; to the right of it, the graph shows predictions if the gap continued to close at the same rate as it did over 1965–2005. It should be stressed that the forward projection is a hypothetical exercise, and it very well might be that the gender wage gap might not shrink any further. In fact, some authors (e.g. Blau and Kahn (2016)) have argued that there is evidence for a slowdown, or perhaps even complete stagnation, in the narrowing of the gender wage gap in recent years.

As female relative wages improve, the division of labor changes. Initially, fertility declines as raising children becomes more costly to the parents. However, fertility bottoms out after an initial decline. While the gender wage gap is closing, the model implies a reallocation of labor across gender. When female wages rise, a household finds it optimal to increase the wife's labor supply and to decrease her time working at home. However, since relative wages have changed, but not relative productivities at home, this reallocation entails an increase in men's home production. Since the opportunity cost of working at home overall increases, the couple also increases their purchases of time-saving inputs by 12.8 percent between 1965 and 2005. The link between higher female relative wages and lower fertility breaks at some point. In the calibrated model fertility stabilizes around 1995 and increases from 2005. As argued in section 5.1 above, the reason is that due to the optimal substitutions between the inputs to home production, at sufficiently low gender wage gaps the implicit cost of having children increases by less than household income.

I will show in the next two sections that both the substitution between male and female time as well as marketization are key in understanding why fertility did not fall further.

5.3 Counterfactual: The Absence of Male Home Labor

In this part I shut down the rise of male home labor. In the existing literature on fertility it is commonly assumed that child-care is a function of female time only. By setting male home productivity to zero, my model nests this as a special case. Notice that in this counterfactual exercise, there will be no convergence in the time use of men and women. Even if wages were equalized, specialization in the household would persist as men are

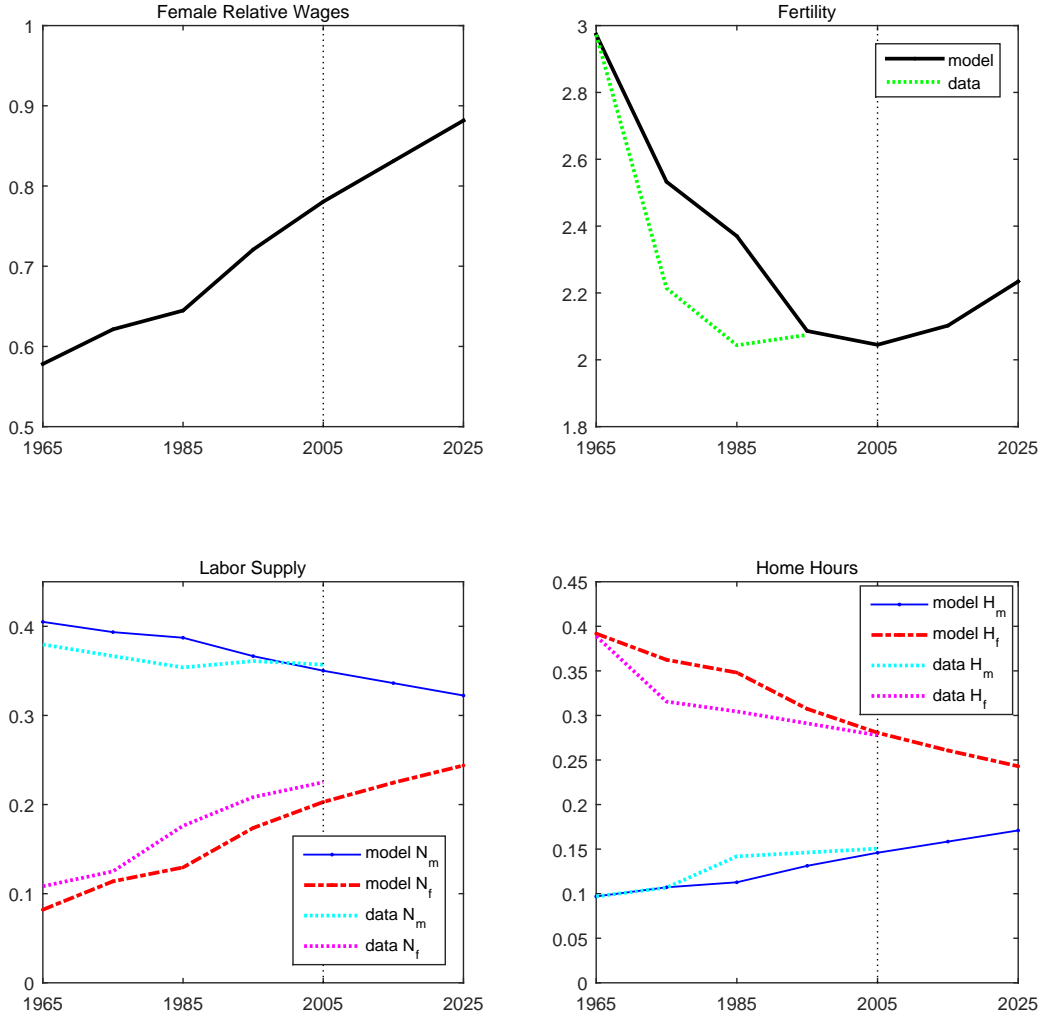


Figure 6: Transition path of the benchmark model

This graph shows the benchmark model’s transition over time that is implied by the narrowing of the gender wage gap (upper-left panel). The dotted vertical line corresponds to the last value of the gender wage gap from the data; to the right of it, the graph shows predictions if the gap continued to close. Home hours and labor supply are shown as fractions of the time endowment of one. The dotted lines show the data.

not productive at home.

As I had calibrated the baseline model under the condition that $z_m + z_f = 1$, I set in the counterfactual $z_m = 0$ and $z_f = 1$. Then I adjust the weight on children in the utility function such that the household still chooses in 1965 the same fertility as in the data.²² This leads to $\phi_b = 0.0168$. Figure 7 shows the transition of the model, when men cannot counteract the fall in women’s home hours. It implies a monotone drop in fertility over the entire series– which is inconsistent with the data. Throughout the transition, as female labor supply rises, parents are having fewer and fewer children as child-care hours continue to fall. As the wife’s income increases, a couple also acquires more time-saving

²²Since by assumption of $z_m = 0$ this model variant is doomed to fail along the dimension of male home hours, notice that I could not apply the calibration strategy of the benchmark model, which has H_m as one of the targets.

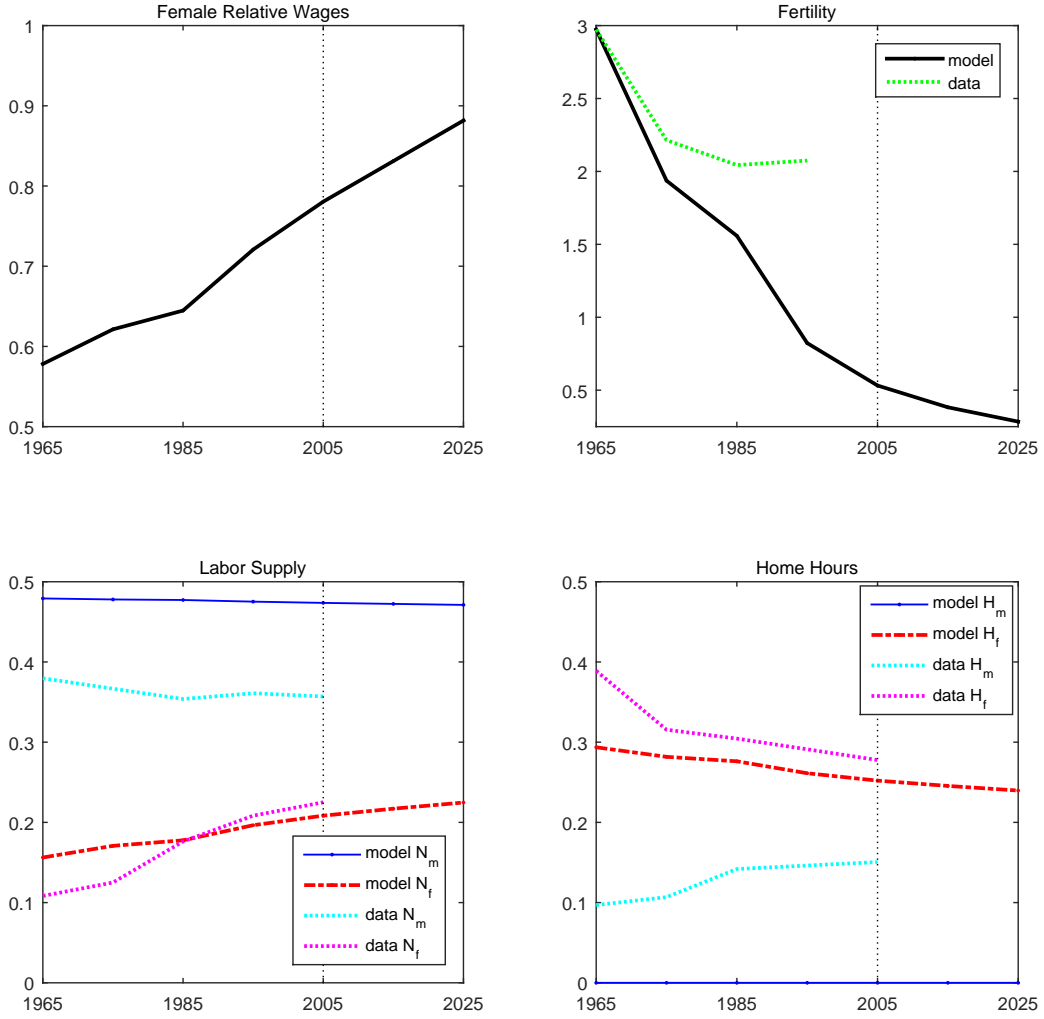


Figure 7: Transition of the model with $z_m = 0$

For the model without male home production, this graph plots the transition implied by the narrowing of the gender wage gap (upper-left panel) over time. The dotted vertical line corresponds to the last value of the gender wage gap from the data; to the right of it, the graph shows predictions if the gap continued to close. Home hours and labor supply are shown as fractions of the time endowment of one. The dotted lines show the data.

inputs, but this marketization of home production is not strong enough to prevent fertility from falling.²³

In fact, the calibrated model with $z_m = 0$ implies a fall in fertility even when female relative wages are close to, or equal to, one. In section 3.3.1, I derived that in the absence of male home production, the optimal number of children falls if $\frac{w_f}{w_m} < \frac{1-\gamma}{\gamma}$, which would require a $\gamma > 0.5$ to imply a stabilization in fertility before the gender wage gap is closed. In the model parameterization, however, $\gamma = 0.29$ and, as a consequence, fertility falls throughout. Since the degree of marketization is too low, and men are not productive at

²³In this counterfactual fertility drops in the long-run to a very low value. The result that fertility keeps falling is not driven by the lack of imposing a restriction on the behavior of fertility on the transition path (which pinned down σ_b as a parameter of the baseline calibration). Section B.1 of the online appendix shows that also under an alternative calibration, which imposes this extra target, fertility falls as long as women's relative wages improve.

home, the cost of children increases by more than household income when female wages go up. The lack of good substitutes for women’s home hours curbs the increase in female labor supply and in total income. As a result, fertility always falls when men do not contribute to home production and the gender wage gap narrows.

5.4 No Marketization of Home Production

In this section I show that marketization alone is not sufficient to generate a recovery of fertility.

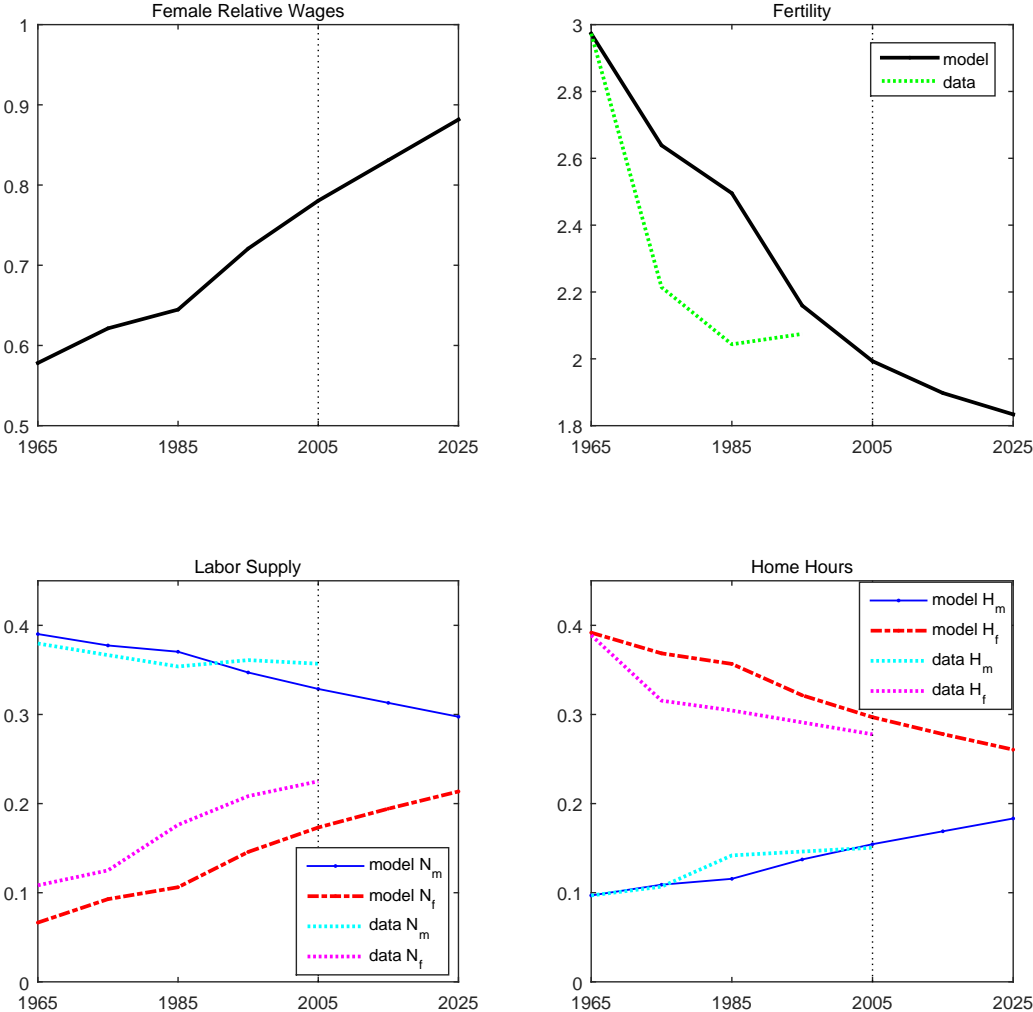


Figure 8: Transition path of the model with $\gamma = 0$

For the model without marketization, this graph plots the transition implied by the narrowing of the gender wage gap (upper-left panel) over time. The dotted vertical line corresponds to the last value of the gender wage gap from the data; to the right of it, the graph shows predictions if the gap continued to close. Home hours and labor supply are shown as fractions of the time endowment of one. The dotted lines show the data.

To rule out marketization of home production, consider a simpler version of the model in which home production is a function of male and female home hours only. This is the special case of $\gamma = 0$. Following the calibration strategy used for the benchmark model,

I adjust some parameters. First, to ensure that this version of the model matches the home hours targeted in 1965, I adjust κ_0 to 0.1976. Then I choose the fertility preference parameters such that the model replicates the 1965 fertility statistics as well as a long-run target for fertility. In the calibration of the benchmark, I required the model to generate on the transition path the same value of minimum fertility as observed in the data (2.044), which the baseline model delivers when $\frac{w_f}{w_m} = 0.7614$. To make the two models comparable, I impose that the model with $\gamma = 0$ generates at this relative wage the same level of fertility, i.e. that $b = 2.044$ when $\frac{w_f}{w_m} = 0.7614$. Together with the 1965 fertility target this requires setting $\phi_b = 0.0329$ and $\sigma_b = 0.9016$.

Figure 8 shows the transition of this model variant. As in the baseline model, the narrowing of the gender wage gap implies a reallocation of home production from women to men, along with a rise in female market labor and a fall in men's. Fertility falls throughout. While the speed of the decline falls as female wages improve, fertility keeps declining. Through the lens of the derivations of section 3.3.1, in the absence of marketization fertility declines when $\frac{w_f}{w_m} < \frac{z_f}{z_m}$. Under the calibrated parameters, women have a comparative advantage in home production even when wages are equalized. As a consequence, the female share in the cost of providing the home good, C_x , is larger than 1/2, whereas the female share in household income is at most 1/2. Higher female wages therefore always increase the cost of having children by more than household income, and fertility falls.

From the two counterfactual exercises with $z_m = 0$ and with $\gamma = 0$, I conclude that neither marketization alone nor the rise of male home production by itself is sufficient to generate the stabilization of fertility, but that the interaction of the two is needed. The magnitude of γ is also important in matching the timing of the fertility transition, as I show in section 6.1.

6 Confronting the Models and the Data

In this section, I compare the predictions of the different versions of the model to each other and to the observed variation in the data. Figure 9 shows the transition paths of the various models and the data for married men and women, and table 4 lists the relative changes from 1965 to 1995, which is the last period for which I can construct married fertility.

The benchmark model performs better than the model version ruling out male participation at home ($z_m = 0$), and than the version not allowing for marketization of home production ($\gamma = 0$). All models imply initially a fall in fertility and that subsequently the decline slows down. But only the full model, which allows for male home production and marketization, implies that fertility stops falling before the gender wage gap has closed.

The $z_m = 0$ -model as well as the $\gamma = 0$ -model fail to generate a bottoming out of fertility before the gender wage gap has closed. From this I conclude that it is the

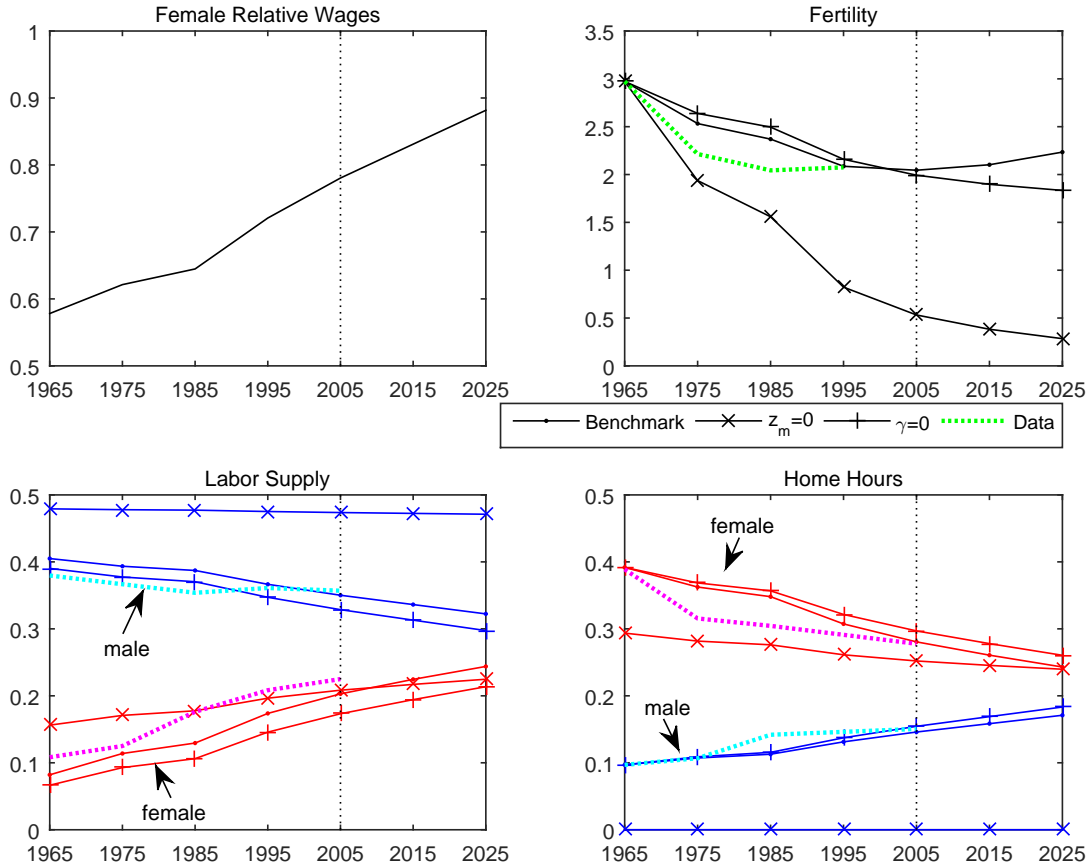


Figure 9: Comparing the models and the data

combination of men contributing more to home production and households purchasing more inputs in the market that explains why fertility did not fall further, despite the continued improvement in female wages. This also suggests that the analytical results derived in section 3.3 for the simplified model generalize to the full model.

In the $z_m = 0$ -model, which rules out male home production, fertility falls initially much faster than in the data. The model with $\gamma = 0$, which does not allow for marketization of home production or child care, overstates the number of children per couple for the first decades, but understates the rise in female labor supply and the drop in female home labor. However, even the benchmark models predicts that fertility would flatten out later than observed in the data. Under the baseline calibration the model predicts fertility to reach its lowest point in 2005, but in the data the lowest value is in 1985. Since then fertility showed a very modest incline, but virtually remained flat.

Thus, both the rise of male time and the acquisition of time-saving inputs into home production are important in understanding why the fertility decline ended. Comparing the $\gamma = 0$ -model to the benchmark, which uses $\gamma = 0.29$, also suggests that the higher the share of market inputs into child care, the earlier fertility flattens out. In section 6.1 below, I show a sensitivity analysis which confirm this.

Table 4: Relative changes (in percent) over 1965–1995

	Data	Model		
		Benchmark	$z_m = 0$	$\gamma = 0$
<i>Fertility</i>	-30.2	-29.8	-72.3	-27.4
N_m	-4.9	-9.5	-0.8	-11.1
N_f	+92.4	111.8	25.8	+119.2
H_m	+50.8	+35.4	0	+41.7
H_f	-25.2	-21.6	-11.0	-17.9

This table shows the relative changes of the variables in the different versions of the model and in the data over the period 1965–1995. The fertility indicator is ever-married 40-44 year old women’s number of children-born shifted back by 16 years.

6.1 Robustness Checks

First, I conduct a sensitivity analysis for the share of market inputs into home production, γ , since in the literature the range of estimates is wide, and in my model the interpretation of γ should include all market inputs into home production, not only appliances. In the benchmark model I set $\gamma = 0.29$, which is larger than the estimate by Benhabib, Rogerson, and Wright (1991) to allow also for a marketization of child care, but below the upper-end value of Greenwood, Rogerson, and Wright (1995), who include housing capital and conclude that market inputs’ share into home production does not differ from the capital share in market output (0.36). Since in my model home production is at the margin varying only with child care, which requires mainly human time, I only consider lower values ($\gamma < 0.36$).

In section 5.4, I have already shown the transition of the model with $\gamma = 0$. As a further sensitivity analysis, I also calibrate²⁴ models with different values for γ and solve for their transition path. Figure 10 shows the results. The higher γ , the earlier the fertility rate stabilizes, bringing the model even closer to the data. In this sense, the choice of $\gamma = 0.29$ in the benchmark calibration is therefore conservative. With $\gamma = 0.33$, the model matches the fertility response over 1965–1995 almost perfectly. However, this higher value of γ also implies that fertility responds more strongly at very low or very high female relative wages.

Second, I perform a robustness checks on σ_b , the elasticity of utility with respect to children. For the benchmark model I chose σ_b such that the model generates the lowest level of fertility seen in the data. Now I lift this restriction and instead use various ad-hoc values for σ_b . All other parameters are taken from the benchmark calibration except ϕ_b which is adjusted such that the model still matches the initial fertility in 1965. As shown in figure 11, the higher σ_b , the less fertility declines during the transition and the more stable it is throughout. To replicate the steep decline of fertility in the 1960s, however, the model needs a value for σ_b that is not too high. A very low value of σ_b , on the other

²⁴To calibrate the models with different γ I again use the calibration strategy of section 4.

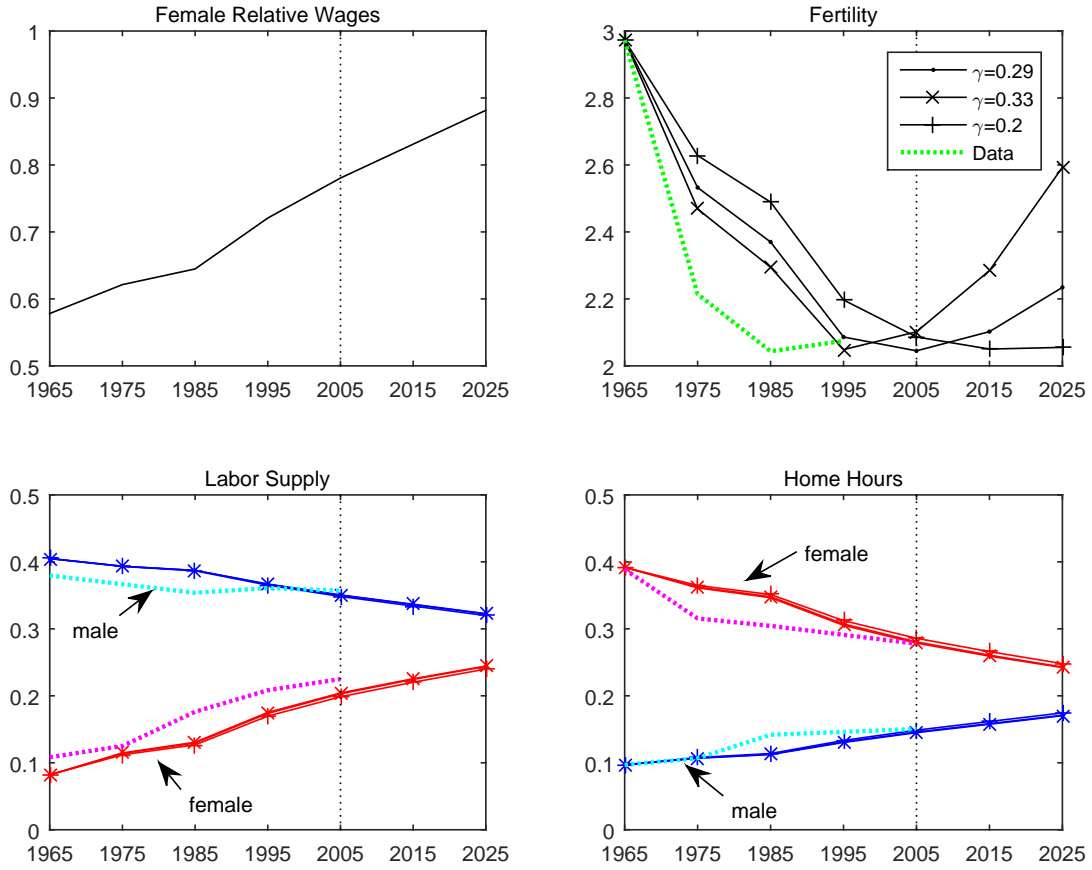


Figure 10: Sensitivity analysis: varying γ

This graph shows the transition of the model under different values for γ . The models are calibrated individually, using the strategy outlined in section 4. The baseline model has $\gamma = 0.29$.

hand, would imply a lower fertility than observed most recently. Most importantly, the value of σ_b has no effect on when the fertility decline stops, but only for the amplitude of negative or positive changes.

To sum up, whether the fertility decline ends, or not, depends on the home production technology. The conclusions from section 3.3 also apply to the full model that allows for a leisure margin. In response to higher female wages, fertility falls when the gender wage gap is large, but when the gap is small and there is a sufficient amount of complementarity in home production and the possibility of some marketization, fertility stabilizes. When γ is high, fertility is predicted to follow a pronounced U-shape when female relative wages improve further, whereas for smaller γ there is only a modest rise.

6.2 The 20th Century

To analyze how my explanation performs over the longer time series, I take the model's benchmark calibration (i.e. the parameters calibrated to 1965 as described above), and simulate the implied transition path when starting the economy in 1900. To do this, drawing on data sources from above is insufficient, as for one the time use surveys started

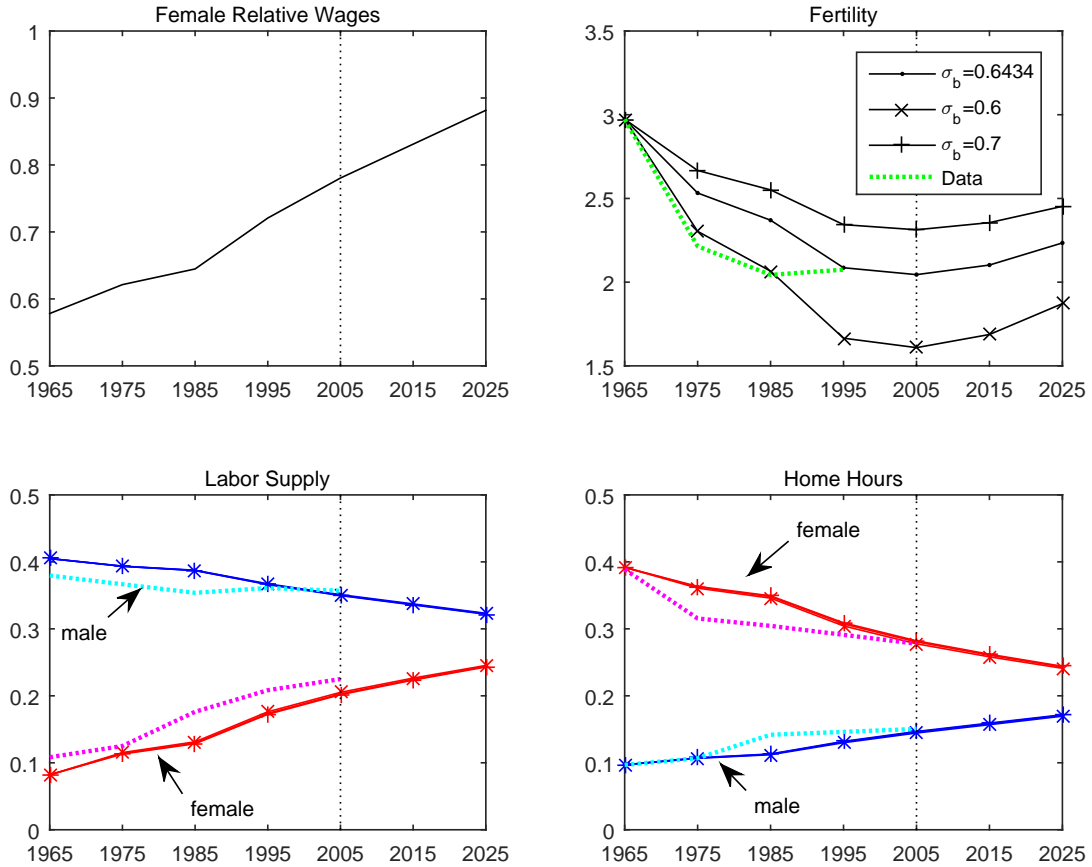


Figure 11: Sensitivity analysis: varying σ_b

This graph shows the transition of the model under different ad-hoc values for σ_b . All parameters but ϕ_b , which is adjusted to replicate the 1965 fertility, are taken from the benchmark calibration. The baseline model has $\sigma_b = 0.6434$.

only in 1965. I therefore make use of additional data sources. For a longer time series on home production I take data from Ramey (2009) and for market hours worked from Ramey and Francis (2009). As a fertility indicator I use Jones and Tertilt (2008)'s data for married women's children ever born. Since they construct this series by mothers' birth cohort, I shift it by 27 years. I do this as Jones and Tertilt report that cohort children ever born when shifted by 27 years displays similar trends to the total fertility rate. All these series are needed to evaluate the model predictions. In addition, I need pre-1965 data on relative wages to feed into the model. These I take from Fernández (2013), which are based on Goldin (1990).

Figure 12 shows the results. The model does well in replicating the overall trends in male and female work in the market and at home. More importantly, it performs well in terms of fertility. At the beginning of the century fertility dropped when female wages improved. But when in the aftermath of the second world war women's relative earnings on average declined, fertility temporarily increased, consistent with the baby boom we see in the data.²⁵ Subsequently, with the further improvement in female relative

²⁵This is similar to the findings of Doepke, Hazan, and Maoz (2015), who argue that WWII implied a

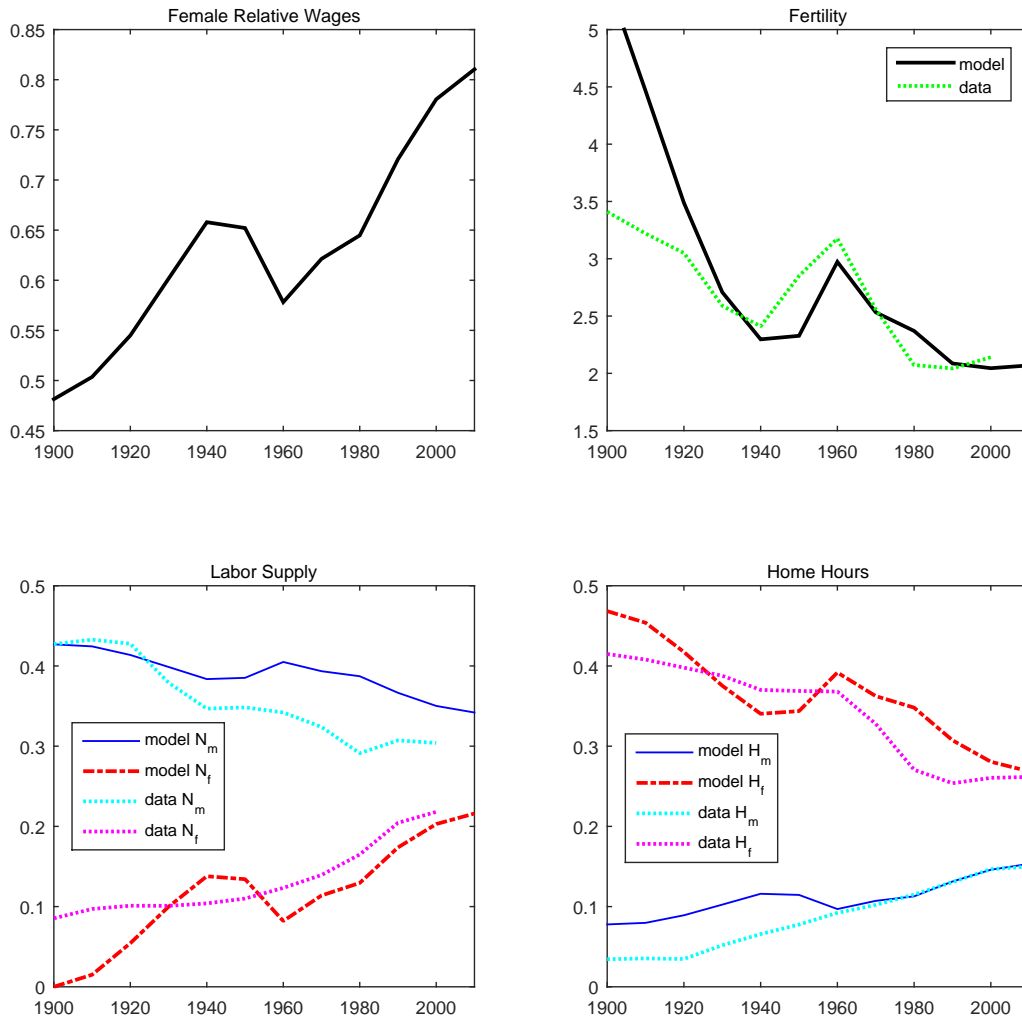


Figure 12: Transition path over the 20th century

This graph shows the model's transition path over the last century that is implied by the narrowing of the gender wage gap (upper-left panel). This is based on the benchmark calibration. Home hours and labor supply are shown as fractions of the time endowment of one. The dotted lines show the data.

wages fertility declined further. Hence, also in the first half of the last century the model performs well in replicating the data. It can explain the secular decline in fertility and yet allow for the baby boom, as both are linked to movements in women's relative wages at times when the gender wage gap was large, and women therefore shouldered most of home production and of the cost of having children. Then the fertility decision depended mostly on the market value of women's time. However, as argued in section 5 the model performs well also in the second half of the last century. Then the gender wage gap was already fairly small and when it shrank further, fertility flattened out as men started to supply more hours to home production.

negative shock to young women's wages after the war due to a competition with older women who had acquired more work experience during the war.

7 Conclusion

In this paper I argue that the common force behind the observed trends in fertility and hours worked is the narrowing of the gender wage gap. I present a model in which having children increases the need for home production, and in which rising female relative wages have not only direct effects on employment, but also reallocate hours worked at home from women to men. Initially, because of the gender wage gap, women shoulder most of home production. When the wage gap shrinks, women's labor supply increases and total home production falls. Since men's wages are considerably higher than their wives, they are not fully offsetting the drop in women's time working at home. Instead, fertility falls. However, the narrower the gender wage gap, the smaller the scope for specialization in the household. When female relative wages improve further, the opportunity cost of having children increases relatively little, as reallocating home hours from women to men is not as costly anymore. I show that the fertility decline ends when female relative wages are still improving, provided that complementarities between both parents' time working at home are sufficiently strong and it is possible for a couple to substitute some of their time for purchased inputs.

Qualitatively the model predictions are consistent with the data for the U.S. After an initial steep decline, fertility stabilizes before the gender wage gap has fully closed. The model also implies the secular rise of male home production and a fall in female time spent working at home. However, the model generates the end of the fertility decline later than observed in the data. I find that a higher use of parental time-saving inputs into home production helps in matching the timing. However, a marketization of home production alone is not sufficient. Both the rise in male home production and marketization are key in explaining why fertility did not decline further. Since in my model the condition for the fertility stabilization depends only on the home production and child care technology, it suggests that cross-country differences in whether the fertility decline ended, or not, might be due to differences in marketization rates.

References

- AGUIAR, M., AND E. HURST (2007): "Measuring Trends in Leisure: The Allocation of Time over Five Decades," *The Quarterly Journal of Economics*, 122(3), 969–1006.
- ATTANASIO, O., H. LOW, AND V. SANCHEZ-MARCOS (2008): "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *American Economic Review*, 98(4), 1517–52.
- BECKER, G. S. (1960): "An Economic Analysis of Fertility," in *Demographic and Economic Change in Developed Countries*, no. 11, pp. 209–231. Princeton University Press.

- BECKER, G. S., K. M. MURPHY, AND R. TAMURA (1990): “Human Capital, Fertility, and Economic Growth,” *Journal of Political Economy*, 98(5), S12–S37.
- BENHABIB, J., R. ROGERSON, AND R. WRIGHT (1991): “Homework in Macroeconomics: Household Production and Aggregate Fluctuations,” *Journal of Political Economy*, 99(6), 1166–87.
- BLAU, F. D., AND L. M. KAHN (2016): “The Gender Wage Gap: Extent, Trends, and Explanations,” NBER Working Papers 21913, National Bureau of Economic Research, Inc.
- BURDA, M. C., D. S. HAMERMESH, AND P. WEIL (2007): “Total Work, Gender and Social Norms,” Discussion Paper 6232, CEPR.
- CHIAPPORI, P.-A. (1988): “Rational Household Labor Supply,” *Econometrica*, 56(1), 63–90.
- DOEPKE, M., M. HAZAN, AND Y. MAOZ (2015): “The Baby Boom and World War II: A Macroeconomic Analysis,” *Review of Economic Studies*, 82(3), 1031–1073.
- EROSA, A., L. FUSTER, AND D. RESTUCCIA (2010): “A Quantitative Theory of the Gender Gap in Wages,” Working Papers 2010-04, Instituto Madrileo de Estudios Avanzados (IMDEA) Ciencias Sociales.
- FERNÁNDEZ, R. (2013): “Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century,” *American Economic Review*, 103(1), 472–500.
- FEYRER, J., B. SACERDOTE, AND A. D. STERN (2008): “Will the Stork Return to Europe and Japan? Understanding Fertility within Developed Nations,” *Journal of Economic Perspectives*, 22(3), 3–22.
- GALOR, O., AND D. N. WEIL (1996): “The Gender Gap, Fertility, and Growth,” *American Economic Review*, 86(3), 374–387.
- (2000): “Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and beyond,” *The American Economic Review*, 90(4), 806–828.
- GOLDIN, C. (1990): *Understanding the Gender Gap: An Economic History of American Women*. Oxford University Press.
- GREENWOOD, J., R. ROGERSON, AND R. WRIGHT (1995): “Household Production in Real Business Cycle Theory,” in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, pp. 157–174. Princeton University Press.

- GREENWOOD, J., AND A. SESHADRI (2005): “Technological Progress and Economic Transformation,” in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf, vol. 1 of *Handbook of Economic Growth*, chap. 19, pp. 1225–1273. Elsevier.
- GREENWOOD, J., A. SESHADRI, AND G. VANDENBROUCKE (2005): “The Baby Boom and Baby Bust,” *American Economic Review*, 95(1), 183–207.
- HAZAN, M., AND H. ZOABI (2015): “Do Highly Educated Women Choose Smaller Families?,” *The Economic Journal*, 125(587), 1191–1226.
- HOFFERTH, S. L., S. M. FLOOD, AND M. SOBEK (2013): “American Time Use Survey Data Extract System,” Discussion paper, Version 2.4 [Machine-readable database]. Maryland Population Research Center, University of Maryland, College Park, Maryland, and Minnesota Population Center, University of Minnesota, Minneapolis, Minnesota.
- JONES, L. E., R. E. MANUELLI, AND E. R. MCGRATTAN (2003): “Why Are Married Women Working So Much?,” Staff Report 317, Federal Reserve Bank of Minneapolis.
- JONES, L. E., AND M. TERTILT (2008): “An Economic History of Fertility in the U.S.: 1826-1960,” in *Frontiers of Family Economics*, ed. by P. Rupert, vol. 1, pp. 165–225. Elsevier.
- KNOWLES, J. A. (2007): “Why Are Married Men Working So Much? The Macroeconomics of Bargaining Between Spouses,” IZA Discussion Papers 2909, Institute for the Study of Labor (IZA).
- (2013): “Why Are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply,” *The Review of Economic Studies*, 80(3), 1055–1085.
- NGAI, L. R., AND C. A. PISSARIDES (2008): “Trends in Hours and Economic Growth,” *Review of Economic Dynamics*, 11(2), 239–256.
- OLIVETTI, C. (2006): “Changes in Women’s Hours of Market Work: The Role of Returns to Experience,” *Review of Economic Dynamics*, 9, 557–587.
- RAMEY, G., AND V. A. RAMEY (2010): “The Rug Rat Race,” *Brookings Papers on Economic Activity*, 41(1), 129–176.
- RAMEY, V. A. (2009): “Time Spent in Home Production in the Twentieth-Century United States: New Estimates from Old Data,” *Journal of Economic History*, 69(1).
- RAMEY, V. A., AND N. FRANCIS (2009): “A Century of Work Leisure,” *American Economic Journal: Macroeconomics*, 1(2), 189–224.

- REGALIA, F., AND J.-V. RÍOS-RULL (1999): “What Accounts for the Increase in Single Households and the Stability in Fertility?,” University of Pennsylvania, mimeo.
- ROBINSON, J., AND G. GODBEY (2008): *Time for Life: The Surprising Ways Americans Use Their Time*. Pennsylvania State University Press.
- SAYER, L. C., S. M. BIANCHI, AND J. P. ROBINSON (2004): “Are Parents Investing Less in Children? Trends in Mothers and Fathers Time with Children,” *American Journal of Sociology*, 110(1), 1–43.
- SIEGEL, C. (2012): “Female Employment and Fertility - The Effects of Rising Female Wages,” Discussion Paper 1156, Centre for Economic Performance.